

Gentran Integration Suite™

Performance and Tuning Guide

Version 4.2

Sterling Commerce
An IBM Company

© Copyright 2007-2008 Sterling Commerce, Inc. All rights reserved.
Additional copyright information is located on the Gentran Integration Suite Documentation Library:
<http://www.sterlingcommerce.com/Documentation/GIS42/homepage.htm>.

Contents

Chapter 1 Monitoring Operations 5

Managing System Logs	5
Monitoring Business Process Threads	15
Monitoring Messages	16
Reviewing System Information	18
Monitoring Node Statuses	26

Chapter 2 Managing Performance 29

Performance Concepts Overview	29
Performance Tuning Methodology	30
Troubleshooting Methodology	32
Managing Performance Statistics	33
Turning On and Off Performance Statistics	43
Creating or Deleting a Performance Statistics Report	43

Chapter 3 Performance Tuning Utility 45

Calculated Recommendations	45
Viewing or Editing Performance Configuration Settings	46

Chapter 4 Manual Performance Tuning 53

Tuning Your ASI Environment	54
Scheduling Policy	58
Queues in an ASI Environment	59
Monitoring Queues using queueWatcher	64
Compressing Cache Contents	66
Setting the Document Body Serialization Threshold Property	67
Allocating JNDI Contexts	68
Changing Persistence Levels	68
Changing System Logging	70
Enabling Trusted Domains for Schemas	71
Enabling the Console Listener	71
Enabling the JMX Event Listener	72
Enhancing EDI Performance	73
Managing Property Files	74
Managing System Recovery	74
Changing the SoftStop Time Limit	76

Changing Advanced File Transfer Settings	76
Working with Large BLOB Data	79
Chapter 5 Resolving Slow System Issues	80
Symptoms, Causes, and Resolutions of a Slow System	80
Conducting a Thread Dump	82
Chapter 6 Resolving Database Issues	89
Full Database Issues	89
Failed Database Connection Requests	92
Tracking JDBC Connections	94
Database Maintenance Check	95
Database Down Check	95
Optimizing the Database	97
Using the Database Statistics Report	98
Slow Performance in Oracle	100
Chapter 7 Resolving Business Process Issues	103
Halted, Halting, Waiting, or Interrupted Business Processes	103
Running a Business Process from the Command Line	105
Chapter 8 Resolving Cache Issues	108
Types of Memory Cache	108
Symptoms and Causes of Inefficient Use of Cache	109
Resolutions for Inefficient Use of Cache	110
Chapter 9 Improving Business Process Execution Times	112
Symptoms and Causes of Poor Business Process Execution Times	112
Resolving Business Process Execution Times	113
Chapter 10 Resolving Memory Issues	115
Correcting Dropped Data Problems	115
Correcting <i>Out of Memory</i> Errors	116
Appendix Requirements Planning Worksheet	117
Index	121

Monitoring Operations

The Operations functions make it possible to monitor the operational status of Gentran Integration Suite, its components and services, to monitor current threads and system logging, and to troubleshoot system problems.

This section covers the following topics:

- ◆ *Managing System Logs* on page 5
- ◆ *Monitoring Business Process Threads* on page 15
- ◆ *Monitoring Messages* on page 16
- ◆ *Reviewing System Information* on page 18
- ◆ *Monitoring Node Statuses* on page 26

Managing System Logs

Gentran Integration Suite has many components, including software applications, Web servers, and database servers. To monitor the activities of each component, Gentran Integration Suite generates log files.

Each operations server on a host has its own operations log file. Log files are created in the *install_dir/logs* directory (or, on Windows, the *install-dir\logs* directory). To prevent the system log files from taking up needed storage space and memory allocations, Gentran Integration Suite generates a log file only when the component runs. This improves the performance of Gentran Integration Suite.

New log files are created once every 24 hours, and normally at midnight. When the time changes to or from Daylight Savings Time, you need to stop and restart Gentran Integration Suite to create daily log files at the same time as before.

If the daily log file size limit is exceeded, another log file is created. If a pre-determined number of logs is exceeded, old logs are automatically deleted. If the interface links to a deleted log file (which displays a blank page), click a link to a newer log file.

Note: If you are working in a clustered environment, the log information that displays is determined by the node you select from the Select Node list.

Log File Naming Conventions

The naming convention for a directory is *install_dir/logs/directory* (or, for Windows, the *install_dir\logs\directory*).

- ◆ *install_dir* or *install-dir* is the name of the installation directory.
- ◆ *logs* is the primary log directory.
- ◆ *directory* is the subdirectory created when you start Gentran Integration Suite. All old log files are moved to the subdirectory for archiving. The naming convention for the old log directory is *logs_mmddyy_hhmmss*.

The naming convention for a log file is *name.log.Dyyyymmdd.Thhmmss_#*.

- ◆ *name* identifies the type of the log file.
- ◆ *log* is the file name extension indicating the type of file.
- ◆ *Dyyyymmdd* is the date in year, month, day format. The D at the beginning of the date means Date.
- ◆ *Thhmmss* is the time in hours, minutes, seconds format. The T at the beginning of the time means Time.
- ◆ *_#* is the increment of the log file. If you attempt to write a log file that already exists, the *_#* is appended to the log file name, allowing you to write the new file and save the integrity of the exiting file.
- ◆ For example, if *mylog.D20041101.T092022* exists and you save a new log file under the same name, the new file becomes *mylog.D20041101.T092022_2*, where the *_2* indicates it is the second in a sequence using the same file name.

Note: The date and time additions in the naming convention may or may not be part of the log file name. For example, *noapp.log* does not include the date and time information, but the *ui.log.Dyyyymmdd.Thhmmss* does include the date and time information.

When you start Gentran Integration Suite, the *install_dir/logs/log* files are archived to a time-stamped subdirectory. When the archive is complete, Gentran Integration Suite writes log files to the *install_dir/logs* directory again.

Viewing Log File Contents

Note: If you are working in a clustered environment, the information that displays is mandated by the node you select from the Select Node list.

To view current log file contents in Gentran Integration Suite:

1. From the **Administration** menu, select **Operations > System > Logs**.
2. Select the appropriate log file.

The log opens.

Note: The Gentran Integration Suite interface displays only the last 2500 lines of a current log file. To view the entire log, you must have Read permission for the file system where Gentran Integration Suite is located. Open the log file (located at the installation path on your hard drive), with a text editor in read-only mode.

To view old log file contents in Gentran Integration Suite:

1. In the `install_dir/logs/log` directory, locate the old log file that you want to view.
2. Open the log file in a text editor in read-only mode.

The log contents display.

Analyzing Log File Contents

The contents of a log file can provide information about system activities and problems. The format for entries written to a log file is `YYYY-MM-DD HH:MM:SS.ss loglevel 'message code' Scope.Subsystem.Name 'information string'`.

- ◆ `YYYY-MM-DD` is the date in year, month, day format.
- ◆ `HH:MM:SS.ss` is the time in hour, minutes, seconds format.
- ◆ `Loglevel` is the condition used to determine the type of information written to the log. The log level conditions are:
 - ◆ `ERROR` – Collects only error conditions.
 - ◆ `INFO` – Collects additional information.
 - ◆ `TIMING` – Collects timing information.
 - ◆ `DEBUG` – Collects detailed information about data flow.
 - ◆ `ALL` – Collects information for all of the above conditions.
- ◆ `message code` describes the activity or problem, using the following format:

The first four digits specify the scope (like Workflow, Ops, Util). The next digit specifies the log severity level (default conventions use 1 for error/exception, 2 for Debug messages, 3 for Warnings and 4 for info/All messages). The next three digits specify the subsystem (like Workflow Queue or Workflow Engine). The last four digits specify the error number.
- ◆ `Scope.Subsystem.Name` is a text description of the affected part of Gentran Integration Suite (like Workflow, Ops, Util); the Gentran Integration Suite subsystem (like Workflow Queue or Workflow Engine); and what occurred (Name)
- ◆ `information string` is a brief description of the activity that occurred.

The following message is an example of this format:

```
[2006-05-30 11:06:55.661] ALL 000440020297
SERVICES.SERVICES_CONTROLLER.INFO_sdi_getName startup: loading HTTP_SEND_ADAPTER
```

At 11:06:55.661 a.m. on May 30, 2006, Gentran Integration Suite was started and attempted to load the HTTP Send adapter (`HTTP_SEND_ADAPTER`). The message also shows the scope (Services); the affected part of Gentran Integration Suite (Services Controller); what occurred (`INFO_sdi_getName`); and the error code (0297).

Changing Log Settings

You can control the logs in Gentran Integration Suite by specifying settings either globally, or for individual log types.

Changing Log Settings Globally

To change property file settings using the customer override property file:

1. In the *install_dir/properties* directory, locate (or create, if necessary) the *customer_overrides.properties* file.
2. Open the *customer_overrides.properties* file in a text editor.
3. Specify settings for the global log properties in the following table using the following format:

```
logService.Property=Value
```

Property	Description
newloggers	Specifies whether to allow new log files to be created when the maximum log file size setting has been exceeded. Valid values: <ul style="list-style-type: none"> ◆ true – Allow new logs to be created (Default) ◆ false – Do not allow new logs to be created
defaultlog	Specifies the name of the default log. Default is systemlogger.
<i>logtype.maxnumlogs</i>	Specifies the maximum number of logs to keep before deleting old logs. Default is 10. Find the <i>logtype</i> for the log file that you want to change in the following table. Examples: <code>defaultlog.maxnumlogs=15</code> – Sets the maximum number of the log type specified as the default log (systemlogger by default) to 15. <code>uilogger.maxnumlogs=20</code> – Sets the maximum number of UI type logs to 20.

The following table provides the *logtype* name for each log file name defined by default in the *log.properties* file. If you have changed the file name, use the original file name from the table to find the *logtype* name.

File Name of Log	<i>logtype</i>
alerterlogger.log	alerterlogger
archive.log	archivelogger
Authentication.log	AuthenticationLogger
cdinterop.log	cdinteroplogger
cdinterop_cdjava.log	cdinteropcjava_logger
ceuinterop.log	ceulogger
common3splogger.log	common3splogger
corbadapter.log	corbadapter
datastore.log	datastore
delete.log	deletelogger

File Name of Log	<i>logtype</i>
ebXML.log	ebXMLlogger
EDIINT.log	EDIINTLogger
event.log	event
ftp.log	ftpllogger
ftpclient.log	psftpclientlogger
http.log	httplogger
httpclient.log	httpclientlogger
jetty.log	jettylogger
lifecycle.log	lifecycleLogger
mailbox.log	mailboxlogger
mgmtdash.log	neo
noapp.log	noapplogger
oftp.log	oftpllogger
ops_exe.log	opslogger
Perimeter.log	PSLogger
pipeline.log	pipelinelogger
report.log	reportlogger
resourcemonitor.log	resourcemonitorlogger
rnif.log	rnlogger
sap.log	saplogger
schedule.log	schedulelogger
Security.log	SecurityLogger
servicesctl.log	sclogger
sftpclient.log	sftpclientlogger
sftpserver.log	sftpserverlogger
si_exe.log	sillogger
sql.log	sqllogger
SyncEngine.log	SyncEngineLogger
system.log	systemlogger
system.log	purgelogger
test.log	testlogger

File Name of Log	<i>logtype</i>
tracking.log	tracking
txtrace.log	txtracelogger
ui.log	uilogger
ui_performance.log	ui_perf_logger
webdav.log	webdavlogger
WebSphereMQSuite.log	wsmqSuiteLogger
webx.log	webxlogger
wf.log	wflogger
wfexception.log	wfexception_logger
wfstatistics.log	wfstatistics

4. Save and close the customer_overrides.properties file.
5. Stop and restart Gentran Integration Suite to use the new values.

Changing Log Settings for an Individual Log Type

Note: If you are working in a clustered environment, the information that displays is mandated by the node you select from the Select Node list.

To change log settings for an individual log type:

1. From the **Administration** menu, select **Operations > System > Logs**.
2. Click the icon next to the type of log for which you want to change the log settings.
3. On the Log Settings page, specify settings for the options in the following table:

Option	Description
Location	Specifies the absolute path for the log file.

Option	Description
Rollover Interval	<p>If newloggers is set to false, this specifies the point at which the oldest lines in the log file are deleted as new lines are created. If newloggers is set to true (default), this is the maximum number of lines allowed in the log file before a new file is created. Select one of the following values:</p> <ul style="list-style-type: none"> ◆ 50000 lines ◆ 100000 lines ◆ 150000 lines ◆ 200000 lines <p>Note: In this case, the term “lines” refers to logical lines (entries) and not to physical lines. For example, the following is a “line”:</p> <pre>[2005-07-11 08:12:07.679] ALL 000440020297 SERVICES.SERVICES_CONTROLLER.INFO_sdi_getName startup: loading HTTP_SEND_ADAPTER FtpConfig.logConfiguration() client configuration: ftpListenPort=[10021] localDataPortCollection=[null] localControlPortCollection=[null] minThreadPoolSize=[3] maxThreadPoolSize=[6] systemCertificateId=[null] caCertificatesIds=[null] passphrase=[*****] cipher=[Strong] sslOption=[SSL_NONE] sslAllowed=[false] sslRequired=[false] sslImplicit=[false] cccAllowed=[false] cccRequired=[false] clusterNodeName=[Node1] perimeterServerName=[local] nonTerminationCharsToTotalCharsRatio=[0.9]</pre>
Logging Level	<p>Specifies the amount of detail to log. Choose one of the following values:</p> <ul style="list-style-type: none"> ◆ On – set the logging level to ALL, which includes debugging (creates larger files) ◆ Off – set the logging level to ERROR (Default), which only logs errors (creates smaller files) <p>Note: Setting the logging level to ALL could generate an excessive amount of debugging information. You should lower the logging level after you have retrieved the debugging information you need.</p>

4. Click **Save**.

Log File Types

Note: If you are working in a clustered environment, the information that displays is mandated by the node you select from the Select Node list.

You can access the following log files through the Gentran Integration Suite interface:

Log Type	Log Name	Description
Central Operations Server		
Operations Security	opsSecurity.log	Used by the security components. Indicates problems with startup, passwords, and passphrases.
Operations Server	ops.log	Used by the operations server.
	opsServer.log	Receives all log messages the operations servers generate at startup.
Gentran:Server for UNIX		
Data Adapter		<p>Logs activity of the Gentran:Server for UNIX adapter. You cannot turn Gentran:Server for UNIX data adapter logging on or off.</p> <p>Note: Gentran:Server for UNIX logs display only if you have Gentran Integration Suite configured for Gentran:Server for UNIX.</p>
Application Logs		
Adapter Server	servicesctl.log	Used by the service controller component.
Alerter	alerterlogger.log	Logs notification failures and Alert own errors in the Alert Service. When debug is turned on, it also logs all Alert information, such as all defined alerters and filters information.
Archive	archive.log	Used by the archive components.
Business Process Exceptions	wfexception.log	Tracks the exceptions that occur while a business process is running.
Business Process Execution	wf.log	Captures information specific to running a business process.
Business Process Policy Statistics	wfstatistics.log	Contains workflow policy statistics generated by the workflow scheduling policy. The actual content depends on the scheduling policy in place, but basically it contains XML time stamps followed by XML records.
Connect:Direct Server Adapter Protocol Layer	cdinterop_cdjava.log	Used by the Connect:Direct Server adapter.
Connect:Direct Server and Requester Adapter and Services	cdinterop.log	Used by the Connect:Direct Server and Requester adapter with related services
Connect:Enterprise Server Adapter and Services	ceuinterop.log	Used by the Connect:Enterprise Server adapter.
Corba Adapter	corbaadapter.log	Used by the CORBA components.

Log Type	Log Name	Description
Dashboard and Community Management	mgmtdash.log	Used by the Dashboard and Community Management components.
Datastore	datastore.log	Contains messages from any of the datastore components. Commonly used to isolate the cause of an abnormal system process termination.
Delete Resources	delete.log	Logs information about resources that have been deleted from Gentran Integration Suite.
Document Tracking	tracking.log	Logs document tracking activities.
ebXML Business Process Execution	ebXML.log	Logs ebXML business process execution activities.
EDIINT AS1 and AS2	EDIINT.log	Used by the EDIINT components.
Event Framework	event.log	Logs event framework activities for events completed in Gentran Integration Suite.
FTP Client Adapter and Services	ftpcient.log	Used by the FTP Client Adapter and related services.
FTP Server	ftp.log	Used by the FTP server components.
G:Server UNIX Lifecycle	lifecycle.log	Used by the Gentran:Server for UNIX Lifecycle components when loading lifecycle records.
Gentran:Server for UNIX Lifecycle Purge Service	system.log	Used by the Gentran:Server for UNIX Lifecycle purge components when purging lifecycle records.
HTTP Client Adapter and Services	httpclient.log	Used by the HTTP Client Adapter and related services.
HTTP Server Adapter	http.log	Used by the HTTP Server Adapter.
Integrator Administration	ui.log	Used by the Gentran Integration Suite interface.
Jetty HTTP Server	jetty.log	Used by the Jetty HTTP Server.
Log.ResourceMonitorLog	resourcemonitor.log	Used by the Resource Monitor.
Mailboxing Subsystem	mailbox.log	Used by the mailbox components in Gentran Integration Suite.
OFTP Administration	oftp.log	Logs OFTP administration activities.
Perimeter Services	Perimeter.log	Used by the perimeter servers components in Gentran Integration Suite.
Pipeline	pipeline.log	Used by the pipeline components.
Report	report.log	Used by the reporting components.
RosettaNet Business Process Execution	rnif.log	Used by the RosettaNet™ components.

Log Type	Log Name	Description
SAP Adapter Administration	sap.log	Used by the SAP® components.
Schedule	schedule.log	Logs scheduling activities.
Security	security.log	Used by the security components. Indicates problems with startup and component licensing.
SFTP Client Adapter and Services	sftpclient.log	Used by the SFTP Client adapter and related services.
SFTP Common Log	common3splogger.log	Logs SFTP security errors.
SFTP Server Adapter	sftpserver.log	Used by the SFTP Server adapter.
SQL Manager	sql.log	Logs queries sent to the database by the SQL Query service.
SyncEngine Components	SyncEngine.log	Logs synchronization engine activities
System	system.log	Used as a general logging service, typically the default system log.
	noapp.log	Used on an application server independent system as a general activity log.
System Output/Error Redirect	noapp.log	Used to provide additional system log information.
Translation Trace Output	txtrace.log	Used as a logging service that helps with map debugging. This log contains debugging messages that show how the translator traversed the maps definition and how the translator matched each block of data against the map.
User Authentication	Authentication.log	Logs user authentication attempts and activities.
Web Extension	webx.log	Used by the Web Extensions components in Gentran Integration Suite.
WebSphereMQ Suite	WebSphereMQSuite.log	Used by the WebSphereMQ Suite Async Receiver adapter and related services.
Windows Service Gentran Integration Suite log	si_exe.log	Log file created by the Gentran Integration Suite Windows service.
Windows Service Ops log	ops_exe.log	Log file created by the Opserver Windows service.

The following table describes log files for Gentran Integration Suite as a Windows service:

Log Name	Description
ScheduleBackup.log	Temporary file that is created when Gentran Integration Suite Windows service stops.

Log Name	Description
Backuplogs.log	Temporary file created when Gentran Integration Suite Windows service stops.
ScheduleStopOps.log	Temporary log file that can be ignored.

The following table describes log files for the DMZ perimeter server:

Log Name	Description
PSLogger.Dyyyyymmdd.Thhmmss	Logs perimeter server information for the DMZ perimeter server.
StartupPS.log	Logs startup activities for the DMZ perimeter server.

Monitoring Business Process Threads

You can monitor threads related to business processes in Gentran Integration Suite using the Thread Monitor. The Thread Monitor lists all threads currently running in the Activity engine. This information may be useful for troubleshooting business processes that appear to have stopped or are hung. You can also stop a thread from the Thread Monitor.

Note: Whenever possible, stop business processes using the options in the Business Process Monitor. These options allow the processes to stop gracefully, maintaining state and status information.

Monitoring Threads


To monitor threads, from the **Administration** menu, select **Operations > Thread Monitor**. Gentran Integration Suite displays the following information for each thread:

- ◆ The state of the thread
- ◆ The ID of the thread (Click the ID number for more information about the thread.)
- ◆ The type of the thread
 - ◆ Business Process – Click the ID number in the ID column to display the Business Process Detail page.
 - ◆ RMI (Remote Management Interface) – Click the ID number in the ID column to display the Business Process Detail page.
 - ◆ Harness – Click the ID number in the ID column to display the Business Process Detail page.
 - ◆ Schedule – Click the ID number in the ID column to display the Schedule settings.
- ◆ The processing priority assigned to the thread
- ◆ The date and time the thread was registered

Note: You can also view thread details from the System Troubleshooter.

Stopping Threads

Caution: Before using the Thread Monitor to stop a thread, you should contact Sterling Commerce Customer Support for assistance. Unexpected results can occur that may affect other processing. Stop or terminate business processes from the Business Process Monitor whenever possible.

To stop a thread, click the stop icon  for the thread. This instructs the thread to immediately stop processing without waiting for any steps to complete or exit with an error. In most cases, the thread will stop immediately. This will not immediately change the display in the Thread Monitor. The thread will eventually be marked as interrupted.

Note: In some cases, it is not possible to stop a thread and your attempt will be unsuccessful. Stopping a thread does not stop the service that is performing that thread. The service is still available for Gentran Integration Suite to use in other business processes.

Setting the Refresh Option for the Thread Monitor

By default, the Thread Monitor is set to automatically refresh every 15 seconds. To turn this option off, clear the check box next to *Automatically refresh every 15 seconds*.

Monitoring Messages

In Gentran Integration Suite, business processes use messages to communicate with each other. Certain business processes create Produce and Consume messages. The Produce business processes produce messages to be consumed. The Consume business processes wait for messages from the Produce business processes.

Sometimes the hand off between the Produce and Consume business processes does not occur—the produced message has no waiting Consume process, or the waiting Consume process waits without success for the produced message. You use the Message Monitor to track unsuccessful message hand offs within Gentran Integration Suite.

Here are some special considerations about messages monitored within Gentran Integration Suite:

- ◆ The Message Monitor pages of the Gentran Integration Suite interface show only messages associated with a failed Produce or Consume activity. Messages that are successfully produced or consumed are not monitored.
- ◆ Messages are tied closely to business processes. You can search for messages by business process name, type of business process (Produce or Consume), and date and time that the associated business process step invokes the Produce or Consume activity. But messages also have identities separate from business processes. Messages have names, and their names can be used as search input.
- ◆ Sometimes Produce and Consume business processes have to choose between messages with the same name. For example, two Produce processes run at different times; both produce a message named OutMsg1. So there are two messages in Gentran Integration Suite named OutMsg1, but with different start date/times. When a Consume process detects the two messages, it consumes the OutMsg1 with the later start date/time. The situation also works in reverse. Two Consume processes request an

identically named message (InMsg1). When a Produce process produces InMsg1, the Produce process produces InMsg1 for the Consume process with the earlier start date/time.

- ◆ An unsuccessful message remains in Gentran Integration Suite until an appropriate Produce or Consume business process step completes the message hand off and uses the message. Unsuccessful messages never expire or time out.

Searching for Messages

To search for unsuccessful messages in Gentran Integration Suite:

1. From the **Administration** menu, select **Operations > Message Monitor**. The Message Monitor page opens.
2. Select one of the following search methods and click **Go!**
 - ◆ Message Name – Display messages whose names contain the specified character or string.
 - ◆ Start Date: and End Date: – Display messages for which the related business process step invoked the Produce or Consume activity at the specified start date and time. Which start dates and times are displayed by default depends on whether Gentran Integration Suite has outstanding unsuccessful messages.
 - If Gentran Integration Suite has no unsuccessful messages, the **Start Date:** field displays the current system date and time minus one minute, and the **End Date:** field displays the current system date and time plus one minute.
 - If there are unsuccessful messages in Gentran Integration Suite, the **Start Date:** field displays the date and time of the unsuccessful message with the earliest start date and time, minus one minute, and the **End Date:** field displays the date and time of the unsuccessful message with the latest start date and time, plus one minute.
 - ◆ Alphabetically or by message type (ALL, Producer, or Consumer) – Display messages initiated by business processes that:
 - Have names starting with a particular letter or number selected from the Alphabetically list. Default is ALL (all business processes).
 - Are either Produce or Consume processes, or both (ALL).
3. Review the information for the listed messages. This information is displayed in the following columns:
 - ◆ Type – Type of business process associated with the message, either Produce or Consume.
 - ◆ ID – ID of the business process associated with the message. Click an ID in this column to display the Business Process Details page for the business process.
 - ◆ Process Name – Name of the associated Produce or Consume business process. Click a process name in this column to display the BPML of the business process.
 - ◆ Message Name – Name of the unconsumed message.
 - ◆ Start Time – Start date and time of the associated step within the business process.

Reviewing System Information

Use the System Troubleshooter page to review system information and to start troubleshooting system issues in Gentran Integration Suite.

From the System Troubleshooting page, you can:

- ◆ View system information on different nodes in a clustered environment

Note: Viewing system information on different nodes from any node in the clustered environment is only available if you are set up in a clustered environment.

- ◆ Refresh system statuses
- ◆ Stop Gentran Integration Suite
- ◆ View database usage statistics
- ◆ View business process queue and usage statistics
- ◆ Terminate a business process
- ◆ View system classpath information
- ◆ View system JNDI tree information
- ◆ View environment statistics, including cache and memory used
- ◆ View perimeter servers information

Note: If you are running DB2 as your database, the database usage statistics always shows as unavailable in the System Troubleshooting page.

The System Troubleshooting page is separated into different areas. The following table describes each area:

Area	Information
Select Node	<p>The select node list enables you to select a node in a clustered environment triggering which node's information to display in the System Troubleshooting page.</p> <p>For example, if you have two nodes in a cluster, Node 1 and Node 2 and you want to view the System Troubleshooting page for Node 2, select Node 2 from the list and the System Troubleshooting page for Node 2 displays. If you want to view Node 1 information, select Node 1 from the list and the System Troubleshooting information for Node 1 displays.</p> <p>Note: The Select Node list displays only if you are working in a clustered environment. Your selection mandates which node's information displays in the remainder of the System Troubleshooting page.</p>

Area	Information
System Status	<p>The System Status area displays the following information and options:</p> <ul style="list-style-type: none"> ◆ Start time ◆ Uptime ◆ Refresh Status ◆ Stop the System ◆ Database usage ◆ Business process queue usage ◆ Business process usage ◆ Classpath information ◆ Java™ Naming and Directory Interface (JNDI) tree structure ◆ Length of time since completion of the following system cleanup processes: <p>Note: To refresh the system status, click Refresh Status.</p>
Gentran:Server	<p>The Gentran:Server area displays the following information:</p> <ul style="list-style-type: none"> ◆ Host ◆ Location ◆ State ◆ Controllers ◆ Adapters <p>Note: This area displays only if you configured your system for Gentran:Server Data Manager support.</p>
Application Status	<p>The Environment area displays the following information:</p> <ul style="list-style-type: none"> ◆ Cluster Node <p>Note: The cluster node name displays only if you are working in a clustered environment.</p> <ul style="list-style-type: none"> ◆ Host ◆ Location ◆ State ◆ Cache usage ◆ Memory in use ◆ List of active controllers ◆ List of all adapters in the system and their status <p>Note: Click an adapter name to view more about its status on the Service Settings page.</p> <ul style="list-style-type: none"> ◆ List of active threads

Area	Information
Database Services	<p>The <i>Database Services</i> area displays the following information:</p> <ul style="list-style-type: none"> ◆ Host ◆ Location ◆ State <p>Note: The name of this area changes with the database that you are running. For example, MySQL Services for a MySQL database.</p>
RMIAPrime Services	<p>The RMIAPrime Services area displays the following information:</p> <ul style="list-style-type: none"> ◆ Host ◆ Location ◆ State
Perimeter Servers	<p>The Perimeter Servers area displays the following information:</p> <ul style="list-style-type: none"> ◆ Cluster Node name (in a clustered environment only) ◆ Whether the perimeter server is on or off ◆ State, either enabled or disabled ◆ Name of the perimeter server ◆ Last Activity <p>Note: The information in this area displays only after you have added a perimeter server to Gentran Integration Suite.</p>

About the System Status Area

The System Status section of the System Troubleshooting displays the following information for the integration server:

Note: If you are working in a clustered environment, the information that displays is mandated by the node you select from the Select Node list.

- ◆ Current date and time
- ◆ Start Time – The date and time the system was last started or restarted
- ◆ Uptime – The amount of time the system has been running since the last start or restart
- ◆ Refresh Status – Refreshes active controllers
- ◆ Stop the System – Stops Gentran Integration Suite using the softstop script
- ◆ DB Usage – Displays database usage statistics
- ◆ Business Process Queue Usage – Displays business process queue usage statistics
- ◆ Business Process Usage – Displays business process usage statistics
- ◆ Force Terminate a Business Process – Enables you to terminate a business process by force.
- ◆ Classpath – Displays the Gentran Integration Suite classpath
- ◆ JNDI Tree – Displays the JNDI tree in Gentran Integration Suite

Refreshing System Status

System status information is collected and reported in real time. To view the most current system status, you must refresh the system status.

To refresh the system status:

1. From the **Administration** menu, select **Operations > System > Troubleshooter**.
2. On the System Troubleshooting page, in the System Status area, click **Refresh Status**.

The system status information refreshes showing the most current information.

Stopping Gentran Integration Suite From the System Troubleshooting Page

You can stop Gentran Integration Suite using the System Troubleshooting page. Stopping Gentran Integration Suite in this manner stops Gentran Integration Suite using the softstop script allowing all business processes to complete before stopping the system.

Caution: Using the Stop the System option stops only the Gentran Integration Suite interface immediately, while all business processes in progress run until complete. After all business processes' current services complete, Gentran Integration Suite stops. To stop the system and all processing immediately, in the *install_dir/bin* directory, run the hardstop script. All processes that have not completed stop and may need to be restarted.

To Stop Gentran Integration Suite using the System Troubleshooting page:

1. From the **Administration** menu, select **Operations > System > Troubleshooter**.
2. On the System Troubleshooting page, in the System Status area, click **Stop the System**.
3. In the message asking if you want to stop Gentran Integration Suite, click **OK**.

The interface stops immediately, but all business processes that are in progress complete before the system stops.

Viewing Database Usage Statistics

Note: Database usage statistics are not available for DB2 and DB2 Z-Series.

Gentran Integration Suite uses pools to store database connections. Database usage statistics show how your database is performing. To change pool settings, you must manually edit the configuration files and restart the system.

To view database usage statistics:

1. From the **Administration** menu, select **Operations > System > Troubleshooter**.
2. In the System Status area, click **Database Usage**.

Note: If you have the `displayGraphics` property located in the *install_dir/properties/ui.properties* file set to `true`, the Database Usage page displays in graphic format; otherwise, the Database Usage page displays in text format. The default is `true` for Linux, Sun, HP, and Windows operating systems. The default is `false` for AIX, OS400, and Z Series operating systems.

The Database Usage page displays, showing the following information from a unit test of the database:

- ◆ Average time it takes to perform the number of database inserts in the unit test

- ◆ Number of inserts performed to the database in the unit test

Note: You can change the value in the `dbAccessLoopCnt` property in the `install_dir/properties/ui.properties.in` file. After you make your changes, in the `install_dir/bin` directory, run the `setupfiles` script.

- ◆ Size of inserts performed to the database in the unit test

Note: You can change the value in the `dbAccessDataSize` property in the `install_dir/properties/ui.properties.in` file. After you make your changes, in the `install_dir/bin` directory, run the `setupfiles` script.

- ◆ Size of the database and the amount of the database used in megabytes

- ◆ Green – Normal range
- ◆ Yellow – Warning range
- ◆ Red – Critical range

- ◆ Number of business processes that are waiting to be archived, indexed, or purged

- ◆ Size in megabytes of the following pools and the number of requests that had to wait for the following pools:

- ◆ `gentranTPPool`
- ◆ `databaseArchivePool`
- ◆ `databasePool`
- ◆ `databasePool_local`
- ◆ `databasePool_NoTrans`
- ◆ `databasePool_Select`
- ◆ `databaseUIPool`

Note: The database portion of the pool names changes depending on the database you are using. For example, if the database is MySQL, you see `mysqlPool`.

Viewing the System Classpath

You can view the system classpath for debugging purposes and to verify whether third-party libraries are available in the classpath.

To view the system classpath:

1. From the **Administration** menu, select **Operations > System > Troubleshooter**.
2. On the System Troubleshooting page, in the System Status area, click **Classpath**.

The system classpath displays.

Viewing the System JNDI Tree

You can view the system JNDI Tree for debugging purposes and to verify whether the expected resources, for example, adapters or pool names, are in the JNDI tree.

To view the system JNDI Tree:

1. From the **Administration** menu, select **Operations > System > Troubleshooter**.
2. On the System Troubleshooting page, in the System Status area, click **JNDI Tree**.

The system JNDI Tree displays showing the JNDI name and class name pairs.

Viewing a Node JNDI Tree in a Clustered Environment

You can view a specific nodes JNDI Tree for debugging purposes and to verify whether the expected resources, for example, adapters or pool names, are in the JNDI tree. This option is available only in a clustered environment.

To view a node JNDI Tree in a clustered environment:

1. From the **Administration** menu, select **Operations > System > Troubleshooter**.
2. In the System Status area, click **node#**, where # is the number of the node you want to view information about.

The node JNDI Tree displays showing the JNDI name, class name pairs, and the node name.

About the Application Status Area

Note: If you are working in a clustered environment, the information that displays is mandated by the node you select from the Select Node list.

The Application Status area of the System Troubleshooting page displays the current operational status of the processing environment for a copy of Gentran Integration Suite. The System Troubleshooting page displays a separate section for each environment. Each section provides the following information:

- ◆ The cluster node name if you are working in a cluster environment

Note: The cluster node list displays only if you are working in a clustered environment. After you set up your cluster, the select cluster node list displays.

- ◆ Host - The name of the host on which a specific environment resides
- ◆ Location - The location or path of the installation
- ◆ State - The state of the environment, either Active or Inactive
- ◆ Memory in use - The amount of memory used by Gentran Integration Suite
- ◆ Active threads - The number of concurrent threads that are active
- ◆ Cache Usage - The cache information used by Gentran Integration Suite
- ◆ Threads - The number of threads
- ◆ Controllers - The state and name of each controller or server in the Gentran Integration Suite environment
- ◆ Adapters - The state, name, and type of each adapter installed in the Gentran Integration Suite environment

Viewing Cache Usage Information

The cache usage report displays these statistics for each cache: count, number of requests, and number of successful hits. Gentran Integration Suite uses caches to hold information that is frequently requested by the

system. To change cache settings, see *Performance Tuning Utility* on page 45. You can view the cache usage information to monitor the use of various cache types.

To view the cache usage:

1. From the **Administration** menu, select **Operations > System > Troubleshooter**.
2. In the Application Status area, click **Cache Usage**.

The Cache Usage report displays the following information for each cache type:

- ◆ Cache name – Name of the cache
- ◆ Count – Number of objects in the cache
- ◆ Requests – Number of times an object was requested from the cache, regardless of success
- ◆ Hits – Number of times an object was requested from the cache and found.

Viewing Threads

To view threads in Gentran Integration Suite:

1. From the **Administration** menu, select **Operations > System > Troubleshooter**.
2. In the Application Status area, click **Threads**.

The Threads report displays.

Refreshing a Controller

You can refresh controllers running in your environment using the System Troubleshooting page.

To refresh a controller:

1. From the **Administration** menu, select **Operations > System > Troubleshooter**.
2. In the Application Status area, next to the controller you want to refresh, click the refresh icon in the Refresh column.

The controller is refreshed.

About the Database Services Area

Note: In a cluster, the database is shared by all nodes in the cluster. The database information displayed in the *Database Services* area is the same across all nodes in the cluster.

The *Database Services* area of the System Troubleshooting page displays the following information about your database:

- ◆ Name of the host on which a database resides
- ◆ Location or path of the database installation
- ◆ State of the database, either Active or Inactive

Note: The name for this area displays showing the name of the database you are running. For example, MySQL Services for a MySQL installation.

About the RMIAPrime Service Area

Note: If you are working in a clustered environment, the information that displays is mandated by the node you select from the Select Node list.

The RMIAPrime Service area of the System Troubleshooting page displays the following information about the RMIAPrime Service:

- ◆ Name of the host
- ◆ Location or path to the RMIAPrime Service
- ◆ State of the RMIAPrime Service, either Active or Inactive

About the Perimeter Servers Area

Note: If you are working in a clustered environment, the information that displays is mandated by the node you select from the Select Node list.

The Perimeter Servers area of the System Troubleshooting page displays the following information about your perimeter servers:

- ◆ Name of the cluster node that the perimeter server is associated
- ◆ State of the perimeter server – Enabled or Disabled
- ◆ Name of the perimeter server
- ◆ Date and time of the last activity the perimeter server performed

Enabling a Perimeter Server

You can enable a disabled perimeter server using the System Troubleshooting page.

To enable a perimeter server:

1. From the **Administration** menu, select **Operations > System > Troubleshooter**.
2. In the Perimeter Servers area, next to the perimeter server that you want to enable, in the On/Off column, select the check box.

Disabling a Perimeter Server

You can disable an enabled perimeter server using the System Troubleshooting page.

To disable a perimeter server:

1. From the **Administration** menu, select **Operations > System > Troubleshooter**.
2. In the Perimeter Servers area, next to the perimeter server that you want to disable, in the On/Off column, clear the check box.

Viewing Perimeter Server Settings

You can view the perimeter server settings to verify accuracy or to plan changes as needed.

To view perimeter server settings:

1. From the **Administration** menu, select **Operations > System > Troubleshooter**.

2. In the Perimeter Servers area, in the Name column, click the name of the perimeter server that you want to view.

Monitoring Node Statuses

If you are working in a clustered environment, information is available about all the nodes in the cluster.

To view node status information:

1. From the **Administration** menu, select **Operations > System > Cluster > Node Status**. The Node Status page displays the following information:

Heading	Description
Name	Shows the name of the node. Click the name to view more details about the node.
URL	Shows the uniform resource locator for the node.
Troubleshooter	Provides a link to the System Troubleshooter.
Token Node	States whether this is the token node. <ul style="list-style-type: none"> ◆ True – This is the token node. ◆ False – This is not the token node.
Creation Time	Shows the date and time that the node was created.
Status	Shows the status of the node. <ul style="list-style-type: none"> ◆ Active – Node is working and available for processing ◆ Node went down – Node is not working and not available for processing ◆ Node completely down – All services of a node are completely stopped ◆ Starting Ops – Node is starting up, but not available for processing <p>Note: If a node went down, in the <i>install_dir/bin</i> directory of the node, run the <i>hardstop.sh</i> (UNIX) or <i>hardstop.cmd</i> (Windows) script and then run the <i>run.sh</i> (UNIX) or <i>startWindowsService.cmd</i> (Windows) script. This stops all processing and restarts the node.</p> <p>Note: Restart the script of the node when its status is <i>Node completely down</i>. You can monitor the status of the restarted node in other nodes of the cluster.</p>

2. To view more details about a node, click the **Name** of the node on the Node Status page. A page will display that contains the following information about that node:

Heading	Description
Node Name	Shows the name of the node.
URL	Shows the uniform resource locator for the node.

Heading	Description
Token	States whether this is the token node. <ul style="list-style-type: none"> ◆ True – This is the token node. ◆ False – This is not the token node.
Creation Time	Shows the date and time that the node was created.
Status	Shows the status of the node. <ul style="list-style-type: none"> ◆ Active – Node is working and available for processing ◆ Node went down – Node is not working and not available for processing ◆ Node completely down – All services of a node are completely stopped ◆ Starting Ops– Node is starting up, but not available for processing <p>Note: If a node went down, in the <i>install_dir/bin</i> directory of the node, run the <i>hardstop</i> script and then run the <i>run</i> script. This stops all processing and restarts the node.</p> <p>Note: Restart the script of the node when its status is <i>Node completely down</i>. You can monitor the status of the restarted node in other nodes of the cluster.</p>
Location	Shows the directory path where the node is installed.
Role	States the role of the node.
Operation Controller Host	Shows the name of the server acting as the Operation Controller Host.
Operation Controller Port	Shows the port number for the Operation Controller.
Gentran Integration Suite Version	Shows which version of Gentran Integration Suite is used for the node.
JVM Version	Shows the version number for the java virtual machine.
JVM Vendor	Shows the vendor that provided the java virtual machine.
	Note: The following information only displays if <i>queueWatcher=true</i> in the <i>install_dir/properties/noapp.properties</i> file.
Multicast Info	Provides information on the communication mechanism across nodes. This is used to provide load factor and status information about a node to all other nodes in a cluster. Clustering uses the JGroups communication toolkit, which supports the following communication methods. The communication protocols below can be configured at deployment time: <ul style="list-style-type: none"> ◆ UDP (User Datagram Protocol) (IP multicast) ◆ TCP
Engine Queue Content (Use the dropdown menu to select the Queue you want to review)	Provides a link to the engine queue content page for the node and queue. The engine queue content page contains details about active threads and items in the queue that are waiting for a thread.

Heading	Description
Distribution Size	Provides a link to the distribution size page. The distribution size page contains the following information: <ul style="list-style-type: none"><li data-bbox="493 333 792 359">◆ Scheduling policy details<li data-bbox="493 380 834 405">◆ Soft Reference cache details<li data-bbox="493 426 857 451">◆ Queue details (for each queue)

For more information about clustering, contact Sterling Commerce Customer Support.

Managing Performance

Gentran Integration Suite provides many ways for you to enhance system performance. You can change the database configuration, allocate more or less memory to different components of the system, or change cache and pool sizes to meet your needs.

This section covers the following topics:

- ◆ *Performance Concepts Overview* on page 29
- ◆ *Performance Tuning Methodology* on page 30
- ◆ *Troubleshooting Methodology* on page 32
- ◆ *Managing Performance Statistics* on page 33
- ◆ *Turning On and Off Performance Statistics* on page 43
- ◆ *Creating or Deleting a Performance Statistics Report* on page 43

Performance Concepts Overview

Performance means different things to different users. Performance can be related to:

- ◆ Latency – Amount of time required to generate a response to a request (speed)
- ◆ Throughput – Amount of data transferred in a specified time period (volume)
- ◆ Scalability – Ability of the system to grow with increased workloads (additional hardware)
- ◆ Load – Ability of the system to continue performance levels with increased workloads

Gentran Integration Suite can be tuned to perform as you want it to. If you prefer more speed in your processing, you can change various parameters to meet your needs. If you need higher volume processing, you can set parameters to accomplish your goals.

Performance issues can typically be addressed with the following actions:

- ◆ Changing performance parameters in the properties files or through the performance tuning utility. For more information, see *Performance Tuning Utility* on page 45.
- ◆ Adding additional hardware.
- ◆ Tuning your business processes to run more efficiently within Gentran Integration Suite.

- ◆ Monitoring databases and archiving databases to free up resources. For more information, see *Viewing Database Usage Statistics* on page 21.
- ◆ Creating a Gentran Integration Suite cluster for load balancing and scaling.

Before You Begin Performance Tuning

Before you can take any performance tuning actions, you must consider capacity planning issues. You may have worked with Sterling Commerce to determine your capacity needs and to install Gentran Integration Suite. If not, the *Requirements Planning Worksheet* on page 117 can help you determine your capacity needs.

The following is a list of some capacity issues that impact performance and tuning:

- ◆ Daily volume requirements, including the average size and number of transactions to be processed.
- ◆ Additional processing requirements, for example translation and straight through processing.
- ◆ Types of pay loads including EDIFACT, XML, and other formats.
- ◆ Translation requirements. For example, translation from EDIFACT to XML.
- ◆ Enterprise Resource Planning (ERP) integration requirements. For example, integration with SAP or PeopleSoft.
- ◆ Number of processors dedicated to Gentran Integration Suite.
- ◆ Memory requirements to meet your processing requirements.
- ◆ Disk space requirements.
- ◆ Hard disk *Redundant Array of Independent Disks (RAID)* level, which are disk drives using two or more drives in combination for fault tolerance and performance. Recommended RAID level for Gentran Integration Suite is Level 5.
- ◆ Database size requirements.

Note: When conducting any performance tuning activity, keep this information nearby for easy reference and future planning needs.

Performance Tuning Methodology

For effective performance tuning, you must be able to identify the performance issue and determine the corrective action. Following a methodology helps focus your efforts and reduces redundant operations.

Use the following methodology to identify and resolve performance issues in Gentran Integration Suite:

1. Fill out the *Requirements Planning Worksheet* to determine the hardware requirements for the processing volume and speeds you want. Use the worksheet to help determine the hardware requirements.

Note: If you purchased Sterling Commerce Consulting Services, you received a report recommending requirements according to your business needs.

2. Verify that your hardware and memory amounts match the hardware and memory recommendations according to your processing time and volume requirements completed in the performance planning phase.
3. Verify that you have patched your system using any patches that fix your performance issue from the Sterling Commerce SOD site.
4. Verify that you have the supported Java Virtual Machine (JVM) on the computer running Gentran Integration Suite and on the DMZ computer if you are running perimeter servers. Both JVM versions must match each other and the requirements for your version of Gentran Integration Suite. The build date and lower release numbers must also match.
5. Verify that you are running the supported version of your operating system.
6. Verify that you are running the supported versions of the Java Database Connectivity (JDBC) drivers,
7. Verify that you have created your business processes using the most current adapters and services, using well structured XPath statements, and with the lowest persistence by step and business process that meets your business needs.
8. Tune Gentran Integration Suite using the Performance Tuning Utility and the information that you provided in the Requirements Planning worksheet (see *Requirements Planning Worksheet* on page 117). The Performance Tuning Utility enables you to tune the cache, memory, and other system components, but not business processes.

For more information about the Performance Tuning Utility, see *Performance Tuning Utility* on page 45.

9. Enable Performance Statistics on the **Operations > System > Performance > Statistics** page. The Performance Statistics reports provide information about business process and activities execution times, database connections and persistence, and business process queue performance. You can use this information to tune your system to meet your volume and speed requirements.

For more information about the Performance Statistics Report, see *Managing Performance Statistics* on page 33.

10. Review your history of incoming and outgoing documents, and place the information on a chart showing daily, weekly, and monthly processing trends. Use these charts to determine your peak volume processing levels according to your historical processing methods.
11. Conduct initial performance testing by running your business processes with sample data that is of equal size of the data that will be processed in production. In addition, run the approximate volume that is your peak anticipated volume processing. This helps you tune your system as closely to your production environment as possible.
12. Review the Performance Statistics Report for processing speeds, volumes, and database connections.
13. Review the other reports, such as the Database Usage report and the Cache Usage report, for areas that can be tuned.
14. Retune Gentran Integration Suite using the Performance Tuning Utility, according to the information you gathered from your initial performance testing.
15. Continue this process until you have met your processing time and volume requirements according to your peak volume processing.
16. Create a new Performance Statistics Report called *Benchmarksdd/mm/yy*.

17. Conduct the same tests that you conducted in step 11.
18. Review the Benchmarks $dd/mm/yy$ Performance Statistics Report. If the statistics do not compare closely to your previous statistics complete steps 11 through 14 again.
19. Compare your monthly or weekly Performance Statistics Reports to this benchmark report to verify that your system is processing business processes efficiently and that your resources are being used efficiently. Using this methodology as a proactive measure may reduce downtime and processing issues before they arise.

Troubleshooting Methodology

After you have planned for and implemented Gentran Integration Suite, you must monitor your system to keep it in running at peak performance and to reduce the risks of delayed processing or loss of data. If you monitor your system reviewing the various log files reports you proactively resolve problems before they cause processing delays or loss of data.

Use the following general troubleshooting methodology to monitor Gentran Integration Suite performance.

1. Verify the hardware and memory used matches the hardware and memory recommendations according to your processing time and volume requirements completed in the performance planning phase.
2. Verify that you have applied any patch that fixes your problem using the patches from the Sterling Commerce SOD site.
3. Verify that you have the supported JVM on the computer running Gentran Integration Suite and on the DMZ computer if you are running perimeter servers. Both JVM versions must match each other and the requirements for your version of Gentran Integration Suite. The build date and lower release numbers must also match.
4. Verify that you are running the supported version of your operating system.
5. Verify that you are running the supported versions of the JDBC drivers.
6. If other applications are running on the same computer as Gentran Integration Suite, are all of the applications experiencing problems, such as running slow?
 - ◆ If so, the problem is most likely the operating system or a hardware problem and not specific to Gentran Integration Suite.
 - ◆ If the only application affected is Gentran Integration Suite, the issue may be Gentran Integration Suite specific. For more information, see *Resolving Slow System Issues* on page 80.
7. Is the database used by Gentran Integration Suite used by other applications? If so, verify whether the other applications are getting quick responses to database queries.
 - ◆ If the only application affected is Gentran Integration Suite, the issue may be Gentran Integration Suite specific. For more information, see *Resolving Database Issues* on page 89.
 - ◆ If the other applications are impacted, the issue may be with your database. For more information, see your database documentation.
8. Are your business processes well designed and do they use well structured XPath statements? For example, verify that your business processes do not use redundant services and adapters, or have

persistence levels set too high. Write XPath statements using relative paths. For example, write `PurchaseOrder/text` instead of `/ProcessData/PurchaseOrder/text()`. In addition, do not use `//` at the beginning of an XPath statement; this causes the entire process data to be searched and slows down the query.

- ◆ If your business processes are well designed and the XPath statements are well structured, the issue may be a code issue. Contact Sterling Commerce Customer Support.
 - ◆ If the business process requires improvements, make the improvements and run the business process again. For more information about improving business processes, see the *Managing Business Processes* topic.
9. If you cannot pinpoint the cause of your performance issue, call Sterling Commerce Customer Support for assistance.

Managing Performance Statistics

The Performance Statistics Report is the key to managing your Gentran Integration Suite performance.

Before you can create or view a performance statistics report, you must turn on performance statistics. See the *Turning On and Off Performance Statistics* on page 43. You can turn on and off statistics to the current report as many times as you need to. When you have performance statistics turned off, the statistics are simply discarded, and when you have performance statistics turned on, the current report resumes collecting statistics. If you do not create a report before you turn on performance statistics, the default report is created automatically, and collects the performance statistics.

Run the Performance Statistics Report to obtain benchmarks directly after you:

1. Install Gentran Integration Suite.
2. Tune Gentran Integration Suite using the Performance Tuning Utility.
3. Complete any manual tuning to the `install_dir/properties/noapp.properties` file.
4. Make improvements to your business processes, including changing the persistence levels.
5. After you have your initial benchmark Performance Statistics Report, create a new report every few weeks or at least one time a month to take a snapshot of how your Gentran Integration Suite is performing in relation to your benchmarks.

If you note variances between your benchmarks and your current report, you can use the information in the report to help troubleshoot and resolve the issue before it causes performance delays.

The most current report is always at the top of the view report list. After you close a report, no further data can be added to the report. Every time you create a new report, the new report becomes the current report.

Viewing the Performance Statistics Report

To view a Performance Statistics Report:

1. From the **Administration** menu, select **Operations > System > Performance > Statistics**.

- On the Performance Statistics Report Manager page, under **View Report**, select a name from the list, and click **Go!**

Reading the Performance Statistics Report

The Performance Statistics Report provides you with information that can help you maintain your Gentran Integration Suite system and help reduce performance issues before they become problems. The Performance Statistics Report can help you identify and analyze bottlenecks in your business processes and the time period that each activity takes to complete.

The report displays business process and system statistics since the time that the report began. If a new report is created, it does not contain historical statistics, only the statistics collected after the report was created.

The following figure shows a sample of the Business Process Statistics portion of the Performance Statistics Report:

Performance Statistics Report

Initial

Business Process Statistics							
Name	Min (ms)	Max (ms)	Average (ms)	Invocations	Total (ms)	Pre-Service (ms)	Post-Service (ms)
CllDevelope	18183	19275	19229.0	2	38458	Invocations per hour: 371.393	
→ 1 DevelopeCll	18474	18672	18573.0	2	37146	152.0	176.0
→ 2 DecisionEngineService	0	0	0.0	2	0	5.0	100.0
EDIDevelope	15142	15142	15142.0	1	15142	Invocations per hour: 237.749	
→ 1 EDIDevelope	0	14059	14059.0	1	14059	0.0	53.0
Schedule_IndexBusinessProcessService	733	1732	1144.0	7	8008	Invocations per hour: 6.979	
→ 1 AssignService	0	1	0.0	7	3	1.0	7.0
→ 2 SystemLockService	15	82	33.0	7	235	0.0	38.0
→ 3 AssignService	4	24	10.0	7	72	1.0	5.0
→ 4 IndexBusinessProcessService	549	1534	809.0	7	5067	0.0	14.0
→ 5 SystemLockService	13	95	22.0	7	167	0.0	12.0
Schedule_PurgeService	447	1226	809.0	7	5067	Invocations per hour: 6.980	
→ 1 AssignService	1	17	3.0	7	22	1.0	6.0
→ 2 SystemLockService	17	128	63.0	7	373	1.0	12.0
→ 3 AssignService	3	29	11.0	7	81	1.0	6.0

The following table describes the columns:

Note: The business process statistics are more varied than what the examples in the Description column might imply.

Statistic Name	Description
Business Process Statistics	
Name	<p>Name of the business process followed by numbered rows containing the individual activities in the business process. The number of the row indicates the step in the business process.</p> <p>For example, in the preceding figure, the business process CIIDeveloper includes two activities:</p> <ul style="list-style-type: none"> ◆ --> 1 DeenveloperCII ◆ --> 2 DecisionEngineService.
Min (ms)	<p>Minimum time in milliseconds (ms) the business process or activity took to complete. If there are multiple invocations, this is the lowest minimum time for all of the invocations.</p> <p>For example, in the preceding figure, the CIIDevelope business process took 19183 ms to complete, while the Deenveloping step (1 DeenveloperCII) took a minimum time of 18474 ms. Both of these are the lowest statistics; there were two invocations of the business process.</p> <p>Use this statistic to track business processes or activities that begin to slow in processing times. If you notice this statistic increasing over your benchmark value, this may indicate a performance issue. If you see an occasional variance for a business or activity, this does not indicate a performance issue. If, however, you notice a continual variance between production statistics and your benchmarks, you must resolve the issue.</p> <p>For more information on slow systems, see <i>Resolving Slow System Issues</i> on page 80.</p> <p>For more information about improving business process execution times, see <i>Improving Business Process Execution Times</i> on page 112.</p>
Max (ms)	<p>Maximum time in milliseconds (ms) the business process or activity took to complete. If there are multiple invocations, this is the highest maximum time for all of the invocations.</p> <p>For example, in the preceding figure, the CIIDevelope business process took 19275 ms to complete, while the Deenveloping step (1 DeenveloperCII) took a maximum time of 18672 ms. Both of these are the highest statistics; there were two invocations of the business process.</p> <p>Use this statistic to track business processes or activities that begin to slow in processing times. If you notice this statistic increasing over your benchmark value, this may indicate a performance issue. If you see an occasional variance for a business or activity, this does not indicate a performance issue. If, however, you notice a continual variance between production statistics and your benchmarks, you must resolve the issue.</p> <p>For more information on slow systems, see <i>Resolving Slow System Issues</i> on page 80.</p> <p>For more information about improving business process execution times, see <i>Improving Business Process Execution Times</i> on page 112.</p>

Statistic Name	Description
Average (ms)	<p>Average processing time for the business process or service to complete.</p> <p>For example, in the preceding figure, the CIIDevelope business process took an average of 19229.0 ms to complete, while the Deenveloping step (1 DeenveloperCII) took an average time of 18573.0 ms. Both of these are the average statistics; there were two invocations of the business process.</p> <p>Use this statistic to track overall averages of business processes or activities that begin to slow in processing times. If you notice this statistic increasing over your benchmark value, this may indicate a performance issue. If you see an occasional variance for a business or activity, this does not indicate a performance issue. If, however, you notice a continual variance between production statistics and your benchmarks, you must resolve the issue.</p> <p>For more information on slow systems, see <i>Resolving Slow System Issues</i> on page 80.</p> <p>For more information about improving business process execution times, see <i>Improving Business Process Execution Times</i> on page 112.</p>
Invocations	<p>Number of times the business process or service was invoked since the report began. The number of invocations should be the same for the business process and each activity in the business process, unless you have a business process that loops several times. In this case, you may set the number of invocations lower than the number of activities in the business process.</p> <p>For example, in the preceding figure, the CIIDevelope business process had 2 invocations, and the Deenveloping step (1 DeenveloperCII) had 2 invocations.</p> <p>Use this statistic to determine if the number of business processes expected are running and if all activities in the business process are running. If you see a variance between the business process invocation number and the activity number, this indicates that a business process may have an error and is in a waiting, interrupted, or halted state. You can also use this statistic to determine volume loads on your system during different processing periods, such as peak and non-peak processing hours.</p> <p>For more information about business processes in a halted, interrupted, or waiting state, see <i>Resolving Business Process Issues</i> on page 103.</p>
Total (ms)	<p>Total time in milliseconds (ms) that the business process or service has taken to process since you created the report. This is the sum total of time of all business process completions.</p> <p>For example, in the preceding figure, the CIIDevelope business process took 38458 ms to process since the report was created and the Deenveloping step (1 DeenveloperCII) has taken a total of 37146 ms to process since the report was created.</p> <p>Use this statistic to see if the total processing time for a business process or activity is increasing substantially over the number of invocations. If this number does increase over the number of invocations, this indicates that the processing times are longer. If the processing time is longer than your benchmark, you must resolve the issue.</p> <p>For more information about improving business execution times, see <i>Improving Business Process Execution Times</i> on page 112.</p>

Statistic Name	Description
Pre-Service (ms)	<p>Total time that the business process engine completed activities before the business process or service ran.</p> <p>For example, in the preceding figure, the DeenvelopeCII activity took 37146 ms for pre-service activities. Every time an activity runs, Gentran Integration Suite must complete system service invocations, before running the activity.</p> <p>Use this statistic to see if the system services are taking long periods of time to run before running the business process. If you see a variance between the report and the benchmark times, this may indicate that there is a database, operating system, or other system issue that needs resolved.</p> <p>For more information about database issues, see <i>Resolving Database Issues</i> on page 89.</p> <p>For more information about operating system issues, see your vendor's documentation.</p> <p>For more information about slow systems, see <i>Resolving Slow System Issues</i> on page 80.</p>
Post-Service (ms)	<p>Total time that the business process engine completed activities after the business process or service ran.</p> <p>For example, in the preceding figure, the DecisionEngineService activity took 100.0 ms for post-service activities. Every time an activity runs, Gentran Integration Suite completes some persistence of information to the database. This persistence takes time to complete, and Post-Service time includes this persistence time.</p> <p>Use this statistic to see if the system services are taking long periods of time to run after running the business process. If you see a variance between the report and the benchmark times, this may indicate that there is a database, operating system, or other system issue that needs resolved.</p> <p>For more information about database issues, see <i>Resolving Database Issues</i> on page 89.</p> <p>For more information about operating system issues, see your vendor's documentation.</p> <p>For more information about slow systems, see <i>Resolving Slow System Issues</i> on page 80.</p>

The following figure shows a sample of the Internal System Statistics portion of the Performance Statistics Report:

Name	Min (ms)	Max (ms)	Average (ms)	Invocations	Total (ms)
Persist a Document to the Database	9	212	111.0	9	1003
Persist a Business Process Step to the Database	5	391	59.0	163	9735
Assign from an XPath Statement	2	43	16.0	29	310
Execute an XPath Statement	1	43	5.0	111	598
Acquire a Database Connection	1	4990	3.0	15102	57722
Put a Business Process on the Queue	1	49	3.0	38	146

Statistic Name	Description
Internal System Statistics	

Name	<p>Name of the system activity, including:</p> <ul style="list-style-type: none"> ◆ Persist a Business Process Step to the Database ◆ Persist a Document to the Database ◆ Assign from an XPath Statement ◆ Acquire a Database Connection ◆ Execute an XPath Statement ◆ Put a Business Process on the Queue <p>For more information about each system activity, see <i>Internal System Statistics</i> on page 38.</p>
Min (ms)	<p>Minimum time in milliseconds (ms) the system activity took to complete. If there are multiple invocations, this is the lowest time for all of the invocations.</p> <p>For example, in the preceding figure, the Persist a Document to the Database activity took 9 ms to complete. This is the lowest statistic of the nine invocations of the activity.</p>
Max (ms)	<p>Maximum time in milliseconds (ms) the system activity took to complete. If there are multiple invocations, this is the highest time for all of the invocations.</p> <p>For example, in the preceding figure, the Persist a Document to the Database activity took 212 ms to complete. This is the highest statistic of the nine invocations of the activity.</p>
Average (ms)	<p>Average processing time for the system activity to complete. This is the average time for all invocations.</p> <p>For example, in the preceding figure, the Persist a Document to the Database activity took an average of 111.0 ms to complete. This is the average processing time of the nine invocations of the activity.</p>
Invocations	<p>Number of times the system activity was invoked since you created the report.</p> <p>For example, in the preceding figure, the Persist a Document to the Database activity was invoked nine times.</p>
Total (ms)	<p>Total time in milliseconds (ms) that the system activity has taken to process since the report began. This is the sum total of time of all system activity completions.</p> <p>For example, in the preceding figure, the Persist a Document to the Database activity took a total of 1003 ms to complete.</p>

Internal System Statistics

The Internal System Statistics part of the Performance Statistics Report includes a great amount of information that you can use to monitor your Gentran Integration Suite system to keep it performing up to the benchmarks you establish.

The following table describes the internal system activity and how to use the statistics to reduce the chance of a performance issue:

Statistic Name	Description
Persist a Document to the Database	<p>Statistics showing the length of time and the number of times that a document was persisted to the database.</p> <p>For more information about persistence, see <i>Changing Persistence Levels</i> on page 68.</p> <p>For example, in the preceding figure, the Persist a Document to the Database had the following statistics:</p> <ul style="list-style-type: none"> ◆ Min (ms) – 9 ◆ Max (ms) – 212 ◆ Average (ms) – 111.0 ◆ Invocations – 9 ◆ Total (ms) – 1003 <p>This statistic is helpful in determining database issues and slow processing issues. When you compare this information to your benchmarks, if you notice:</p> <ul style="list-style-type: none"> ◆ The Min (ms), Max (ms), and Average (ms) times are increasing, this may indicate that the database is becoming full, or you have a connection leak. For more information about databases, see <i>Resolving Database Issues</i> on page 89. ◆ The number of invocations increasing, this may indicate that you have your persistence level set too high. For more information on slow systems, see <i>Changing Persistence Levels</i> on page 68. ◆ The number of invocations is small and the Min (ms) and Max (ms) times are increasing, this may indicate that you are persisting large documents to the database, which can be moved to disk to save database space. For more information about databases, see <i>Resolving Database Issues</i> on page 89.

Statistic Name	Description
Persist a Business Process Step to the Database	<p>Statistics showing the length of time and the number of times that a business process step was persisted to the database.</p> <p>For more information about persistence, see <i>Changing Persistence Levels</i> on page 68.</p> <p>For example, in the preceding figure, the Persist a Business Process Step to the Database had the following statistics:</p> <ul style="list-style-type: none"> ◆ Min (ms) – 5 ◆ Max (ms) – 391 ◆ Average (ms) – 59.0 ◆ Invocations – 163 ◆ Total (ms) – 9735 <p>This statistic is helpful in determining database issues and slow processing issues. When you compare this information to your benchmarks, if you notice:</p> <ul style="list-style-type: none"> ◆ The Min (ms), Max (ms), and Average (ms) times are increasing, this may indicate that the database is becoming full, or you have a connection leak. For more information about databases, see <i>Resolving Database Issues</i> on page 89. ◆ The number of invocations increasing, this may indicate that you have your persistence level set too high. For more information on slow systems, see <i>Changing Persistence Levels</i> on page 68. ◆ The number of invocations is small and the Min (ms) and Max (ms) times are increasing, this may indicate that you are persisting large documents to the database, which can be moved to disk to save database space. For more information about databases, see <i>Resolving Database Issues</i> on page 89.
Assign from an XPath Statement	<p>Statistics showing the length of time and the number of times that an assign activity completed from an XPath statement in a business process.</p> <p>For example, in the preceding figure, the Assign from an XPath Statement had the following statistics:</p> <ul style="list-style-type: none"> ◆ Min (ms) – 2 ◆ Max (ms) – 43 ◆ Average (ms) – 10.0 ◆ Invocations – 29 ◆ Total (ms) – 310 <p>This statistic is helpful in determining if you have well-structured XPath statements in your business processes. When you compare this information to your benchmarks, if you notice:</p> <ul style="list-style-type: none"> ◆ The Min (ms), Max (ms), and Average (ms) times are increasing, this may indicate that the XPath statement is not written efficiently, which may slow your process times. Write XPath statements using relative paths. For example, write <code>PurchaseOrder/text</code> instead of <code>/ProcessData/PurchaseOrder/text()</code>. In addition, do not use <code>//</code> at the beginning of an XPath statement, because this causes the entire process data to be traversed. For more information about databases, see <i>Improving Business Process Execution Times</i> on page 112.

Statistic Name	Description
Execute an XPath Statement	<p data-bbox="467 262 1419 317">Statistics showing the length of time and the number of times that an XPath statement ran in a business process.</p> <p data-bbox="467 327 1338 382">For example, in the preceding figure, the Assign from an XPath Statement had the following statistics:</p> <ul data-bbox="467 401 711 604" style="list-style-type: none"><li data-bbox="467 401 634 428">◆ Min (ms) – 1<li data-bbox="467 443 659 470">◆ Max (ms) – 43<li data-bbox="467 485 711 512">◆ Average (ms) – 5.0<li data-bbox="467 527 691 554">◆ Invocations – 111<li data-bbox="467 569 683 596">◆ Total (ms) – 598 <p data-bbox="467 621 1419 676">This statistic is helpful in determining if you have well-structured XPath statements in your business processes. When you compare this information to your benchmarks, if you notice:</p> <ul data-bbox="467 690 1419 835" style="list-style-type: none"><li data-bbox="467 690 1419 835">◆ The Min (ms), Max (ms), and Average (ms) times are increasing, this may indicate that the XPath statement is not written efficiently, which may slow your process times. Write XPath statements using relative paths. For example, write <code>PurchaseOrder/text</code> instead of <code>/ProcessData/PurchaseOrder/text()</code>. In addition, do not use <code>//</code> at the beginning of an XPath statement, because this causes the entire process data to be traversed. <p data-bbox="505 850 1386 905">For more information about databases, see <i>Improving Business Process Execution Times</i> on page 112.</p>

Statistic Name	Description
Acquire a Database Connection	<p>Statistics showing the length of time and the number of times that a database connection was made.</p> <p>For example, in the preceding figure, the Acquire a Database Connection had the following statistics:</p> <ul style="list-style-type: none"> ◆ Min (ms) – 1 ◆ Max (ms) – 4990 ◆ Average (ms) – 3.0 ◆ Invocations – 16102 ◆ Total (ms) – 57722 <p>This statistic is helpful in determining if you have database issues, a resource leak, or you need to increase the number of database pools. When you compare this information to your benchmarks, if you notice:</p> <ul style="list-style-type: none"> ◆ The Min (ms), Max (ms), and Average (ms) times are increasing, this may indicate that the database is quite active, if it is used by other applications outside of Gentran Integration Suite. ◆ The number of invocations increasing, this may indicate that you have your persistence level set too high causing too much to be stored to the database, or your cache levels are not tuned correctly causing the data to be stored and retrieved from the database instead of from cache. <p>For more information on slow systems, see <i>Changing Persistence Levels</i> on page 68. For more information about databases, see <i>Resolving Cache Issues</i> on page 108.</p> <ul style="list-style-type: none"> ◆ The number of invocations is low and not increasing, this may indicate that you have a resource leak that is not releasing previously used database connections to be used by other threads, or you do not have enough database pools set in the Performance Tuning Utility. <p>For more information on resource leaks, see <i>Full Database Issues</i> on page 89. For more information on database pools, see <i>Calculated Recommendations</i> on page 45.</p>
Put a Business Process on the Queue	<p>Statistics showing the length of time and the number of times that a business process was placed on the queue.</p> <p>For example, in the preceding figure, the Put a Business Process on the Queue had the following statistics:</p> <ul style="list-style-type: none"> ◆ Min (ms) – 1 ◆ Max (ms) – 49 ◆ Average (ms) – 3.0 ◆ Invocations – 38 ◆ Total (ms) – 146 <p>This statistic is helpful in determining if you are using the queue instead of your cache in processing. When you compare this information to your benchmarks, if you notice:</p> <ul style="list-style-type: none"> ◆ Increasing invocation times, this may indicate that you are not using your cache efficiently. <p>For more information about databases, see <i>Resolving Cache Issues</i> on page 108.</p>

Turning On and Off Performance Statistics

You may find that you want to have performance statistics turned on during specific time periods and turned off during other time periods. For example, turn on the performance statistics when you are testing or when you are running your weekly or monthly report to compare to your benchmarks, and turn off performance statistics during day-to-day operation to save your system resources for processing.

Turning On Performance Statistics

To turn on performance statistics:

1. From the **Administration** menu, select **Operations > System > Performance > Statistics**.
2. On the Performance Statistic Report Manager page, under **On/Off**, select the check box next to **Enable Performance Statistics**.

Performance Statistics are turned on and ready to be captured in a report.

Turning Off Performance Statistics

To turn off performance statistics:

1. From the **Administration** menu, select **Operations > System > Performance > Statistics**.
2. On the Performance Statistic Report Manager page, under **On/Off**, clear the check box next to **Enable Performance Statistics**.

Performance Statistics are turned off.

Creating or Deleting a Performance Statistics Report

Creating a New Performance Statistics Report:

After you have enabled performance statistics, you need to create a new report before you can view it.

1. From the **Administration** menu, select **Operations > System > Performance > Statistics**.
2. On the Performance Statistic Report Manager page, under **Create New Report**, type a name for the new report in the **Name** field, and click **Go!**

The new report is created and the name displays in the list under **View Report**. The View Report list displays the most current report at the top of the list with the oldest report at the bottom of the list.

Deleting a Performance Statistics Report:

You may need to delete Performance Statistics Reports to solve storage issues or to simply remove old reports from Gentran Integration Suite.

1. From the **Administration** menu, select **Operations > System > Performance > Statistics**.
2. On the Performance Statistic Report Manager page, under **Delete Reports**, next to **Delete All Inactive Reports**, click **Go!**

All reports except the most current report are deleted from Gentran Integration Suite.

Performance Tuning Utility

Gentran Integration Suite provides a utility that calculates the majority of tuning parameter values for you. Using the Performance Tuning Utility, you can easily tune Gentran Integration Suite. Performance tuning enables you to configure Gentran Integration Suite for the most efficient performance for your specific needs. You can change the database settings, memory allocations, and pool settings, so that Gentran Integration Suite performs as you want it too, however, you cannot use the utility to tune your business processes.

This section covers the following topics:

- ◆ *Calculated Recommendations* on page 45
- ◆ *Viewing or Editing Performance Configuration Settings* on page 46

Calculated Recommendations

Gentran Integration Suite calculates recommended settings according to the number of CPUs and the amount of memory you specify. These performance properties are stored in the tuning.properties file in the *install_dir*/properties directory. The calculations are made according to formulas in the tuningFormulas.properties file in the *install_dir*/properties directory.

The following table shows the default and calculated performance tuning properties for a 6 CPU system:

Performance Tuning Property (Formula)	User Identified Value	User Identified Value
Number of CPU's in system (shipped value =4)	2	6
Physical memory (MB) allocated to Gentran Integration Suite (shipped value = 1024)	768	2048
The following properties are calculated according to the number of CPUs and physical memory	Calculated Value	Calculated Value

Desired Global Threads (4 * processor number)	8	24
In memory cache size (MB) for small contexts (0.25 * Physical Memory)	192	512
Disk cache size (MB) (2 * Physical Memory)	1536	512
Transactional pool connections (initial) ((4 * processor number) + 5)	13	44
Transactional pool connections (max) ((4 * processor number) + 30)	38	44
Non-transactional pool connections (initial) ((4 * processor number) + 5)	13	29
Non-transactional Pool Connections (max) ((4 * processor number) + 5)	13	29
JVM long-lived memory (min) (Any number greater than 512 is calculated to 2048)	768	2048
JVM long-lived memory (max) (Any number greater than 512 is calculated to 2048)	768	2048
JVM short-lived memory (min) (0.33333 * Physical Memory)	256	683
JVM short-lived memory (max) (0.33333 * Physical Memory)	256	683
Number of BP steps executed before returning to queue (No formula)	100	100
Initial steps in the first execute cycle (No formula)	2	2
Translation maps Cache (No formula)	100	100
Envelopes Cache (No formula)	500	500
Other EDI Cache (No formula)	100	100

Viewing or Editing Performance Configuration Settings

Viewing Performance Configuration Settings

You may not need to change performance configurations; you may only need to view the performance configuration to verify that a performance configuration setting is correct.

To view your performance configuration:

1. From the **Administration** menu, select **Operations > System > Performance > Tuning**.
2. On the Performance Tuning page, under **View**, next to **View Performance Configuration**, click **Go!**
3. On the Performance Settings page, review the performance configuration settings.

Click **Return** to return to the Performance Tuning page.

Editing Performance Configuration Settings

You may find that you need to allocate more or less system memory to specific components, more or less database pools, or more or less cache to specific components of Gentran Integration Suite to reach the level of performance you want. You can use the performance tuning utility to accomplish your goals.

You should configure performance based on your processing and volume requirements. If you have used Sterling Commerce Consulting Services, review the recommended requirements report. If you have not used Sterling Commerce Consulting Services, refer to the *Requirements Planning Worksheet* on page 117 and tune Gentran Integration Suite according to those requirements.

You can use the performance tuning utility to edit the Gentran Integration Suite performance configuration using recommended settings calculated by the system or, if you have specific needs, you can edit the existing performance configurations without the recommendations displaying.

To edit your performance configuration, perform the following steps:

1. From the **Administration** menu, select **Operations > System > Performance > Tuning**.
2. On the Performance Tuning page, under **Edit**, next to **Edit Performance Configuration**, click **Go!**
3. In the Performance Tuning UI lock enabled message, click **OK**.

The Performance Tuning UI lock releases when you click **Cancel** or **Finish** during the editing process. If the lock is not released or if you close the page without clicking **Cancel** or **Finish**, you must use the Lock Manager function to release the lock.

4. On the Edit Mode page, select one of the following options:
 - ◆ **Generate recommended settings** – Gentran Integration Suite calculates recommended settings based on the number of CPUs and amount of memory you specify. See *Calculated Recommendations* on page 45 for more information about calculated recommended settings. The performance tuning pages will display both the current setting and the recommended setting for each property. You can either accept the recommended setting or enter a different value in the field under the **Recommended** column, based on your performance requirements.
 - ◆ **Edit settings** – Gentran Integration Suite will display the current setting for each property. You can either keep the current setting or enter a new value based on your performance requirements.
5. Click **Next**.
6. On the System page, do you want to use the current or recommended settings?
 - ◆ If Yes, click **Next**.

- ◆ If No, complete the following fields and click **Next**:

Field	Description
Number of CPU(s)	Number of actual CPUs in the system. Gentran Integration Suite does not scan your system to obtain this information, therefore you must enter a value in this field. This field changes PROCESSORS in the tuning.properties file in the <i>install_dir</i> /properties directory. Shipped value is 2. Note: This field is only displayed if you selected Generate recommended settings on the Edit Mode page.
Physical memory (MB) allocated to Gentran Integration Suite	Amount of memory allocated for use in processing Gentran Integration Suite operations. This field changes MEMORY in the tuning.properties file in the <i>install_dir</i> /properties directory. Shipped value is 768 MB.

If you selected **Generate recommended settings** on the Edit Mode page, the remaining recommended settings are calculated according to the number of CPUs and allocated memory you specify.

7. On the BP Queue page, do you want to use the current or recommended settings?
 - ◆ If Yes, click **Next**.
 - ◆ If No, complete the following fields and click **Next**:

Field	Description
Desired Global Threads	Number of active business processes that can run concurrently. This field changes GLOBAL_THREADLIMIT and JMS_LISTNERS in the tuning.properties file in the <i>install_dir</i> /properties directory. Shipped value is 8. Typically, you need higher values with more CPUs, and lower values with fewer CPUs. Too many threads causes <i>thrashing</i> , which is an issue of the hard disk being used too much for virtual memory and slowing down performance.
Disk cache size (MB)	Maximum amount of disk space used to store business process context when it is being moved from memory. This limit ensures that Gentran Integration Suite does not consume all available disk space. This field changes JMS_PAGING_MAX in the tuning.properties file in the <i>install_dir</i> /properties directory. Shipped value is 1536 MB
In memory cache size (MB) for small contexts	Amount of memory allocated to the business process queue. After Gentran Integration Suite reaches this limit, Gentran Integration Suite writes all messages, except active messages, to disk. This field changes JMS_PAGING_HIGH in the tuning.properties file in the <i>install_dir</i> /properties directory. Shipped value is 192 MB. Increasing this value too much reduces the amount of memory available for other components of Gentran Integration Suite, such as caches and general processing.

8. On the Database Pool page, do you want to use the current or recommended settings?
 - ◆ If Yes, click **Next**.
 - ◆ If No, complete the following fields and click **Next**:

Transactional pool connections are used by the workflow engine to manage database operations during a workflow step. Commands issued to a transactional connection are not committed to the database

until the end of the workflow step, allowing rollback if the workflow step fails. Non-transactional pool connections are used when rollback is not required.

Field	Description
Transactional pool connections (initial)	<p>Initial number of connections that are reserved and available for critical data requests (for example, persisting workflow data) when Gentran Integration Suite starts. If Gentran Integration Suite uses all connections concurrently, Gentran Integration Suite creates a new connection until it reaches the maximum number of connections.</p> <p>The pool size must be large enough to process business processes without waiting for a connection. In addition, if you have many File System adapters, you may need to increase this property to improve performance.</p> <p>This field changes MIN_TRANS_POOL in the tuning.properties file in the <i>install_dir</i>/properties directory. Shipped value is 13.</p>
Transactional pool connections (max)	<p>Maximum number of connections that are reserved and available for critical data requests that Gentran Integration Suite can use concurrently. After a connection is no longer in use, the connection is returned to the available pool connections and can be reused. After Gentran Integration Suite reaches the value of this field, Gentran Integration Suite must wait for a connection to be returned to the pool before processing the request.</p> <p>The pool size must be large enough to process business processes without waiting for a connection. In addition, if you have many File System adapters, you may need to increase this property to improve performance.</p> <p>This field changes MAX_TRANS_POOL in the tuning.properties file in the <i>install_dir</i>/properties directory. Shipped value is 38.</p>
Non-transactional pool connections (initial)	<p>Initial number of connections that are reserved and available for general data requests (for example, retrieving data) when Gentran Integration Suite starts. If Gentran Integration Suite uses all connections concurrently, Gentran Integration Suite creates a new connection until it reaches the maximum number of connections.</p> <p>The pool size must be large enough to process business processes without waiting for a connection. In addition, if you have many File System adapters, you may need to increase this property to improve performance.</p> <p>This field changes MIN_NONTRANS_POOL in the tuning.properties file in the <i>install_dir</i>/properties directory. Shipped value is 13.</p>
Non-transactional Pool Connections (max)	<p>Maximum number of connections that are reserved and available for general data requests that Gentran Integration Suite can use concurrently. After a connection is no longer in use, the connection is returned to the available pool connections and can be reused. After Gentran Integration Suite reaches the value of this field, Gentran Integration Suite must wait for a connection to be returned to the pool before processing the request.</p> <p>The pool size must be large enough to process business processes without waiting for a connection. In addition, if you have many File System adapters, you may need to increase this property to improve performance.</p> <p>This field changes MAX_NONTRANS_POOL in the tuning.properties file in the <i>install_dir</i>/properties directory. Shipped value is 13.</p>

If you have purchased additional adapters, you may need to increase the database pools to reduce the chances that a business process or adapter will have to wait for a database connection.

9. On the Memory page, do you want to use the current or recommended settings?
 - ◆ If Yes, click **Next**.

- ◆ If No, complete the following fields and click **Next**:

Field	Description
JVM long-lived memory (initial)	Initial amount of JVM memory that Gentran Integration Suite reserves for long-lived objects when Gentran Integration Suite starts. This field changes INIT_HEAP in the tuning.properties file in the <i>install_dir</i> /properties directory. Shipped value is 768.
JVM long-lived memory (max)	Maximum amount of JVM memory that Gentran Integration Suite can use for long-lived objects. This field changes MAX_HEAP in the tuning.properties file in the <i>install_dir</i> /properties directory. Shipped value is 768.
JVM short-lived memory (initial)	Initial amount of JVM memory that Gentran Integration Suite uses for short-lived objects. This field changes INIT_AGE in the tuning.properties file in the <i>install_dir</i> /properties directory. Shipped value is 256. Note: JVM short-lived memory has a fast trash collection rate. To avoid reduced performance, set the JVM short-lived memory to one third that of the JVM long-lived memory.
JVM short-lived memory (max)	Maximum amount of JVM memory that Gentran Integration Suite uses for short-lived objects. This field changes MAX_AGE in the tuning.properties file in the <i>install_dir</i> /properties directory. Shipped value is 256. Note: JVM short-lived memory has a fast trash collection rate. To avoid reduced performance, set the JVM short-lived memory to one third that of the JVM long-lived memory

10. On the BP Execution page, do you want to use the current or recommended settings?

- ◆ If Yes, click **Next**.
- ◆ If No, complete the following fields and click **Next**:

Field	Description
Number of BP steps executed before returning to queue	Number of business process steps that are run before returning to the JMS queue. This field changes BP_STEPS in the tuning.properties file in the <i>install_dir</i> /properties directory. Shipped value is 100. Typically, this property should match the number of steps in your largest business process. If your business processes contain loops, you may need to increase this property to improve performance. If you find that you need to increase this number, you should streamline your business processes to remove unnecessary loops. If you set this value too low, business processes may not be able to complete processing before they are returned to the queue allowing another business process to complete some or all of its processing. This slows processing times and causes bottlenecks in Gentran Integration Suite. For example, if you have five business processes with five activities each, and you set the number of business process steps executed before returning to queue at one, the first business process completes one activity, then the second business completes one activity. This process continues, until all five business have completed the first activity. Then the process begins again with each business process completing the second activity. This continues until all business processes have completed all five activities. This is one example of how Gentran Integration Suite processing is slowed if properties are not tuned correctly.

Field	Description
Initial steps in the first execute cycle	Number of steps executed before returning a business process to the queue on its first execution cycle. This field changes ASYNC_BP in the tuning.properties file in the <i>install_dir/properties</i> directory. Shipped value is 2.

11. On the Cache page, do you want to use the current or recommended settings?

- ◆ If Yes, click **Next**.
- ◆ If No, complete the following fields and click **Next**:

Field	Description
Translation Maps	Amount of cache used to store translation maps. This field changes MAP_CACHE in the tuning.properties file in the <i>install_dir/properties</i> directory. Shipped value is 100. Note: Increasing the cache size for items that are not used frequently may degrade performance. Increase the cache sizes for items that are used frequently to improve performance.
Envelopes	Amount of cache used to store envelopes. This field changes ENVELOPE_CACHE in the tuning.properties file in the <i>install_dir/properties</i> directory. Shipped value is 500. Note: Increasing the cache size for items that are not used frequently may degrade performance. Increase the cache sizes for items that are used frequently to improve performance.
Other EDI	Amount of cache used to store EDI data. This field changes EDI_CACHE in the tuning.properties file in the <i>install_dir/properties</i> directory. Shipped value is 100. Note: Increasing the cache size for items that are not used frequently may degrade performance. Increase the cache sizes for items that are used frequently to improve performance.

12. On the Confirm page, review the changes you made to the performance configuration, and click **Finish**.

Gentran Integration Suite saves the edited performance configurations.

13. From the **Administration** menu, select **Operations > System > Troubleshooter**.

14. On the System Troubleshooting page, click **Stop the System**.

Stop the System runs the softstop script allowing all business processes in progress to complete before stopping the system.

15. In the *install_dir/bin* directory, complete one of the following actions depending on your operating system:

- ◆ UNIX or Linux – Run setupfiles.sh
- ◆ Windows – Run setupfiles.cmd

The performance configuration changes populate the initial setup files that are used during startup.

16. In the *install_dir/bin* directory, start Gentran Integration Suite by running one of the following scripts depending on your operating system:

- ◆ UNIX or Linux – Run run.sh

- ◆ Windows – Run startWindowsService.cmd or click the icon.

Gentran Integration Suite starts and the updated performance configuration settings are applied.

Manual Performance Tuning

You may find that you need to tune some application server independent (ASI) properties that are not tuned in the performance tuning utility. You can manually tune these properties in the *install_dir/properties/noapp.properties* file in Gentran Integration Suite.

This section covers the following topics:

- ◆ *Tuning Your ASI Environment* on page 54
- ◆ *Scheduling Policy* on page 58
- ◆ *Queues in an ASI Environment* on page 59
- ◆ *Monitoring Queues using queueWatcher* on page 64
- ◆ *Compressing Cache Contents* on page 66
- ◆ *Setting the Document Body Serialization Threshold Property* on page 67
- ◆ *Allocating JNDI Contexts* on page 68
- ◆ *Changing Persistence Levels* on page 68
- ◆ *Changing System Logging* on page 70
- ◆ *Enabling Trusted Domains for Schemas* on page 71
- ◆ *Enabling the Console Listener* on page 71
- ◆ *Enabling the JMX Event Listener* on page 72
- ◆ *Managing Property Files* on page 74
- ◆ *Managing System Recovery* on page 74
- ◆ *Changing the SoftStop Time Limit* on page 76
- ◆ *Changing Advanced File Transfer Settings* on page 76

Note: For large volume operations of FTP, Mailbox, Advanced File Transfer, and File Broker, the number of open file descriptors should be set to at least 4096, using the command `ulimit -n 4096` (recommended value is unlimited).

Tuning Your ASI Environment

You can tune your application server independent (ASI) system using the properties found in the *install_dir/properties/noapp.properties* file.

Properties that you may need to tune include:

Property	Description
SchedulingPolicyName	Set of rules that Gentran Integration Suite uses make decisions about how to manage the workload. Valid values are: <ul style="list-style-type: none"> ◆ BasicSchedulingPolicy ◆ FairShareSchedulingPolicy For more information, see <i>Scheduling Policy</i> on page 58.
MemCacheSize	Size, in megabytes, of the in-memory cache used for small objects. This cache can speed up execution by eliminating the need for the data to be retrieved from the database or disk.
DiskCacheSize	Maximum size of the disk cache, so that your cache does not increase beyond your available disk space.
DiskCachePath	Directory to use for caching objects.
MemCacheThreshold	Threshold size, in bytes, for caching a context in the in-memory cache versus the disk cache. In general, the distribution of context sizes tends to look vaguely like an “M” with one cluster of small contexts and another cluster of larger contexts. If this value is set large, the value for MemCacheSize should also be large. A moderate multiple of the disk block size seems to work very well in many cases.
MaxThreads	Total number of concurrent threads that Gentran Integration Suite is allowed to use. This number may be checked against the licensed number of threads.
AE_ExecuteCycle.#	Number of steps for a business process to complete prior to returning to the queue. Higher values will accelerate individual business process execution, while lower values will provide smoother multi-tasking. Interactive use favors a lower number of steps while batch processing favors a higher number of steps. This value can be different for each queue. The .# indicates the queue number.
AE_ExecuteCycleTime.#	Maximum time period, in milliseconds, that a business process can use a thread, before releasing it to be used for another business process. This value will override the value set for AE_ExecuteCycle . It is intended to ensure that a series of unusually slow steps will not tie up a thread completely. This value can be different for each queue. The .# indicates the queue number.
QueueDepth.#	Maximum number of business processes that can be in the queue at one time. Generally, this value should be left at the default value of 10000, unless you anticipate having more than 10,000 business processes in the queue at the same time. This value can be different for each queue. The .# indicates the queue number.

Property	Description
MaxPoolSize.#	<p>Maximum number of threads executed for the specific queue. The MaxThreads value overrides this value if this value is set higher than the MaxThreads value. This value can be different for each queue. The .# indicates the queue number.</p> <p>Note: Setting all queues MaxPoolSize parameters to the maximum threads available does not mean you will have faster processing. Depending on your system resources, setting the MaxPoolSize value too high may cause the queues to back up and degrade performance. A good place to start is 4 times the number of your CPUs. Interactive use tends to favor more threads, while batch or document processing tends to favor less threads.</p>
CacheThreshold.#	<p>Number of business processes that must be in the queue before any business processes are cached. In addition, this value is the minimum number of business processes in the queue before any rescheduling occurs. This value can be different for each queue. The .# indicates the queue number.</p> <p>In general, setting this value high improves performance by keeping more business process contexts in memory when they are placed in the queue. To estimate the amount of memory that will be consumed, multiply this value by the average size of your business process contexts. You can obtain the average context size using the Gentran Integration Suite System Troubleshooter. For more information, see <i>Reviewing System Information</i> on page 18.</p> <p>There is a tradeoff between performance and memory consumption. Setting this value too high can:</p> <ul style="list-style-type: none"> ◆ Leave Gentran Integration Suite with insufficient memory in some circumstances ◆ Cause some business processes to remain in the queue too long without being examined for rescheduling. <p>Note that, in most cases, even if this value is set low, business process contexts will usually be recovered from one of the in-memory caches and not from the disk. The performance impact is usually seen as the context is placed in the queue.</p>
MinPoolSize.#	<p>Minimum number of threads reserved for the specific queue. The sum total of the queues MinPoolSize values must be equal to or less than the MaxThreads value. This value can be different for each queue. The .# indicates the queue number.</p>
JavaPriority.#	<p>Java priority of the threads running business processes. This enables you to set some queues to run more slowly to reduce the issues associated with heavily used computers being unresponsive to the interface. This value can be different for each queue. Not all JVMs handle this the same way, and some do not respect it at all (particularly the IBM AIX JVM). The .# indicates the queue number.</p>
EnableDeadlines.#	<p>Enables or disables deadline support for this queue. Having deadlines enabled controls both notifications and the execution order in the queue. If enabled, business processes with deadlines are executed before those without deadlines. Valid values are:</p> <ul style="list-style-type: none"> ◆ false – Disables deadline support for the queue. ◆ true – Enables deadline support for the queue. <p>This value can be different for each queue. The .# indicates the queue number.</p>
Rescheduling.#	<p>Enables or disables rescheduling support for this queue. Valid values are:</p> <ul style="list-style-type: none"> ◆ false – Business processes are not rescheduled in this queue. ◆ true – Business processes are rescheduled in this queue. <p>This value can be different for each queue. The .# indicates the queue number.</p>

Property	Description
ReschedulingInterval.#	Time interval, in milliseconds, to wait before the rescheduler passes through the queue to reschedule business processes. This value can be different for each queue. The .# indicates the queue number.
MaxWaitTime.#	<p>Maximum time, in milliseconds, a business process can be in the queue without it being reviewed for rescheduling. If a business process is rescheduled, the business process is moved forward in the queue. This value can be different for each queue. The .# indicates the queue number.</p> <p>Example: MaxWaitTime.8=60000 sets the maximum wait time for queue number 8 to 60 seconds.</p>
ResourceAllocation.#	<p>Amount of resources to allocate to this specific queue for fair share scheduling. For example, if you have only two queues configured, queue 1 for a few small business processes and queue 2 for many large business processes, you can set the amount of resources available to each queue by a percentage. You may set queue 1 with a resource allocation of 10 and queue 2 as 90. This ensures that queue 2 has more resources for processing; this improves processing times and ensures efficient use of system resources.</p> <p>Note: Set the resource allocations according to the percentages of use for each queue and the importance of processing completed on each queue. If you have higher priority items on one queue, increase the resource allocation to that queue to increase processing capability. Apply lower percentages of resources to lower priority or less used queues to keep resources free for higher priority processing.</p> <p>This value can be different for each queue. The .# indicates the queue number.</p>
InitialCycles.#	Number of cycles to execute the first time a business process gets to execute. Normally, this value should be 5 or less. It is intended to facilitate business processes (particularly web services) for which normal processing is quite short, but for which there is longer processing in special cases. It also allows a mix of short and long business processes in a queue, favoring the shorter ones. This value can be different for each queue. The .# indicates the queue number.
StealThreads.#	<p>Enables or disables the ability to steal threads from other queues. Valid values:</p> <ul style="list-style-type: none"> ◆ true – Enables the ability to steal threads ◆ false – Disables the ability to steal threads <p>Queues that are configured to steal threads from other queues (StealThreads property set to true) only steal from queues that have the AllowStealing property set to true. This value can be different for each queue. The .# indicates the queue number.</p>
AllowStealing.#	<p>Enables or disables the ability of other queues to steal threads from this queue. Valid values:</p> <ul style="list-style-type: none"> ◆ true – Allow other queues to steal threads ◆ false – Do not allow other queues to steal threads <p>Queues that are configured to steal threads from other queues (StealThreads property set to true) only steal from queues that have the AllowStealing property set to true. This value can be different for each queue. The .# indicates the queue number.</p>

Property	Description
persistence_level	<p>Persistence (storage to the database) level for Gentrans Integration Suite. Set the persistence_level value to PERSISTENCE_NONE during day-to-day processing.</p> <p>Valid values are:</p> <ul style="list-style-type: none"> ◆ PERSISTENCE_DEFAULT (defined in noapp.properties) ◆ PERSISTENCE_FULL – From user interface tracking, every step will be persisted in full, which includes all documents, status report (if any), and instance data. ◆ PERSISTENCE_MINIMAL/PERSISTENCE_STEP_STATUS – From user interface tracking, every step will be persisted, but document, status report (if any) and instance data will not be persisted in some of the steps (for example, assign, sleepservice). Most of the services that create documents will be persisted in full. ◆ PERSISTENCE_NONE/PERSISTENCE_BP_START_STOP – From user interface tracking, the first step and last step will be persisted in full. Some of the services will not be persisted (for example, assign, sleepservice). Services that are defined as full will be persisted in full. Documents created via a stream and stored in the file system will be persisted. ◆ PERSISTENCE_OVERRIDE_MINIMAL/PERSISTENCE_STEP_STATUS_ONLY – The first step will be persisted in full. The last step will be persisted in minimal. Documents created via a stream and stored in the file system will be persisted. ◆ PERSISTENCE_OVERRIDE_NONE – From user interface tracking, the first step will be persisted in full and the last step will be persisted in minimal. Some of the services will not be persisted (for example, assign, sleepservice). Services that are defined as full will not be persisted. Documents created via a stream and stored in the file system will be persisted. ◆ PERSISTENCE_OVERRIDE_NONE_KEEP_ERRORS/PERSISTENCE_BP_START_STOP_ONLY – The first step will be persisted in full and the last step will be persisted in minimal. Documents created via a stream and stored in the file system will be persisted. ◆ PERSISTENCE_WF_NONE (can only be used with workflow definition in sync mode) – From user interface tracking, nothing will be persisted. The business process runs in the caller's thread. If it is running in the Gentrans Integration Suite engine, if the service is creating document with document stream or if FS is document storage, then the document will be persisted. If it is running in embedded engine, document stream and file system document storage is not allowed. ◆ PERSISTENCE_WF_ERROR_ONLY – From user interface tracking, nothing will be persisted. Unless the service is doing document stream FS as document storage, only the first step will be persisted or when the business process errors out, then the first step (initial workflow context) and error step will be persisted. The error step will be persisted as minimal to avoid the resume problem. <p>If there is an onfault in the business process and the business process errors out, then the first step (initial workflow context), error step, and last step from the onfault will be persisted. If the onfault block runs fine, if onfault block errors out too, then error step in onfault will be persisted in minimal instead of last step in onfault. If there is step persisted, the first step initial workflow context always persisted in full.</p>

Scheduling Policy

Gentran Integration Suite uses scheduling policies to manage workload. A scheduling policy is a set of rules that Gentran Integration Suite uses to make decisions about how to manage the workload. Workload management decisions include, but are not limited to:

- ◆ Number of threads to run at one time. More threads running at one time can mean more work being completed at one time. You want to assign the number of threads efficiently, so you do not slow processing times by trying to run all business processes at one time.
- ◆ Order to run business processes.
- ◆ Number of steps for a business process to run before releasing the thread for use by other business processes.
- ◆ Caching requirements when the business process is returned to the queue.
- ◆ Workload distribution, if you are working in a clustered environment.

Scheduling policy choices include:

- ◆ **BasicScheduling Policy** – Recommended for consistent workloads that do not have data processing peaks or change in data or processing types.
- ◆ **FairShareScheduling Policy** – Recommended for mixed workloads that include both batch and online processing, and in environments that have data processing peaks. (Default).

Changing a Scheduling Policy

To change the scheduling policy in Gentran Integration Suite:

1. Open the `install_dir/properties/noapp.properties` file in a text editor.
2. In the `noapp.properties` file, locate the following code:

```
# SCHEDULING POLICY CONFIGURATION
```
3. Under `#SCHEDULING POLCY CONFIGURATION`, locate the following code:

```
SchedulingPolicyName=
```
4. Next to the `SchedulingPolicyName` parameter, type one of the following code samples depending on your choice of either the basic or the fair-share scheduling policy:
 - ◆ Basic scheduling policy:

```
com.sterlingcommerce.woodstock.workflow.queue.BasicSchedulingPolicy
```
 - ◆ Fair-share scheduling policy:

```
com.sterlingcommerce.woodstock.workflow.queue.FairShareSchedulingPolicy
```
5. Save the `install_dir/properties/noapp.properties` file. Changing the name or the location of the `noapp.properties` file may cause your changes to not take effect.

Queues in an ASI Environment

Gentran Integration Suite uses fair-share scheduling, which is in-memory queuing and caching, to enable you to set a global thread limit and to allocate resources on a queue-by-queue basis. This helps manage your mixed workloads with flexibility. The *global thread limit* is the maximum number of business processes that can run simultaneously in Gentran Integration Suite. You can change the thread limit and queue resource allocations in the *install_dir/properties/noapp.properties* file in the Queue Configuration ALL Policies and Queue Configuration FairShareSchedulingPolicy sections.

You can set both a maximum and minimum thread number for each queue you use. Gentran Integration Suite includes nine queues (Q1-Q9). When you check in a business process, you assign a queue to process the business process.

Each queue can always use its minimum thread number, but can never exceed its maximum thread number. The sum of the minimum numbers should be less than the **MaxThread** limit.

The following example explains the fair-share scheduling process:

1. You open the *install_dir/noapp.properties* file in a text editor and set the global number of threads as 16 by typing **16** as the new **MaxThreads** value.
2. You determine that queue Q1 is for small items and you set the **MinPoolSize** value to **10** and the **MaxPoolSize** value to **16**. This means that this queue always has 10 threads available for use, but if no other queues are using threads, this queue can use all 16 threads.
3. You determine that according to the number of small items requiring processing, that queue Q1 requires a resource allocation of 70 percent of the total resources and you type **70** as the new **ResourceAllocation** value.
4. You determine that queue Q8 is for batch processing and you set the **MinPoolSize** value to **1** and the **MaxPoolSize** value to **4**.
5. You determine that according to the number of batch items requiring processing, that queue Q8 requires a resource allocation of 30 percent of the total resources and you type **30** as the new **ResourceAllocation** value.
6. During processing, different conditions arise:
 - ◆ If there is no batch processing for queue Q8, but a great amount of processing for queue Q1, then queue Q1 uses all 16 threads to complete processing faster.
 - ◆ If there is some batch processing for queue Q8, but also some processing for queue Q1, then queue Q1 uses up to 12 threads to complete processing faster, while queue Q8 uses up to the 4 threads to complete processing faster. These numbers are according to the percentages applied to the **MaxThread** value of 16. In the case of queue Q1 with 70 percent of resource allocations, it is $16 * 0.70 = 11.2$. In the case of queue Q8 with 30 of resource allocations, it is $16 * 0.30 = 4.8$ rounded to 5, but because **MaxPoolSize** is set to **4**, only 4 threads can be used, enabling the 12th thread to be used by queue Q1.
 - ◆ If there is no processing for queue Q1, but a great amount of processing for queue Q8, then queue Q8 uses all 4 threads to complete processing faster, because the **MaxPoolSize** for queue Q8 is **4**.

Tuning Queues

To tune the queues in Gentran Integration Suite:

1. Open the *install_dir/properties/noapp.properties* file in a text editor.
2. In the *noapp.properties* file, under # POLICY CONFIGURATION -ALL Policies, locate the **MaxThreads** property and change the value to the value you determine is best for your level of processing.

Note: Setting the **MaxThreads** value higher does not mean you will have faster processing. Depending on your system resources, setting the **MaxThreads** value too high may degrade performance. Set **MaxThreads** according to your processing volumes in relation to your number of CPUs.

3. In the `noapp.properties` file, under `# QUEUE CONFIGURATION`, Queue # - ALL Policies, where # is queue number, locate the following properties and change the values for each as appropriate. Complete this step for each queue that you are using.

Property	Description
AE_ExecuteCycle.#	Number of steps for a business process to complete prior to returning to the queue. Higher values will accelerate individual business process execution, while lower values will provide smoother multi-tasking. Interactive use favors a lower number of steps while batch processing favors a higher number of steps. This value can be different for each queue. The .# indicates the queue number.
AE_ExecuteCycleTime.#	Maximum time period, in milliseconds, that a business process can use a thread, before releasing it to be used for another business process. This value will override the value set for AE_ExecuteCycle . It is intended to ensure that a series of unusually slow steps will not tie up a thread completely. This value can be different for each queue. The .# indicates the queue number.
QueueDepth.#	Maximum number of business processes that can be in the queue at one time. Generally, this value should be left at the default value of 10000, unless you anticipate having more than 10,000 business processes in the queue at the same time. This value can be different for each queue. The .# indicates the queue number.
MaxPoolSize.#	<p>Maximum number of threads executed for the specific queue. The MaxThreads value overrides the MaxPoolSize value if the MaxPoolSize value is set higher than the MaxThreads value. This value can be different for each queue. The .# indicates the queue number.</p> <p>Note: Setting all queues MaxPoolSize parameters to the maximum threads available does not mean you will have faster processing. Depending on your system resources, setting the MaxPoolSize value too high may cause the queues to back up and degrade performance. A good place to start is 4 times the number of your CPUs. Interactive use tends to favor more threads, while batch or document processing tends to favor less threads.</p>
CacheThreshold.#	<p>Number of business processes that must be in the queue before any business processes are cached. In addition, this value is the minimum number of business processes in the queue before any rescheduling occurs. This value can be different for each queue. The .# indicates the queue number.</p> <p>In general, setting this value high improves performance by keeping more business process contexts in memory when they are placed in the queue. To estimate the amount of memory that will be consumed, multiply this value by the average size of your business process contexts. You can obtain the average context size using the Gentran Integration Suite System Troubleshooter. For more information, see <i>Reviewing System Information</i> on page 18.</p> <p>There is a trade-off between performance and memory consumption. Setting this value too high can:</p> <ul style="list-style-type: none"> ◆ Leave Gentran Integration Suite with insufficient memory in some circumstances ◆ Cause some business processes to remain in the queue too long without being examined for rescheduling. <p>Note that, in most cases, even if this value is set low, business process contexts will usually be recovered from one of the in-memory caches and not from the disk. The performance impact is usually seen as the context is placed in the queue.</p>

MinPoolSize.#	<p>Minimum number of threads reserved for the specific queue. The sum total of the queues MinPoolSize values must be equal to or less than the MaxThreads value. The MinPoolSize value can be different for each queue. The .# indicates the queue number.</p> <p>Note: Setting the MinPoolSize parameter too low may cause queues to back up and degrade performance.</p>
JavaPriority.#	<p>Java priority of the threads running business processes. This enables you to set some queues to run more slowly to reduce the issues associated with heavily used computers being unresponsive to the interface. This value can be different for each queue. Not all JVMs handle this the same way, and some do not respect it at all (particularly the IBM AIX JVM). The .# indicates the queue number.</p>
EnableDeadlines.#	<p>Enables or disables deadline support for this queue. Having deadlines enabled controls both notifications and the execution order in the queue. If enabled, business processes with deadlines are executed before those without deadlines. Valid values are:</p> <ul style="list-style-type: none"> ◆ false – Disables deadline support for the queue. ◆ true – Enables deadline support for the queue. <p>This value can be different for each queue. The .# indicates the queue number.</p>
Rescheduling.#	<p>Enables or disables rescheduling support for this queue. Valid values are:</p> <ul style="list-style-type: none"> ◆ false – Business processes are not rescheduled in this queue. ◆ true – Business processes are rescheduled in this queue. <p>This value can be different for each queue. The .# indicates the queue number.</p>
ReschedulingInterval.#	<p>Time interval, in milliseconds, to wait before the rescheduler passes through the queue to reschedule business processes. This value can be different for each queue. The .# indicates the queue number.</p>
MaxWaitTime.#	<p>Maximum time a business process can be in the queue without being rescheduled. If a business process is rescheduled, the business process is moved forward in the queue. This value can be different for each queue. The .# indicates the queue number.</p>

4. In the `noapp.properties` file, under `# QUEUE CONFIGURATION`, Queue # - `FairShareSchedulingPolicy`, where # is queue number, locate the following properties and change the value for each as appropriate. Complete this step for each queue you are using.

Property	Description
<code>ResourceAllocation.#</code>	<p>Amount of resources to allocate to this specific queue for fair share scheduling. This ensures that queue # has more resources for processing, which improves processing times and ensures efficient use of system resources.</p> <p>Note: Set the resource allocations according to the percentages of use for each queue and the importance of processing completed on each queue. If you have higher priority items on one queue, increase the resource allocation to that queue to increase processing capability. Apply lower percentages of resources to lower priority or less used queues to keep resources free for higher priority processing.</p> <p>This value can be different for each queue. The <code>.#</code> indicates the queue number.</p>
<code>InitialCycles.#</code>	<p>Number of cycles to execute the first time a business process gets to execute. Normally, this value should be 5 or less. It is intended to facilitate business processes (particularly web services) for which normal processing is quite short, but for which there is longer processing in special cases. It also allows a mix of short and long business processes in a queue, favoring the shorter ones. This value can be different for each queue. The <code>.#</code> indicates the queue number.</p>
<code>StealThreads.#</code>	<p>Enables or disables the ability to steal threads from other queues. Valid values:</p> <ul style="list-style-type: none"> ◆ <code>true</code> – Enables the ability to steal threads ◆ <code>false</code> – Disables the ability to steal threads <p>Queues that are configured to steal threads from other queues (StealThreads property set to <code>true</code>) only steal from queues that have the AllowStealing property set to <code>true</code>. This value can be different for each queue. The <code>.#</code> indicates the queue number.</p>
<code>AllowStealing.#</code>	<p>Enables or disables the ability of other queues to steal threads from this queue. Valid values:</p> <ul style="list-style-type: none"> ◆ <code>true</code> – Allow other queues to steal threads ◆ <code>false</code> – Do not allow other queues to steal threads <p>Queues that are configured to steal threads from other queues (StealThreads property set to <code>true</code>) only steal from queues that have the AllowStealing property set to <code>true</code>. This value can be different for each queue. The <code>.#</code> indicates the queue number.</p>

5. Save the `noapp.properties` file under the same name in the `install_dir/properties` directory.

Caution: Saving the `noapp.properties` file under a different name or in a different location will cause the changes not to be applied to Gentran Integration Suite.

6. Stop Gentran Integration Suite.
7. In the `install_dir/bin` directory, complete the following action depending on your operating system:
- ◆ UNIX or Linux – Run `setupfiles.sh`
 - ◆ Windows – Run `setupfiles.cmd`
8. Start Gentran Integration Suite.

The changes are applied to Gentran Integration Suite and override the changes made in the *install_dir/properties/tuning.properties* file.

Monitoring Queues using queueWatcher

The queueWatcher is a troubleshooting tool in Gentran Integration Suite. The queueWatcher performs the following functions:

- ◆ Monitors Gentran Integration Suite queues
- ◆ Checks node workload information
- ◆ Provides details on active or waiting threads

Enabling queueWatcher

To enable queueWatcher, do the following:

1. Navigate to the *<install_dir>/properties* directory.
2. Add the following line in the *noapp.properties* file:

```
queueWatcher=true
```
3. Add the following line in the *customer_overrides.properties* file:

```
noapp.queueWatcher=true
```
4. Restart Gentran Integration Suite to activate queueWatcher.

Note: If Gentran Integration Suite is running on multiple nodes (clustered environment), enable queueWatcher on all nodes.

Accessing queueWatcher

To access queueWatcher, do the following:

1. Open your web browser to *http://Host:Port/queueWatcher*, where *host:port* is the IP address and port number where Gentran Integration Suite resides on your system. The queueWatcher page appears.

2. The queueWatcher page displays the following information:

Heading	Description
Monitor Threads for All Queues	<p>Shows the current activities on various queues to determine the queue sizing. Click to view more details.</p> <ul style="list-style-type: none"> ◆ QueueName – Shows the queue name. ◆ Min – Minimum number of threads available for the queue. The threads will be honored even if they are higher than MaxThreads (global maximum queue threads). The minimum number of threads cannot be higher than the maximum number. The fairness calculation does not apply for minimum threads. ◆ Used – Number of business processes currently running in a thread. ◆ Calc – Fairshare thread calculation for the queue. Fairshare is based on concurrent activities on all queues and is dynamically updated. ◆ Pool – Number of threads in a queue's pool. The threads timeout if they are not used. ◆ Max – Maximum number of threads used by the queue. ◆ QueueDepth – Number of business processes waiting for a thread in the queue. ◆ List of Working Threads – List of business processes currently running on a thread.
Wait Queue	Shows the workflows listed on individual queues.
Queue_1 – Queue_9	Shows running and waiting (for available thread) business processes.
View Content Cache Entries	<p>Shows the content cache entries.</p> <ul style="list-style-type: none"> ◆ Soft Reference Cache Slots in use – The workflow context is saved into this queue and can be recovered from it. If required, the garbage collector can acquire more heap space from this queue. The workflow contexts are not serialized on this queue. ◆ In Memory Cache Bytes in use – This memory cache holds the workflow contexts with a size lesser than the configured threshold if it has space. The workflow contexts are serialized on this queue. ◆ Disk Cache Bytes in use – This cache holds workflow contexts larger than the defined threshold. The workflow contexts are serialized on this queue.
View the list of Work Flow IDs that recover would see in the queue	Shows the workflow ID when it is run or moved to another node in the cluster.
View System Information, such as a list of threads and JVM settings	Shows the default values of the JVM properties and current threads in the JVM.
View Cluster Multicast Data	Shows the Gentran Integration Suite view of other nodes in the cluster.
View Scheduling Policy Statistics (XML)	Shows an XML document containing information covered by other queues in queueWatcher.
Pause All Queues	Pauses all the queues. You cannot pause individual queues.
Restart All Queues	Restarts all queues. You cannot restart individual queues.

Heading	Description
Worker Threads (ASI and HAR – Interrupt will stop the thread !!)	Shows a list of threads registered as ASI (when acquired by the dispatcher) or HAR (when it starts running).
View Memory	Shows heap usage in the system. More business processes can run if heap space and CPU resources are available.
List Manager Properties	Shows the list of properties from noapp.properties file.
List Default Queue Configuration Parm	Shows the parameters set for all the nine queues.
List Active Queue Configuration Parm	Shows the current queue configuration.
List DB Pool Information	Shows the active view of database pools. It is dynamically updated. <ul style="list-style-type: none"> ◆ CurrentSize – The number of current connections to the database from this pool. ◆ MaximumSize – The number of connections a pool can acquire. This value is made of two values – maximum size and buffered size. ◆ PoolRequests – The number of requests received by the pool. ◆ WaitRequests – The number of requests waiting for a connection. ◆ BufferRequests – The number of requests exceeding maximum size. ◆ DeleteRequests – The number of available connections. ◆ BadItemCount – The number of bad connections returned
Config Queue	Configure the queue parameters to tune performance. The parameters are not persisted and reset when Gentran Integration Suite restarts.
Reset Queue	Resets the queue to default values.
Step Monitor	Shows the list of business processes and workflow contexts in the queue.

Compressing Cache Contents

You can compress cache contents from queues before the cache is written to disk. Compressing this information increases performance, by requiring fewer bytes to be written to and retrieved from storage. After the cache contents are compressed and written to disk, the contents are decompressed after the information is read back from the disk.

To compress cache contents before the cache is written to disk:

1. Open the *install_dir*/properties/workflows.properties file in a text editor.
2. In the workflows.properties file, locate the **compressObj** property and change the value to **true**. The default value is **false**.
3. Save the workflows.properties file under the same name in the *install_dir*/properties directory.
4. Stop Gentran Integration Suite.

Caution: Saving the `workflows.properties` file under a different name or in a different location will cause the changes not to be applied to Gentran Integration Suite.

5. In the `install_dir/bin` directory complete one of the following actions depending on your operating system:
 - ◆ UNIX or Linux – Run `setupfiles.sh`
 - ◆ Windows – Run `setupfiles.cmd`
6. Start Gentran Integration Suite.

The changes are applied to Gentran Integration Suite and the cache contents are compressed before being written to disk and decompressed after being read from the disk.

Setting the Document Body Serialization Threshold Property

When the size of the documents increase, there is a significant increase in the use of resources for serialization/deserialization of the business process context. When you install Gentran Integration Suite, the default value is `documentInlineSerializationThreshold = 102400` bytes.

If the business process requires the use of the body at almost every step, the number of database reads may use more resources than serialization/deserialization. In this case, set the value higher. If, however, the documents are large and used infrequently, set the value lower.

To increase performance, in the `install_dir/properties/jdbc.properties.in` file, complete the following actions:

1. In the `install_dir/properties/jdbc.properties.in` file, add the following line:


```
stream_threshold_bytes = &tune.noapp.documentInlineSerializationThreshold;
```
2. Set the value of `tune.noapp.documentInlineSerializationThreshold` in the `install_dir/properties/tuning.properties` file according to your document size.
3. From the **Administration** menu, select **Operations > System > Troubleshooter**.
4. On the System Troubleshooting page, click **Stop the System**.

Stop the System runs the `softstop` script allowing all business processes in progress to complete before stopping the system.
5. Run the `install_dir/bin/setupfiles` script. The performance configuration changes populate the initial setup files that are used during startup.
6. In the `install_dir/bin` directory, start Gentran Integration Suite by running one of the following scripts depending on your operating system:
 - ◆ UNIX or Linux – Run `run.sh`
 - ◆ Windows – Run `startWindowsService.cmd` or click the icon.

Gentran Integration Suite starts and the updated performance configuration settings are applied.

Allocating JNDI Contexts

Several components of Gentran Integration Suite use Java Naming and Directory Interface (JNDI) to locate objects. You may find that you need to change the allocation of JNDI contexts in Gentran Integration Suite to enhance performance. In Gentran Integration Suite, the default value is 50.

To change the JNDI allocation:

1. From the **Administration** menu, select **Operations > System > Troubleshooter**.
2. On the System Troubleshooting page, click **Stop the System**.
Stop the System runs the softstop script allowing all business processes in progress to complete before stopping the system.
3. In the *install_dir/properties* directory, open the tuning.properties file in a text editor.
4. Locate the **tune.noapp.jdni.contextpoolsize** property, and change the value to the level you need.
5. In the *install_dir/bin* directory, complete one of the following actions depending on your operating system:
 - ◆ UNIX or Linux – Run setupfiles.sh
 - ◆ Windows – Run setupfiles.cmdThe changes populate the initial setup files that are used during startup.
6. In the *install_dir/bin* directory, start Gentran Integration Suite by running one of the following actions depending on your operating system:
 - ◆ UNIX or Linux – Run run.sh
 - ◆ Windows – Run startWindowsService.cmd or click the icon.Gentran Integration Suite starts and the updated performance configuration settings are applied.

Changing Persistence Levels

Gentran Integration Suite can persist all data or only a small amount of data for each activity processed. *Persistence level* means the level of detail written to the database as the business process runs. The default value for Gentran Integration Suite is Full persistence. You may find that you need to change the persistence level to a lower level to enhance performance.

When you design your business processes review your persistence requirements and set persistence levels appropriately to reduce the chances of your database filling up and slowing your processing and system. You may have some business processes that you want to persist all steps and data, and other business processes that you do not need to persist any steps or data.

There are three levels at which you can set persistence levels:

- ◆ Global – by changing the value of the **persistence_level** property in the *install_dir/properties/noapp.properties* file. It is more efficient to set persistence levels at either the

business process or activity level. The default value is **PERSISTENCE_FULL**. You cannot set the value to **PERSISTENCE_WF_NONE** or **PERSISTENCE_ERROR_ONLY**.

- ◆ Business process – by changing the persistence level during the check-in process. The business process persistence level overrides the global persistence level.
- ◆ Activity level – by changing the persistence level in the service or adapter. The activity persistence level overrides the business process persistence level and the global persistence level.

Note: Changing the persistence level at the Activity level is rarely necessary. If you need to change the persistence level of an individual service or adapter, contact Sterling Commerce Customer Support for assistance.

Changing Persistence at the Global Level

To edit the global persistence levels:

1. Open the *install_dir*/properties/noapp.properties file in a text editor.
2. In the property file, locate the **persistence_level** parameter, and change the value to the persistence level you want. Possible values are:
 - ◆ **PERSISTENCE_FULL** – Saves a complete copy of process data along each step of the process.
 - ◆ **PERSISTENCE_MINIMAL** – Saves all steps in a business process and selected copies process data.
 - ◆ **PERSISTENCE_NONE** – Saves the first and last steps of a business process, any steps with an override persistence level, and no copies of business process data.

Note: Use **PERSISTENCE_NONE** for day-to-day processing.

3. Stop Gentran Integration Suite.
4. In the *install_dir*/bin directory, complete one of the following actions depending on your operating system:
 - ◆ UNIX or Linux – Run setupfiles.sh
 - ◆ Windows – Run setupfiles.cmd
5. Start Gentran Integration Suite.

Changing Persistence at the Business Process Level

To edit the persistence level for a business process:

1. From the **Administration** menu, select **Business Processes > Manager**.
2. Type all or part of the name for the business process you want to edit into the **Search** box and click **Go!**
3. Click **source manager** next to the business process you want to edit.
4. Click **edit** next to the business process to lock the business process and access the Editor.
5. On the BPML Specification page, type a new description that helps you identify this version of the business process.

6. Click **Next**.
7. On the Process Levels page, indicate the level of persistence you want for the business process:
 - ◆ Full - Retain all data for this business process, including associated documents, activities, and all process data.
Note: This choice uses the most system resources of any persistence setting.
 - ◆ Step Status (Engine May Override) - Retain process data and documents for services configured to always persist data for this business process. Gentran Integration Suite may persist additional data that is requested by the process engine or service parameters.
 This setting is recommended for most processes since it helps optimize performance and provides a level of process recovery suitable for most business needs.
 - ◆ BP Start Stop (Engine May Override) - Retain additional data as requested by the process engine or service parameters.
 - ◆ System Default - Assume that for the data, configuration is already defined in Gentran Integration Suite to retain data.
 - ◆ Step Status Only - Every step in the business process runs with minimal persistence, even if a service or adapter is preconfigured for full persistence. Gentran Integration Suite does not persist additional data.
 Choose this setting if you would like to persist only status information for each step; you are certain your process will execute successfully without persisted data; and your recovery needs do not depend on step data being saved.
 - ◆ BP Start Stop Only - Retains only error information when you are certain your process will execute successfully without persisted data, and when your recovery needs do not depend on process step data being saved.
 - ◆ Zero - Do not retain any data.
 - ◆ Error Only - Retain process information only when an error occurs. If nothing is wrong and streamed document data is not used, nothing is retained.
8. Click **Next**.
9. On the Deadline Settings page, click **Next**.
10. On the Life Span page, click **Next**.
11. On the Default Version page, click **Next**.
12. On the Confirm page, click **Finish**.

Changing System Logging

Gentran Integration Suite logs various types of messages to the log files on disk. The default is to log errors only, nothing else. For more information about system logs, see *Managing System Logs* on page 5.

If you have your system logging level set to **ALL**, you may find that you do not need all of the logging messages that are collected, because of resource limitations.

If you have your system logging level set to **ERROR**, you may find that you need additional messages from more detailed logging.

For information on how to change system logging levels, see *Changing Log Settings* on page 7.

Enabling Trusted Domains for Schemas

Gentran Integration Suite has DTDs and schemas in a repository for different XML components. When a DTD or schema is required, Gentran Integration Suite searches the database for the appropriate file. If Gentran Integration Suite cannot find the appropriate file, you can configure the `install_dir/properties/noapp.properties` file to have Gentran Integration Suite go to trusted Web sites and download the appropriate DTD or schema.

To configure Gentran Integration Suite to visit trusted Web sites for DTD and schema data:

1. In the `install_dir/properties` directory, locate the `noapp.properties.in` file.
2. Locate the `schemaResolver.trustedDomain.#` property and type the URL for the Web site. For example, `schemaResolver.trustedDomain.1 = http://www.gdsregistry.org`.
3. Save the `noapp.properties.in` file under the same name and in the same location.
4. Stop Gentran Integration Suite.
5. In the `install_dir/bin` directory, complete one of the following actions depending on your operating system:
 - ◆ UNIX or Linux – Run `setupfiles.sh`
 - ◆ Windows – Run `setupfiles.cmd`
6. Start Gentran Integration Suite.

The changes are applied and you can now download DTD and schema data from the trusted Web site.

Enabling the Console Listener

You can enable the console listener to append the events identified in the events viewer to the `noapp.log` during processing. If you do not enable the console listener, the events are not appended to the `noapp.log` file and, consequently, are not captured in one location for future use. You would need to view each component log file individually to see the events that are logged.

To enable the console listener:

1. In the `install_dir/properties` directory, locate the `listenerStartup.properties.in` file.
2. Locate the `#Listener.Class.1=com.sterlingcommerce.woodstock.event.listeners.console.ConsoleListener` line.

3. Delete the # symbol in front of the preceding line.
4. Save the listenerStartup.properties.in file under the same name and in the same location.
5. Stop Gentran Integration Suite.
6. In the *install_dir/bin* directory, complete one of the following actions depending on your operating system:
 - ◆ UNIX or Linux – Run setupfiles.sh
 - ◆ Windows – Run setupfiles.cmd
7. Start Gentran Integration Suite.

The changes are applied and you can now view the events as they are appended to the noapp.log file during processing.

Enabling the JMX Event Listener

You can enable the JMX (Java Management Extensions) Event Listener to capture Sterling Commerce platform events and do necessary notification and handling based on your custom JMX event handling requirements. You can plug in standard JMX tools to monitor Sterling Commerce platform and its application components.

If you have an existing infrastructure for your enterprise applications to do JMX-based management and event monitoring, you can have a uniform way established to handle event notifications from Gentran Integration Suite/Sterling Platform along with your other applications through your JMX framework using the JMX console/tools.

To enable the JMX Event Listener:

1. In the *install_dir/properties* directory, locate the jmx.properties file.
2. Set the jmx.start and jmx.event properties to **true**.
If necessary, restart the Sterling Commerce platform to pick up these changes.
3. In the *install_dir/properties* directory, locate noapp.properties file.
4. Set the startup.class4 property to **com.sterlingcommerce.woodstock.jmx.event.SPEventJMXAgentStartup**.

This will start the JMX Agent when the Sterling Commerce platform starts.

To see the MBeans configured in the Sterling Commerce platform, go to <http://host:port> where the port is the HTMLAdapterServer port configured in the jmx.properties file.

5. Configure a business process to send an event at startup, completion or another time.
6. At the Sterling Commerce platform JMX console, you can view platform events by looking at the MBean named “name=eventMonitor” in the EventAgent domain. This shows just the most recent event that was received.

Using external JMX tools that support remote JMX connectivity, it is also possible to connect remotely to the Event Listener JMX agent. The JMX agent will be started with a remote connector listening on

the `jmx.rmi.port` port as specified in `jmx.properties`. The URL for remote connection will take the form: `/jndi/rmi://[HOST]:[jmx.rmi.port]/events`. Use this URL to connect to the JMX agent and receive all system wide events.”

It is possible to listen to select events, rather than all of them. By adding a (comma separated) list of event names to the `jmx.eventListener.forward` property, the JMX agent will restrict JMX notifications to only those events specified in the list.”

Enhancing EDI Performance

You can boost EDI performance in Gentran Integration Suite by adjusting the tracking level in the `enveloping.properties` file. This performance boost is done at the expense of Tracking and Search functionality. The tracking level setting affects the following EDI functionality:

- ◆ EDI Correlation Search
- ◆ EDI Document Tracking
- ◆ EDI Reporting

To modify the tracking level:

1. In the `install_dir/properties` directory, locate the `enveloping.properties.in` file.
2. Locate the **TRACKING_LEVEL=*setting*** line.
3. Change the *setting* value to achieve the desired EDI performance and functionality as shown in the following table:

setting	Description
none	Provides the largest EDI performance boost with the least tracking and search functionality. EDI Correlation Search, EDI Document Tracking and EDI Reporting are nonfunctional.
basic	Provides an EDI performance boost while also providing search functionality. EDI Correlation Search is functional. EDI Document Tracking is nonfunctional. EDI Reporting is partially functional; you can use the EDI Outbound Acknowledgement report type.
full	Default setting. Provides the lowest EDI performance with the highest search and tracking functionality. EDI Correlation Search, EDI Document Tracking and EDI Reporting are fully functional.

4. Save the `enveloping.properties.in` file under the same name and in the same location.
5. Stop Gentran Integration Suite.
6. In the `install_dir/bin` directory, complete one of the following actions depending on your operating system:
 - ◆ UNIX or Linux – Run `setupfiles.sh`
 - ◆ Windows – Run `setupfiles.cmd`

7. Start Gentran Integration Suite.

The changes are applied to the `enveloping.properties` file and the new settings will determine the default EDI performance and functionality in your business processes.

The default EDI performance settings set in the `enveloping.properties` file can be overridden for certain EDI services by using a **TRACKING_LEVEL** parameter. This allows you to obtain maximum EDI performance in some business processes and maximum search and tracking functionality in others. See the *Services and Adapters* topic for more information about EDI services.

Note: Document tracking is turned off by default in the system-defined EDI business processes. If you define an EDI business process and turn Document Tracking on, that will override the **TRACKING_LEVEL** settings in both the `enveloping.properties` file and the EDI service parameter.

Managing Property Files

The property files in Gentran Integration Suite help you set up, customize and diagnose problems with the installation and operation of Gentran Integration Suite. You can maintain your customized property file settings with a customer override property file (`customer_overrides.properties`) that overrides default property settings in the property files. The customer override property file is not changed during the installation of Gentran Integration Suite upgrades or patches. To prevent having your customized settings overwritten, you should use the customer override property file whenever possible rather than directly editing the Gentran Integration Suite property files or the associated `.in` files.

Caution: Property files directly affect the operation of Gentran Integration Suite. Contact Sterling Commerce Customer Support before changing any properties.

For information about specific property files, including the `customer_overrides.properties` file, refer to the *Property Files* documentation.

Managing System Recovery

Gentran Integration Suite is designed to automatically reset itself to a previous, stable state following a system crash. After the reset is complete, the system is up and operational.

After a system crash, restore the server where Gentran Integration Suite is installed to a stable state and restart Gentran Integration Suite. After Gentran Integration Suite has been restarted, the system invokes the Recovery business process to search for and requeue any unprocessed transactions. If the database goes down, Gentran Integration Suite stops pulling jobs from queues. The system recovers these jobs when the database becomes available.

Gentran Integration Suite now uses multi-threaded operation to restart or resume business processes following a system crash. Unfinished business processes, up to a pre-configured number, are stored in a pool. A pre-configured number of threads each receive a batch of business processes to restart or resume.

The size of the batch is also pre-configured. When a thread finishes restarting or resuming a batch of business processes, it requests another batch from the pool. This cycle continues until the pool is empty.

By default, the system uses 10 threads that each resume or restart 10 business processes and then request and process another 10 up to a total of 1000 business processes for all threads. You can configure the system settings to adjust recovery performance to your particular needs.

To configure the system recovery process settings, perform the following actions:

1. In the *install_dir/properties* directory, locate (or create, if necessary) the *customer_overrides.properties* file.
2. Open the *customer_overrides.properties* file in a text editor.
3. Add override statements as needed using the following format for each new line:

```
bprecovery.property=new_value
```

The applicable properties are shown in the following chart:

Property	Description
maxAutorecoveryCount	<p>Maximum number of business processes that will be resumed or restarted by the Recovery business process each time it runs. Default is 1000.</p> <p>This defines the maximum size of the pool used to hold the business processes that need to be resumed or restarted. Each time a thread requests a batch of business processes to restart or resume, the pool size is reduced by the size of the batch (specified by batchSize).</p>
numberOfThreads	<p>Number of threads that will be used simultaneously to restart or resume business processes. Default is 10.</p> <p>Each thread will start or resume a batch of business processes. The size of the batch is specified by batchSize. When the thread is finished with the batch, it will request another batch from the pool. This process will continue until the pool of business processes that need to be restarted or resumed is empty.</p> <p>If you frequently have very large numbers of business processes running at the same time, a larger numberOfThreads value will enhance recovery performance.</p> <p>Note: Increasing the number of threads will speed up the recovery process, but will use more system resources. Decreasing the number of threads will free up system resources, but will slow the recovery process. Adjust this value based upon your business needs.</p>
batchSize	<p>Maximum number of business processes that will be resumed or restarted by each thread before it requests more business processes from the pool. Default is 10.</p> <p>Note: Increasing the batch size will speed up the recovery process, but will use more system resources. Decreasing the batch size will free up system resources, but will slow the recovery process. Adjust this value based upon your business needs.</p>

For example, assume that you want to change the number of threads from the default value of 10 to 20 and the batch size from 10 to 5. To do so, add the following lines to the *customer_overrides.properties* file:

```
bprecovery.numberOfThreads=20
```

```
bprecovery.batchSize=5
```

4. Save and close the `customer_overrides.properties` file.
5. Stop and restart Gentran Integration Suite to use the new values.

For assistance in determining the proper settings for your business needs, contact Sterling Commerce Customer Support.

Changing the SoftStop Time Limit

By default, the time limit for a softstop of Gentran Integration Suite is 5 minutes. At the end of the time limit, the system performs a hardstop that immediately shuts down Gentran Integration Suite. You can change the time limit for a softstop by overriding the default value.

To override the default softstop time limit:

1. In the `install_dir/properties` directory, locate (or create, if necessary) the `customer_overrides.properties` file.
2. Open the `customer_overrides.properties` file in a text editor.
3. Add the following line:

```
noapp.shutdown_timeout=new_value
```

`new_value` - The new time limit, in seconds, for a softstop of Gentran Integration Suite.

For example, assume that you want to change the softstop time limit from the default value of 300 seconds (5 minutes) to 600 seconds (10 minutes). To do so, add the following line to the `customer_overrides.properties` file:

```
noapp.shutdown_timeout=600
```

4. Save and close the `customer_overrides.properties` file.
5. Stop and restart Gentran Integration Suite to use the new values.

Changing Advanced File Transfer Settings

Advanced File Transfer (AFT) Visibility performance can be tuned by overriding three properties in the `visibility.properties` file. The following settings can be adjusted:

- ◆ The looping interval of the AFT Worker thread
- ◆ The capacity of the AFT database listener event queue
- ◆ The amount of AFT coverage to enable

Note: For large volume operations of FTP, Mailbox, Advanced File Transfer, and File Broker, the number of open file descriptors should be set to at least 4096, using the command `ulimit -n 4096` (recommended value is unlimited).

The AFT Worker thread stores AFT events in the database. It consumes all current events in an input event queue, and then sleeps for the specified batching interval. This allows events to build up so that they can be

more efficiently batched into the database. You can improve AFT visibility performance by adjusting the batching interval. The batching interval controls how much latency is programmed into the backend in order to batch up the events.

You can also specify the size of the event input queue. By default, the AFT backend will store up to 2048 AFT events waiting for the AFT Database worker thread to return and process them. If the thread does not return before Gentran Integration Suite fills up the event queue, the event queue will block client threads until space is available in the queue. The clients cannot continue with their tasks (for example, copying a file) until the AFT Database worker thread empties the queue. If the AFT Database worker thread returns and the event queue is full to capacity, a message will be written to the visibility log file as shown in the following example:

```
DmiVisEventWarehouse.Worker - ALERT ISSUED - AFT Warehouse input BoundedQueue has reached capacity! Client threads are now blocking! The AFT Warehouse will temporarily suspend. The sleep interval and work full time to handle this spike in event activity.
```

At this point, the Worker thread will suspend and sleep operations and service the event queue. The Worker thread will monitor the number of events it is retrieving from the input queue. When the queue is no longer full to capacity as the Worker thread is removing events for processing, the Worker thread will resume normal sleep intervals after the current events are processed and write a message into the visibility log file as shown in the following example:

```
DmiVisEventWarehouse.Worker - ALERT RESCINDED - AFT Warehouse input BoundedQueue is operating below capacity! The normal batch sleep interval is now re-established.
```

By adjusting the size of the event queue, you can improve AFT Visibility performance.

Lastly, you can improve performance by adjusting the amount of AFT Visibility coverage that is enabled. You can specify settings ranging from None, which disables AFT Visibility, to All, which enables all AFT Visibility components. Specifying the proper settings for your needs allows you to disable unneeded components to improve performance, while preserving core file transfer and data flow capabilities.

To specify AFT Visibility settings, perform the following steps:

1. Change to the *install_dir/properties* directory and locate (or create, if necessary) the *customer_overrides.properties* file. For more information, see *Managing Property Files* on page 74.
2. Open the *customer_overrides.properties* file in a text editor.

3. Add the property settings that you want to use from the following table, using the format:

`dmvisibility.PROPERTY_NAME=PROPERTY_VALUE`

PROPERTY_NAME - The name of the property as shown in the Property column of the table.

PROPERTY_VALUE - The value you want to assign to the property.

For example, assume that you want to change the persistent batching interval from the default value of 2000 to a value of 3000. To do so, add the following line to the `customer_overrides.properties` file:

`dmvisibility.persistent_batching_interval=3000`

Property	Description
<code>persistent_batching_interval</code>	Specifies the interval, in milliseconds, between batching AFT events to the database. Default is 2000 milliseconds (2 seconds). Example: <code>persistent_batching_interval=2000</code>
<code>event_input_queue_capacity</code>	Specifies the capacity of the bounded queue (AFT database listener event queue) through which all AFT events must pass enroute to the database. Default is 2048. Example: <code>event_input_queue_capacity=2048</code> Note: If the value of this property is set too high, it may negatively impact system performance due to excessive memory usage. This property should be tuned with the persistent_batching_interval property for optimum results. Note: For large volume operations of FTP, Mailbox, Advanced File Transfer, and File Broker, the number of open file descriptors should be set to at least 4096, using the command <code>ulimit -n 4096</code> (recommended value is unlimited).
<code>visibility_coverage</code>	Specifies which AFT Visibility components are enabled. Valid values: <ul style="list-style-type: none"> ◆ All – Enable all AFT Visibility components. This is a summation of CommBase, Authentication, Authorization, NonFileXfer and AdminAudit. ◆ CommBase – Enable the basic level of visibility coverage that includes Communication Sessions, File Transfers and Process File Events. AFT Routing events and records are also included in this setting. ◆ Authentication – Enable Authentication events and records. ◆ Authorization – Enable Authorization events and records. ◆ NonFileXfer – Enable Non File Transfer events and records (such as FTP CWD). ◆ CommAll – Enable CommBase, Authorization, Authentication, and NonFileXfer. ◆ AdminAudit – Enable Administrative Audit trail events and records. ◆ None – Disable AFT Visibility. Example: <code>visibility_coverage=All</code> Visibility coverage property settings can be also combined using commas as delimiters. For example: <code>visibility_coverage=CommBase,Authentication,AdminAudit</code>

4. Save and close the `customer_overrides.properties` file.

5. Stop and restart Gentran Integration Suite to use the new values.

Working with Large BLOB Data

Data stored in BLOBs (binary large objects) is sometimes too large to include in a XAPI (XML API) output XML document. Examples of such data are large documents and map objects. To ensure that the amount of data in the output XML from the XAPI is of reasonable size (for memory and performance reasons), excessively large BLOB data is requested “out-of-band.”

This data is retrieved through a streaming process that uses the key/ID from the output XML to access the BLOB in a request that is separate from the main XAPI request. The streaming process identifies the large data as a Document object (instead of just raw data), which allows the data to be stored in either the database or file system, and also be encrypted.

To determine when to use the out-of-band method, the Gentran Integration Suite enables you to set a size threshold value for BLOBs. The default threshold value for a BLOB is 1 MB. Setting this value to a different value can help customers who have certain average size documents and need to view them or download them using the user interface.

To control this threshold value, you can:

- ◆ Initially set the value by adding the `outOfBandThreshold` property to the `xapi.property.in` file.
- ◆ Modify the value by modifying the `outOfBandThreshold` property in the `xapi.property.in` file.

If you try to get a BLOB that is over the threshold value, a token is assigned. This token is a pointer into the `DATA_TABLE`. For documents, this token is the `doc_id`. All entity-based XAPIs support this functionality. For base resources, only the `getDocumentDetails` XAPI currently supports this. Other base resource XAPIs will need to specifically implement this for it to be enforced.

Large BLOB data also affects how Gentran Integration Suite databases access BLOB data. Gentran Integration Suite databases can use BLOB data, but cannot store it. Instead, the databases access the BLOB data from the `DATA_TABLE` table(s) that contain the actual BLOB data.

Resolving Slow System Issues

You may find that your Gentran Integration Suite begins working slowly for a variety of reasons. An increase in volume or improperly designed business processes may cause slower than expected performance.

This section covers the following topics:

- ◆ *Symptoms, Causes, and Resolutions of a Slow System* on page 80
- ◆ *Conducting a Thread Dump* on page 82

Symptoms, Causes, and Resolutions of a Slow System

Symptoms of a slow system may include:

- ◆ Slow transaction processing
- ◆ Unusually long delays in starting a business process
- ◆ Unusually large numbers of business processes waiting in the queues
- ◆ Slow interface log in
- ◆ Slow business process searches
- ◆ Slow system services such as indexing and purging
- ◆ Slow interface

Potential Causes of a Slow System

Determining the cause of a slow system can require you to investigate many areas of the system and how you are implementing Gentran Integration Suite.

Potential causes of a slow system include:

- ◆ Improperly designed business processes. For more information, see the *Managing Business Processes* topic.
- ◆ Various system services being turned off. For more information, see *Service Disabled* on page 81.

- ◆ Improperly tuned performance properties. For more information, see *Improperly Tuned Performance Properties* on page 82.

Determining the Cause of a Slow System

To determine the cause of a slow system, you will complete some, if not all, of the following tasks, depending on your situation:

- ◆ Review your operating system and database performance if other applications are using these resources. If other applications using the same resources are slow, this indicates that the performance issue is not specific to Gentran Integration Suite. Review your operating system or database applications for performance issues.
- ◆ Review the log files for errors and processing information. If your system is slow, good places to start troubleshooting are the system.log and the noapp.log in the *install_dir/logs* directory.
For more information, see *Managing System Logs* on page 5 and *Reviewing System Information* on page 18.
- ◆ Review the Gentran Integration Suite Internal System Statistics report from the **Operations > System > Performance > Statistics** page.
This report shows the number of invocations and the processing time for each business process and system activity. You can use this information to identify activities that are slowing down and impacting performance.
For more information, see *Managing Performance Statistics* on page 33.
- ◆ Conduct a thread dump to identify business process bottlenecks in Gentran Integration Suite.
For more information, see *Conducting a Thread Dump* on page 82.

Slow System Resolutions

There are many causes and resolutions for a slow system. Some of the most common are discussed below.

Improperly Designed Business Processes

Improperly designed business processes can cause performance issues with Gentran Integration Suite. These issues can be averted by implementing well-designed business processes.

Service Disabled

If a service that Gentran Integration Suite uses for processing is disabled, the following issues may arise:

- ◆ Business processes may start to fail or wait until the service becomes available.
- ◆ Indexing, archiving, or purging activities may not occur causing the database to fill up.

Both of these issues cause your system to run slowly.

For example, if the index service is disabled, there is no way for the data in the database to be flagged for archival or purging. If this happens, archival and purging activities may run, but no data is archived or purged. This results in the database storing too much information and slowing your system.

To resolve this issue, start the disabled service.

Improperly Tuned Performance Properties

Gentran Integration Suite may run slowly if you have not tuned the system performance using the Performance Tuning Utility or manually tuned the system performance in the *install_dir/properties/noapp.properties* file.

You may need to make changes to the appropriate performance tuning properties either manually or with the Performance Tuning Utility. For more information, see *Performance Tuning Utility* on page 45, or *Manual Performance Tuning* on page 53.

Conducting a Thread Dump

A thread dump is a snapshot of every thread that is running in Gentran Integration Suite at the time the thread dump is generated. The thread dump shows exactly what each thread is doing at the time the report is captured. You can conduct a thread dump to determine bottlenecks in Gentran Integration Suite. Typically, problem threads have a state of **waiting**, while threads with a state of **runnable** or **R** indicate a healthy thread.

For information about reading a thread dump, see *Reading a Thread Dump* on page 87.

Conducting a Thread Dump in UNIX (HP-UX or Sun Solaris)

To conduct a thread dump in UNIX (HP-UX or Sun Solaris):

1. Access the computer on which Gentran Integration Suite is installed.
2. Change directories to the *install_dir*.
3. On the command line, type **cat noapp.pid**.
This lists the parent thread process ID.
4. Type **ps -ef | grep noapp.pid**.
This returns all child process IDs associated to the parent process ID.
5. Type **kill -QUIT noapp.pid child.pid**.
You must include both the *parent pid* and the *child pid*.
6. The thread dump is placed in the *noapp.log* file in the *install_dir/logs* directory.

Conducting a Thread Dump in UNIX (IBM AIX)

To conduct a thread dump in UNIX (IBM AIX):

1. Access the computer on which Gentran Integration Suite is installed.
2. Change directories to the *install_dir/bin* directory.
3. Open the *tmp.sh* file in a text editor.
4. Comment the following lines in the *tmp.sh* file by placing a **#** at the beginning of the line.

- ◆ `DISABLE_JAVADUMP=true`
- ◆ `IBM_NOSIGHANDLER=true`
- ◆ `export DISABLE_JAVADUMP`
- ◆ `export IBM_NOSIGHANDLER`

This is required to enable java dumps.

5. Close and save the *install_dir/bin/tmp.sh* file.
6. Change directories to the *install_dir*.
7. On the command line, type **cat noapp.pid**.
This lists the parent thread process ID.
8. Type **ps -ef | grep noapp.pid**.
This returns all child process IDs associated to the parent process ID.
9. Type **kill -QUIT noapp.pid child.pid**.
You must include both the *parent pid* and the *child pid*.
10. The thread dump is placed in the *javacore*.txt* file.
11. When you are finished conducting thread dumps, change to the *install_dir/bin* directory.
12. Open the *tmp.sh* file in a text editor.
13. Uncomment the following lines in *tmp.sh* by removing the `#`.
 - ◆ `#DISABLE_JAVADUMP=true`
 - ◆ `#IBM_NOSIGHANDLER=true`
 - ◆ `#export DISABLE_JAVADUMP`
 - ◆ `#export IBM_NOSIGHANDLER`
14. Close and save the *install_dir/bin/tmp.sh* file.

Conducting a Thread Dump in Linux

To conduct a thread dump in Linux:

1. Access the computer on which Gentran Integration Suite is installed.
2. Change directories to the *install_dir*.
3. On the command line, type **killall -QUIT java**.

This thread dump is completed.

Depending on your JVM, the thread dump is placed in one of the following locations:

- ◆ *noapp.log* in the *install_dir/logs* directory
- ◆ *javacore*.txt*

Conducting a Thread Dump in Windows

Note: Before conducting a thread dump in Windows, you must contact Sterling Commerce Customer Support to obtain the following three files:

- ◆ debug_ops.cmd
- ◆ debug_noapp.cmd
- ◆ debug_passphrase.cmd

To conduct a thread dump in Windows:

1. Access the computer on which Gentran Integration Suite is installed.
2. Open the file debug_ops.cmd in a text editor.
3. Replace all occurrences of the string “D:\SterlingCommerce31\mysql4” with the correct path to your Gentran Integration Suite installation. For example, if Gentran Integration Suite is installed on your C drive in a directory named SterlingCommerceMSSQL, replace the above string with “C:\SterlingCommerceMSSQL”. Leave the file open.
4. Repeat steps 1 and 2 for the files debug_noapp.cmd and debug_passphrase.cmd. Leave both files open.
5. From the Windows Taskbar, select **Start > Run** to open the **Run** box.
6. Type **regedit** in the **Open** box and click **OK** to open the Registry Editor.
7. Replace the classpath and path information in the debug_passphrase.cmd file with information from the registry using the following steps:
 - a. In the Registry Editor, go to HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Gentran Integration Suite at port Gentran Integration Suite *_Install_Port*\Parameters. For example, if you installed Gentran Integration Suite at port 30000, the value for Gentran Integration Suite *_Install_Port* would be 30000.
 - b. Go to registry entry JVM Option Number 0 and double-click **classpath**
 - c. In the **Value Data** box, highlight the portion of the value following the equal sign and press **Ctrl + C** to copy the value to the clipboard.
 - d. In the debug_passphrase.cmd file, find the portion that begins with CLASSPATH= and highlight everything following the equal sign up to the next **set** command.
 - e. Press **Ctrl + V** to replace the existing classpath information with the value you copied from the registry.
 - f. Return to the Registry Editor. Go to registry entry Path and double-click **path**.
 - g. In the **Value Data** box, highlight the entire value and press **Ctrl + C** to copy the value to the clipboard.
 - h. In the debug_passphrase.cmd file, find the portion that begins with PATH= and highlight everything following the equal sign up to the next **set** command.
 - i. Press **Ctrl + V** to replace the existing path information with the value you copied from the registry.
 - j. Save and close the debug_passphrase.cmd file.

8. Replace the classpath and path information in the debug_noapp.cmd file with information from the registry using the following steps:
 - a. In the Registry Editor, go to HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Gentran Integration Suite NoApp at port Gentran Integration Suite *_Install_Port*\Parameters. For example, if you installed Gentran Integration Suite at port 30000, the value for Gentran Integration Suite *_Install_Port* would be 30000.
 - b. Go to registry entry JVM Option Number 0 and double-click **classpath**
 - c. In the **Value Data** box, highlight the portion of the value following the equal sign and press **Ctrl + C** to copy the value to the clipboard.
 - d. In the debug_noapp.cmd file, find the portion that begins with CLASSPATH= and highlight everything following the equal sign up to the next **set** command.
 - e. Press **Ctrl + V** to replace the existing classpath information with the value you copied from the registry.
 - f. Return to the Registry Editor. Go to registry entry Path and double-click **path**.
 - g. In the **Value Data** box, highlight the entire value and press **Ctrl + C** to copy the value to the clipboard.
 - h. In the debug_noapp.cmd file, find the portion that begins with PATH= and highlight everything following the equal sign up to the next **set** command.
 - i. Press **Ctrl + V** to replace the existing path information with the value you copied from the registry.
 - j. Save and close the debug_noapp.cmd file.
9. Replace the classpath and path information in the debug_ops.cmd file with information from the registry using the following steps:
 - a. In the Registry Editor, go to HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Gentran Integration Suite Opsserver at port Gentran Integration Suite *_Install_Port*\Parameters. For example, if you installed Gentran Integration Suite at port 30000, the value for Gentran Integration Suite *_Install_Port* would be 30000.
 - b. Go to registry entry JVM Option Number 0 and double-click **classpath**
 - c. In the **Value Data** box, highlight the portion of the value following the equal sign and press **Ctrl + C** to copy the value to the clipboard.
 - d. In the debug_ops.cmd file, find the portion that begins with CLASSPATH= and highlight everything following the equal sign up to the next **set** command.
 - e. Press **Ctrl + V** to replace the existing classpath information with the value you copied from the registry.
 - f. Return to the Registry Editor. Go to registry entry Path and double-click **path**.
 - g. In the **Value Data** box, highlight the entire value and press **Ctrl + C** to copy the value to the clipboard.
 - h. In the debug_ops.cmd file, find the portion that begins with PATH= and highlight everything following the equal sign up to the next **set** command.

- i. Press **Ctrl + V** to replace the existing path information with the value you copied from the registry.
 - j. Save and close the debug_ops.cmd file.
10. Close the Registry Editor.
11. Stop all Gentran Integration Suite services.
12. Start MSSQL services.
13. Open a Command Prompt window and adjust the display settings to provide 2000 lines using the following steps:
 - a. Right-click in the Title Bar of the Command Line window and select **Properties > Layout**.
 - b. In the Screen Buffer Size section, enter **2000** in the **Height** field.
 - c. Click **OK**.
 - d. Select **Modify shortcut that started this window**, and click **OK**.
14. In the Command Prompt window, from the command line, type **debug_ops.cmd** and press **Enter**. Wait for the ops.log file to be created in the *install_dir\SI\logs* directory and verify that there are no errors in the log.
15. Open a new Command Prompt window. From the command line, type **debug_noapp.cmd** and press **Enter**. Wait for the Noapp.log file to be created in the *install_dir\SI\logs* directory and verify that there are no errors in the log. You should see a message similar to “Waiting for root pass phase to be set”.
16. Open a new Command Prompt window. From the command line, type **debug_passphrase.cmd** and press **Enter**. Wait for the pass.log, system.log*, and servicesctl.log* files to be created in the *install_dir\SI\logs* directory and verify that there are no errors in the logs.
17. Start Gentran Integration Suite.
18. Conduct the Thread Dump using the following steps:
 - a. Go to the debug_noapp.cmd Command Prompt window.
 - b. Press **Ctrl + Break**.
 - c. Open the Noapp.log file in the *install_dir\SI\logs* directory and scroll to the bottom.
 - d. Verify that the thread information is in the log file.

Conducting a Thread Dump in iSeries

To conduct a thread dump in iSeries:

1. Access the computer on which Gentran Integration Suite is installed.
2. On the iSeries command line, type **WRKACTJOB**.
3. Locate the two QP0ZSPWP jobs that are running in Gentran Integration Suite and note the user associated with these jobs.
4. Press **F11** twice to locate the job numbers.
5. On the iSeries command line, type **DMPJVM** and press **F4**.

6. At the Dump Java Virtual Machine (DMPJVM) panel, type **QP0ZSPWP** as the job name and *user* for the user prompt and *jobnumber* for the number prompt.
7. Press **Enter**. The thread dump is completed.
8. Repeat step 6 for the second job number.
9. On the command line, type **WRKSPLF** and locate the two QDMPJVM spooled files. These files are the output from your thread dump.

Conducting a Thread Dump in zSeries

To conduct a thread dump in zSeries:

1. Access the computer on which Gentran Integration Suite is installed.
2. Change directories to the *install_dir*.
3. On the command line, type **killall -QUIT java**.

This thread dump is completed.

Depending on your JVM, the thread dump is placed in one of the following locations:

- ◆ noapp.log in the *install_dir/logs* directory
- ◆ javacore*.txt

You can review the appropriate file to determine the processing bottlenecks in Gentran Integration Suite. Review each of the identified waiting processes for possible reasons for the waiting state. For example, a business process may contain a loop, improperly coded XPath statements, or use older services and adapters that are not as efficient as newer ones.

Reading a Thread Dump

A thread dump contains the following components:

- ◆ List of all threads running in the virtual machine (VM).
- ◆ List of all created monitors, including the current owner and the number of threads waiting for that monitor to be released.
- ◆ List of all special monitors used internally in the VM.

Each thread has an associated state. The following states may be associated with a thread:

- ◆ For other than IBM JVMs:
 - ◆ runnable – Running or executable thread.
 - ◆ object.wait () – Thread is waiting on a condition variable.
- ◆ For IBM JVM:
 - ◆ R – Running or executable thread.
 - ◆ S – Suspended thread
 - ◆ CW – Thread is waiting on a condition variable.

- ◆ MW – Thread is waiting on a monitor lock.
- ◆ MS – Thread is suspended waiting on a monitor lock. An MS state indicates a VM error.

The following example shows a thread in a runnable state and a thread in a waiting state:

```
"RMI TCP Accept-0" daemon prio=5 tid=0x015931f0 nid=0x1a runnable [bbbff000..bbbffc28]
  at java.net.PlainSocketImpl.socketAccept(Native Method)
  at java.net.PlainSocketImpl.accept(PlainSocketImpl.java:353)
  - locked <0xcc709620> (a java.net.PlainSocketImpl)
  at java.net.ServerSocket.implAccept(ServerSocket.java:448)
  at java.net.ServerSocket.accept(ServerSocket.java:419)
  at sun.rmi.transport.tcp.TCPTransport.run(TCPTransport.java:334)
  at java.lang.Thread.run(Thread.java:534)
"RMI RenewClean-[10.117.2.193:58586,com.sterlingcommerce.woodstock.util.frame.jndi.AddrClientFactory@c837cd]" daemon prio=5 tid=0x00b5f6f8 nid=0x13 in Object.wait() [bc97f000..bc97fc28]
  at java.lang.Object.wait(Native Method)
  - waiting on <0xcc709728> (a java.lang.ref.ReferenceQueue$Lock)
  at java.lang.ref.ReferenceQueue.remove(ReferenceQueue.java:111)
  - locked <0xcc709728> (a java.lang.ref.ReferenceQueue$Lock)
  at sun.rmi.transport.DGCClient$EndpointEntry$RenewCleanThread.run(DGCClient.java:500)
  at java.lang.Thread.run(Thread.java:534)
```

Resolving Database Issues

After you run Gentran Integration Suite for some time, you may find that the database begins to fill up or becomes full. Your business processes may be designed to store more data than necessary in your database, or you may have old data, which can be archived or purged.

For information about auto extending your database, see your database documentation.

This section covers the following topics:

- ◆ *Full Database Issues* on page 89
- ◆ *Failed Database Connection Requests* on page 92
- ◆ *Tracking JDBC Connections* on page 94
- ◆ *Database Maintenance Check* on page 95
- ◆ *Database Down Check* on page 95
- ◆ *Optimizing the Database* on page 97
- ◆ *Using the Database Statistics Report* on page 98
- ◆ *Slow Performance in Oracle* on page 100

Full Database Issues

You may find that your Gentran Integration Suite database becomes full for a variety of reasons. When your database becomes full, it can cause performance issues that appear to be performance issues related to memory or other resource usage.

If you have your database set to auto extend, the database may attain 99 percent capacity before extending automatically. In this situation, the 99 percent capacity is not an indication of a performance issue.

After you run Gentran Integration Suite for some time and you take out from the database approximately the same amount of data that you put into the database, the database size becomes consistent.

If other applications are using the same database that Gentran Integration Suite uses, verify that the other applications are experiencing slow database returns as well, before making any adjustments to the database. If the other applications are not experiencing database issues, the performance issue may be related to another cause, and not the database.

Symptoms and Causes of a Full Database

Symptoms of a full database may include:

- ◆ Database does not accept new data.
- ◆ Database Usage Report shows your database usage is increasing to capacity.
- ◆ Gentran Integration Suite interface is slow.
- ◆ Business process execution times are slow.

Determining the cause of a full or filling database may require you to review your business processes and your archival and purging practices.

Causes of a full or filling database may include:

- ◆ Persistence levels set to high globally, in your business processes, or for each activity in a business process, causing too much data to be persisted to the database.
For more information about changing persistence levels, see *Changing Persistence Levels* on page 68.
- ◆ Indexing business process is not working properly, causing no data to be flagged for archival or purging and no data being removed from the database. The indexing business process runs every 10 minutes.
- ◆ Archival business process is not working properly, causing no data to be removed from the database and archived to another location. The archive business process checks for the archive flag and then the date the archive is to be completed. When the archive date is reached, the archive business process archives the record.
- ◆ Purging business process is not working properly causing no data to be purged from the database. The purge business process checks for the purge flag and then the date the purge is to be completed. When the purge date is reached, the purge business process purges the record from the database.
- ◆ Large or old files are being stored in the database that can be archived elsewhere, causing the database to fill up.
- ◆ Size of the database is not optimized for your processing requirements, causing the database to fill up prematurely.
- ◆ Data retention period is set too high, causing data to be kept in the database longer than necessary.
- ◆ BPREcovery business process is not running or failing causing business processes to remain in the database in an interrupted state.
- ◆ IndexBusinessProcess business process is not running or failing causing no business processes to be flagged for archival or purging.

To determine the cause of a full database, you will complete some, if not all, of the following tasks, depending on your situation:

- ◆ Monitor the Database Usage Report in Gentran Integration Suite to determine the levels of database usage. If the usage level increases to the high end of its capacity, you may need to take corrective action unless you have your database set to auto-extend (increase in capacity when a specified level is reached).

For more information about auto-extending your database, see your database documentation.

For more information about Database Usage Report, see *About the Database Services Area* on page 24.

- ◆ Review the archive.log and the database-specific logs in the **Operations > System > Logs** page for information on archiving and purging activities, and errors.

Resolving Full Database Issues

The following are resolutions for a full database, depending on the identified cause:

Changing Persistence Levels

Persistence is the activity of storing data in your database. Since you are storing data in the database, persisting large amounts of data can cause your database to fill up. You may have some business processes that you want to persist all steps and data to the database and other business processes that you do not need to persist any steps or data to the database.

When you design your business processes review your persistence requirements and set persistence levels appropriately to reduce the chances of your database filling up and slowing your processing and system.

For information on how to change persistence levels, see *Changing Persistence Levels* on page 68.

Completing Indexing Activities

If your Index business process is scheduled, but not running, no data is flagged for archival or purging, and your database can fill up. Review the index business process and verify that the schedule is defined and turned on in the business process, and that business processes are being indexed.

To verify that the index business process is running, review the archive.log file for data that was purged or archived. This log file will state whether each data purge or data archive was successful. If the archive.log file shows no data being archived or purged, your index business process may not be running, causing your database to fill up. If the indexing business process is running and failing, contact Sterling Commerce Customer Support for assistance.

Completing Archiving Activities

If your archival business process is scheduled, but is not running, your database can fill up because data is not being archived and moved off the database tables to another storage location. Review the archival business processes and verify that the schedule is defined and turned on in the business process, and that business processes are archived after a reasonable amount of time.

To verify that the archival business process is running, review the archive.log file for data that was purged or archived. This log file will state whether each data purge or data archive was successful. If the archive.log file shows no data being archived or purged, your archival business process may not be running, causing your database to fill up. If the archival business process is running and failing, contact Sterling Commerce Customer Support for assistance.

Completing Purging Activities

If your purging business process is scheduled, but is not running, your database can fill up because data is not being removed from the database tables. Review the purging business processes and verify that the schedule is defined and turned on in the business process, and that business processes are purged after a reasonable amount of time.

To verify that the purging business process is running, review the archive.log file for data that was purged or archived. This log file will state whether each data purge or data archive was successful. If the archive.log file shows no data being archived or purged, your purging business process may not be running, causing your database to fill up. If the PurgeService business process is running and failing, contact Sterling Commerce Customer Support for assistance.

Reducing the Amount of Large or Old Data Stored in the Database

If you have old data and large files using database resources and causing your database to fill up, archive the large file or old files to another file system instead of in the database tables, so they can be retrieved if needed.

Optimizing the Size of Your Database

You may find that your initial database size requirements are different from your requirements as your business and processing needs change. If your database is filling and your volume is increasing:

- ◆ Increase the size of your database to reflect your processing needs.
- ◆ Compress your database, depending if your database has this feature, to save more data to the same size database.

For more information about compressing your database, see your database documentation.

Optimizing the Data Retention Period

You may find that you are retaining information in the database for longer periods of time than necessary. Set the data retention period in the database to lower levels as appropriate.

For more information about setting data retention periods in your database, see your database documentation.

BPREcovery Business Process is Not Running or is Failing

If your BPREcovery business process is not running or is failing:

1. Contact Sterling Commerce Customer Support for assistance with the BPREcovery business process
2. Run the corrected BPREcovery.
3. Restart, resume, or terminate interrupted business processes.

Note: You can complete a mass termination of business processes if there are many in an interrupted state.

Failed Database Connection Requests

You may find that your business process throughput slows down in Gentran Integration Suite. *Data bursts* (increased data volume) and *connection* or *cursor leaks* (unreleased database connections), can result in an increased number of failed pool requests.

Symptoms and Causes of Failed Database Connection Requests

Symptoms of an increasing number of failed pool requests may include:

- ◆ Slow business process throughput (volume being processed in a specific time period).
- ◆ Increasing number of failed pool requests on the Database Usage report on the **Operations > System > Troubleshooter** page. This report shows the number of pool requests and the number of failed requests.

Determining the cause of an increasing number of failed pool requests may require you to investigate many areas of the system and how you are implementing Gentran Integration Suite.

Causes of an increasing number of failed pool requests may include:

- ◆ Increased data volume, causing more business processes to compete for the connections.
- ◆ Data bursts – High rates of data in short transmission periods.
- ◆ Connection or cursor leaks – Database connections that are opened for a business process and then kept open by that business process instead of releasing it back to Gentran Integration Suite to be used by other business processes. This causes fewer available connections to the database.
- ◆ Improperly tuned performance properties – Fewer database pools configured than required by your processing volumes.

For more information about tuning Gentran Integration Suite, see *Calculated Recommendations* on page 45.

To determine the cause of failed database connection requests:

- ◆ Monitor the Database Usage report on the **Operations > System > Troubleshooting** page. This report shows the number of failed requests for each pool type. You can use this information to determine which pool settings to change in the Performance Tuning Utility.

For more information about the database usage report, see *Reviewing System Information* on page 18.

- ◆ Monitor database tools specific to your database.

For more information, see your database documentation.

Resolving Database Connection Issues by Increasing the Database Pool Size

You can change the size of the database pools to allow for more connections to the database, to resolve the following issues:

- ◆ Increased volume levels
- ◆ Data bursts
- ◆ Connection or cursor leaks

You can increase the size of the specific database pools in Gentran Integration Suite using the **Operations > System > Performance > Tuning** page.

For more information, see *Performance Tuning Utility* on page 45.

Tracking JDBC Connections

Gentran Integration Suite provides the capability to track your JDBC connections. This is useful when troubleshooting high connection usage or running out of connections at the database. These problems may be caused by connection leaks that occur when a component requests a connection from the pool and does not return it. When enabled, JDBC tracking will track all active JDBC connections.

If JDBC tracking is enabled, you can view or save a JDBC report to help you troubleshoot database connection problems. You can view the report for your own troubleshooting, or save the report and send it to Sterling Commerce Customer Support to aid in resolving the problem. The JDBC report provides current database stack dump information at the time that the JDBC report is viewed or saved.

To Enable or Disable JDBC Tracking

1. From the **Administration** menu, select **Operations > JDBC Monitor**.
2. In the **On/Off** box, select one of the following:
 - ◆ If JDBC Tracking is Disabled, click **(Enable)** to enable JDBC Tracking.
 - ◆ If JDBC Tracking is Enabled, click **(Disable)** to disable JDBC Tracking.

JDBC Tracking enabled through the JDBC Monitor will be reset to disabled if Gentran Integration Suite is restarted. This feature helps prevent degraded performance if JDBC Tracking is accidentally left enabled after you are finished troubleshooting. You should always disable JDBC Tracking when you are finished troubleshooting.


If you to enable JDBC tracking for extended troubleshooting so that it remains enabled if is restarted, you can use the internal **use tracking** parameter. To control JDBC Tracking with the **use tracking** parameter, perform the following steps:

1. Stop Gentran Integration Suite.
2. In the *install_dir/properties* directory, locate the *jdbc.properties.in* file.
3. Open the *jdbc.properties.in* file in a text editor.
4. In the *jdbc.properties.in* file, locate the **use tracking** parameter.
 - ◆ To enable JDBC tracking, change *useTracking=false* to *useTracking=true*
 - ◆ To disable JDBC tracking, change *useTracking=true* to *useTracking=false*
5. Save the *jdbc.properties.in* file.
6. Restart Gentran Integration Suite to apply the change.

Always disable JDBC tracking when you are finished troubleshooting to improve performance.

To View or Save a JDBC Report

1. From the **Administration** menu, select **Operations > JDBC Monitor**.
2. In the **View** box, next to **View JDBC Report**, click one of the following:

- ◆ To view a JDBC Report, click the View icon: . Click **F5** to refresh the report, as needed.
- ◆ To save a JDBC Report, click (**Download**).

Database Maintenance Check

Gentran Integration Suite performs a scheduled maintenance check on the database every Monday at 4:00 A.M. The system checks to see if any of the following conditions have occurred:

- ◆ The database has become more than 80 percent full
- ◆ The database has grown more 30 percent since the last maintenance check
- ◆ The database needs to be indexed

If any of the above conditions are found, the system sends an event notification about the condition. By default, it sends an e-mail to the system administrator.

This maintenance check is performed by the DB Monitor service and its associated Schedule_DBMonitorService business process.

You can use the Database Statistics Report to view information about business processes in the database. See *Using the Database Statistics Report* on page 98 for more information.

If you are notified that the database needs to be indexed, contact your database administrator or use the db_optimization_tool script to rebuild the indexes. For more information about the db_optimization_tool script, see *Optimizing the Database* on page 97.

Database Down Check

Gentran Integration Suite performs a scheduled check on the database connection every 60 seconds to verify that the database has not gone down. These checks are performed by the DB Resource Monitor and are logged in the resource_monitor.log file.

If the check finds that the database has gone down, the system administrator is sent an e-mail notification that the database is unreachable. The date and time of the check is also provided as shown in the following example:

```
Database Unreachable
```

```
The database was unreachable at 2005.07.19 14:29:36 EDT.
```

```
E-mails will be sent at increasing intervals of
1, 5, 15, 30, 45, and 60 minutes until Gentran Integration Suite is shut
down or the database can be reached.
```

```
Event Information:
```

```

eventType: ResourceMonitor.DBResourceMonitor.ExceptionSQLException.3
ExceptionLevel: Exceptional
timestamp: 1121797776902
datetime: 2005.07.19 14:29:36 EDT
Status: SQL Exception
host: myhost
node: node1
subject: Database Server Unreachable (Node- node1 Host- myhost)

```

Additional e-mail notifications are sent at 1, 5, 15, 30, 45, and 60 minutes unless Gentran Integration Suite is shut down or the database can be reached. If the database can be reached, the system administrator is sent an e-mail notification that the database connection was successful as shown in the following example:

```
Database Connection Successful
```

```
Database was reachable at 2005.07.19 14:34:47 EDT.
```

Event Information:

```

eventType: ResourceMonitor.DBResourceMonitor.SUCCESS.1
ExceptionLevel: Exceptional
timestamp: 1121798087408
datetime: 2005.07.19 14:34:47 EDT
Status: Success
host: myhost
node: node1
subject: Database Connection Successful (Node- node1 Host- myhost)

```

Configuring the DB Resource Monitor

The DB Resource Monitor can be configured to perform the database down check at intervals other than the default interval of one minute. To change the interval, perform the following actions:

1. In the *install_dir/properties* directory, locate (or create, if necessary) the *customer_overrides.properties* file.
2. Open the *customer_overrides.properties* file in a text editor.
3. Add the following override statement:

```
resourceMonitor.DBResourceMonitor.delay=new_interval
```

The value for *new_interval* is the interval, in milliseconds, at which you want the database down check to run. For example, the default value of 60000 runs the check every 60 seconds, or once a minute. Assume that you want to run the check every 10 minutes. You must change the value to 600000 (600000 milliseconds = 600 seconds = 10 minutes). To do so, add the following line to the *customer_overrides.properties* file: `resourceMonitor.DBResourceMonitor.delay=600000`

4. Save and close the *customer_overrides.properties* file.
5. Stop and restart Gentran Integration Suite to use the new values.

For assistance in determining the proper settings for your business needs, contact Sterling Commerce Customer Support.

Optimizing the Database

Gentran Integration Suite provides a database optimization script that allows you to optimize your database by performing tasks such as analyzing tables and rebuilding indexes.

To run the database optimization script, from the command line, run one of the following commands:

- ◆ For UNIX, run `install_dir/bin/db_optimization_tool.sh options`
- ◆ For Windows, run `install_dir\bin\db_optimization_tool.cmd options`

The available options are shown in the following table:

Option	Description
-l	View a list of all index rebuild or analyze database SQL statements. Must be used with -i or -a , or both. Uses the current date and time as the cutoff for expired data, unless the -t option is used. Either -l or -r is required.
-r	Run all index rebuild or analyze database SQL statements. Must be used with -i or -a . Uses the current date and time as the cutoff for expired data, unless the -t option is used. Either -l or -r is required.
-i	Use a list of the table indexes that need to be rebuilt (optimized). Must be used with -l or -r . Either -i or -a is required.
-a	Use a list of the tables to analyze. Must be used with -l or -r . Either -i or -a is required.
-t	Date and time cutoff to look for expired data. Format: <code>yyyyMMdd-HH:mm:ss.SSS</code> . Default is current date. This date and time will be used as the starting point to check for expired tables. If the table has expired since the last time it was rebuilt, then that table is added to a list for rebuild or index analyzing. Must be used with -l or -r . Optional.
-o	Specifies a file name to save output messages to. Provide a full path to the file. If not used, output will only show on the screen. Must be used with -l or -r . Optional.
-d	Check all tables. If not used, the check will only be done on tables that were recorded on the database after the last rebuilds. Must be used with -l or -r . Optional.
-p	Print the stack trace if there is an exception. If used with the -o option, the stack trace will print to a file; otherwise it will display on the screen. Must be used with -l or -r . Optional.
-h or -?	View the help screen

Examples

To view a list of the table indexes that need to be rebuilt and save the output to a file named myList, go to the `install_dir/bin` directory and run the following command:

```
db_optimization_tool.sh -l -i -o myList
```

To rebuild the indexes for all tables, run the following command:

```
db_optimization_tool.sh -r -i -d
```

Using the Database Statistics Report

The Database Statistics Report provides statistical information about business processes in the database. You might want to review the Database Statistics Report if you experience full database or failed database connection problems. You may also need a Database Statistics Report because you received automated notification from Gentran Integration Suite that the database health may need to be checked.

Viewing the Database Statistics Report

To view a Database Statistics Report:

1. From the **Administration** menu, select **Operations > Reports**.
2. On the Reports page, under **Search**, enter `DBStats` in the **Name** box, and click **Go!**

Note: The DBStats report can also be found under **List**. However, it cannot be found using the Type list under **Search**.

3. Click **Source Manager**.
4. Select the appropriate report format from the drop-down menu and click **Execute**.

Reading the Database Statistics Report

The Database Statistics report contains several sections that each provide different statistical information about business processes in the database.

Business Process Runs in Active System

The Business Process Runs in Active System section lists all business process definitions that have had at least one instance executed. The following fields are provided:

- ◆ Business Process Name – The name of the business process instance.
- ◆ Number of Runs - The number of times the instance has run.
- ◆ Number of Persisted Steps - The number of steps persisted per business process definition.
- ◆ Average Number of Persisted Steps – The average number of steps persisted.

Business Process Runs in Active System - Last 24 Hours

The Business Process Runs in Active System - Last 24 Hours section lists all business process definitions that have had at least one instance executed in the last 24 hours. It contains the same fields as *Business Process Runs in Active System* on page 98.

Business Process Runs in Active System - Interval of Last 48 to 24 Hours

The Business Process Runs in Active System - Interval of Last 48 to 24 Hours section lists all business process definitions that have had at least one instance executed during the previous day (24 to 48 hours before the current time). It contains the same fields as *Business Process Runs in Active System* on page 98.

Number of BPs by Removal Method

The Number of BPs by Removal Method section shows the number of business processes that are flagged to be archived, purged, or indexed. It contains the following fields:

- ◆ Archive Flag – Shows one of the following conditions:
 - ◆ To Be Archived – Business processes flagged to be archived
 - ◆ To Be Purged – Business processes flagged to be purged
 - ◆ Archived, to be Purged – Business processes that have been archived, and are flagged to be purged
 - ◆ To Be Indexed – Business processes flagged to be indexed
- ◆ Count - Total number of business processes with the stated Archive flag

Number of Eligible BPs by Removal Method

The Number of Eligible BPs by Removal Method section shows the number of eligible business processes that are flagged to be archived, purged, or indexed. It contains the same fields as *Number of BPs by Removal Method* on page 99. The Number of BPs by Removal Method To Be Indexed count minus the Number of Eligible BPs by Removal Method To Be Indexed count provides the number of business processes that for whatever reason are not eligible for indexing and therefore will not be purged until they become eligible for indexing.

Archive Dates by Removal Method

The Archive Dates by Removal Method section shows the date range of business processes that are flagged to be archived, purged, or indexed. It contains the following fields:

- ◆ Archive Flag – Shows one of the following conditions:
 - ◆ To Be Archived – Business processes flagged to be archived
 - ◆ To Be Purged – Business processes flagged to be purged
 - ◆ Archived, to be Purged – Business processes that have been archived, and are flagged to be purged
- ◆ Min – Date and time of the first expired business process
- ◆ Max – Date and time of the last expired business process

Table Row Counts

The Table Row Counts section shows the number of rows for each table in the active system. It contains the following fields:

- ◆ Table Name – Name of the table
- ◆ Count – Number of rows in the table

Table Row Counts - Last 24 Hours

The Table Row Counts - Last 24 Hours section shows the number of rows for each table in the active system during the last 24 hours. It contains the same fields as *Table Row Counts* on page 99.

Table Row Counts - Interval of Last 48 to 24 Hours

The Table Row Counts - Interval of Last 48 to 24 Hours section shows the number of rows for each table in the active system during the previous day (24 to 48 hours prior to the current time). It contains the same fields as *Table Row Counts* on page 99.

Table Row Counts - Index

The Table Row Counts - Index section shows the number of rows that have not been indexed for each table in the active system. It contains the same fields as *Table Row Counts* on page 99.

Table Row Counts - Archive

The Table Row Counts - Archive section shows the number of rows eligible for archiving for each table in the active system. It contains the same fields as *Table Row Counts* on page 99.

Table Row Counts - Purge

The Table Row Counts - Purge section shows the number of rows eligible for purging for each table in the active system. It contains the same fields as *Table Row Counts* on page 99.

Number of Rows with No Matching Records in Archive Info

The Number of Rows with No Matching Records in Archive Info section shows the number of rows in a table that are orphaned (have no matching records in the archive data). It contains the following fields:

- ◆ Table Name – Name of the table
- ◆ Orphaned Record Count – Number of orphaned rows

No Matching Workflow ID (Null)

The No Matching Workflow ID (Null) section shows the number of records in DATA_TABLE with NULL workflow IDs or -1 workflow IDs. It contains the following fields:

- ◆ Number of BPs with No Matching Records in Data Table, and Count – Total number of workflows with -1 Workflow IDs in the data table
- ◆ No Workflow ID Assigned, and Count – Number of business processes that have a null workflow ID

Slow Performance in Oracle

Enabling BLOB Caching

By default, Gentran Integration Suite configures Oracle with BLOB caching disabled when you choose Oracle as your database during Gentran Integration Suite installation. You can significantly improve performance by enabling the cache on the BLOB data object.

To enable the cache:

1. From the Gentran Integration Suite Administration Menu, select **Operations > System > Support Tools > SQL Manager**

2. In the text box, type:


```
alter table data_table modify lob (data_object) (cache)
```
3. Click **Execute**
4. When asked if you want to proceed, click **OK**

Checking for Cursor Leaks in Oracle Databases

Gentran Integration Suite provides a utility called `checkCursorLeaks` that checks for cursor leaks generated by Gentran Integration Suite in an Oracle database. This utility is useful for troubleshooting database issues when cursor leaks are suspected. The `checkCursorLeaks` utility output provides information useful for identifying cursor leaks. This information includes:

- ◆ User name
- ◆ SQL text
- ◆ Count
- ◆ Total number of cursor leaks.

To check for cursor leaks in an Oracle database:

1. Change to the `install_dir/bin` directory.
2. Run one of the following commands:
 - ◆ `./checkCursorLeaks.sh` (UNIX or Linux)
 - ◆ `checkCursorLeaks.cmd` (Windows)

Note: The `checkCursorLeaks` utility is only supported for Oracle databases.

Preventing Stale Connections to the Oracle Database

In some cases, a firewall or router between the Gentran Integration Suite server and the Oracle server may stop the connection between the two servers if they are idle for a period of time. The connections may not drop, but just become unusable.

If you experience this situation, it can be prevented by simulating traffic between the servers even when they are not being used. The simulated traffic will prevent the connection from being stopped due to idleness.

To prevent stale connections between Gentran Integration Suite and the Oracle database, perform the following steps:

1. On the Oracle server, change to the `$ORACLE_HOME/network/admin` directory.
2. Open the `sqlnet.ora` file in a text editor.
3. Add, or change, the **`sqlnet.expire_time`** parameter as shown in the following line:

```
sqlnet.expire_time = 10
```

This sends a probe through the connection every 10 minutes. This probe verifies that the client/server connection is active.

4. Save and close the `sqlnet.ora` file.

5. Restart the listener on the Oracle server.
6. Restart Gentran Integration Suite.

Although very small, a probe packet generates additional network traffic that could affect the performance of the network, depending on the number of connections. If network performance is adversely affected, try using a different expiration time. For example, try 15 or 30 minutes instead of 10.

Resolving Business Process Issues

You may find that the number of business processes in a halting, halted, waiting, or interrupted state increases. This indicates that either performance is not up to your requirements and corrective action is required, or you have business processes that have errors and need manual attention.

This section covers the following topics:

- ◆ *Halted, Halting, Waiting, or Interrupted Business Processes* on page 103
- ◆ *Running a Business Process from the Command Line* on page 105

Halted, Halting, Waiting, or Interrupted Business Processes

In Gentran Integration Suite, business processes that fail receive a state of halted. Failed business processes are not completed because business processes in a completed state can be indexed, and archived, or purged. If Gentran Integration Suite gave a failed business process a state of completed, the failed business processes could be archived or purged without corrective action being taken. Giving a failed business process a state of halted enables you to take manual corrective action without the business process being archived or purged.

A business process in a Halting, Halted, Interrupted_Man, or Interrupted_Auto state requires immediate attention because the business process has stopped processing. In Gentran Integration Suite, business processes remain in a halted or interrupted state until some action is taken on the business process.

When you notice a halted or interrupted business process, you have two options:

- ◆ Terminate the business process.
- ◆ Restart the business process.

A business process in a waiting state may not require immediate attention; the business process may be waiting on a resource or a document from another business process before completing. If you notice numerous business processes in a waiting state, this may indicate a performance issue that requires attention.

When you notice a business process in a waiting state, you have three options:

- ◆ Allow the business process to remain in the waiting state if it is waiting on resources or a service or activity that is disabled, but will be enabled.

- ◆ Terminate the business process.
- ◆ Restart the business process.

Symptoms of Halted, Waiting, or Interrupted Business Processes

Symptoms of an increasing number of business processes in a halting, halted, waiting, or interrupted state may include:

- ◆ Slow system performance
For more information on resolving a slow system, see *Resolving Slow System Issues* on page 80.
- ◆ Database getting full or having performance issues
For more information on resolving a database issues, see *Resolving Database Issues* on page 89.
- ◆ Business processes complete with errors, which places them in a halted state
- ◆ Business Process Usage report shows an increasing number of business processes in a halted, halting, waiting, or interrupted state

Potential Causes of Halted, Waiting, or Interrupted Business Processes

Determining the cause of an increasing number of business processes in a halting, halted, waiting, or interrupted state may require you to investigate many areas of the system and how you are implementing Gentran Integration Suite.

Causes of an increasing number of business processes in a halting, halted, waiting, or interrupted state may include:

- ◆ System, business process, or activity schedules are disabled. For example, if a business process requires an output from, or access to, a different service or business process that is scheduled to work, but the schedule is not turned on, this places the business process in a halted or waiting state.
- ◆ System errors. For example, Java, JVM, out of memory errors, or operating system errors may cause a business process to halt or be interrupted. Check your business process logs for causes of the halted or interrupted business processes. If the logs show JVM errors, Java errors, or operating system errors, review your operating system documentation for resolutions.
- ◆ Improperly designed business processes. For example,
 - ◆ Using the Wait service in a business process for time periods of less than one minute, instead of using the Sleep service. This can cause a business process to be placed into a waiting state instead of an active state, until the Wait service completes.
 - ◆ Using Produce and Consume services instead of the Invoke Subprocess service to invoke subprocesses. Using Produce and Consume services places a business process into a waiting state, waiting for the Produce service to create the document and for the Consume service to use the document.
- ◆ Gentran Integration Suite stops running. For example, your site experiences a power outage and you must restart Gentran Integration Suite after power is restored. Your business processes at the time of the power outage may be in halted, interrupted, or waiting states after the recovery operations run, depending on the activities being completed at the time of the outage.

Determining the Cause of Halted, Waiting, or Interrupted Business Processes

To determine the cause of an increasing number of business processes in a halting, halted, waiting, or interrupted state:

- ◆ Review the Business Process Usage report on the **Operations > System > Troubleshooting** page. This report shows the number of business processes in the different states. You can click the number next to the state to view detailed information about the process, which may help troubleshoot any problems.
- ◆ Review appropriate system and business process schedules to verify that they are turned on. For example, if you notice many business processes are halting and each of these business processes is dependent on the schedule of another business process or service, this indicates that the scheduled business process or service may not be turned on.
- ◆ Review the Performance Statistics report for information related to the business process execution times. Increasing execution times for key business processes or activities may indicate that a business process is not efficiently designed, or that a resource leak may have occurred.

For more information about the Performance Statistics report, see *Turning On and Off Performance Statistics* on page 43.

- ◆ Review appropriate log files in the **Operations > System > Logs** page. The following log files may provide troubleshooting information:
 - ◆ `Archive.log` – Provides information about successful and unsuccessful archiving and purging activities.
 - ◆ `wf.log` – Provides information about business process states, errors, and processing.
 - ◆ Adapter and service logs – Provide information on specific adapter or service activities and errors.
 - ◆ `system.log` – Provides information about the general system.
 - ◆ `noapp.log` – Provides information about the application server. Review business processes to verify that they are effectively designed to meet your business needs.

Running a Business Process from the Command Line

Genran Integration Suite provides a utility called `workflowLauncher` that allows you to launch a business process from the command line. The utility launches a business process and provides the status when it has completed, or when a timeout has occurred (whichever comes first).

To launch a business process from the command line:

1. Change to the `install_dir/bin` directory.
2. Run one of the following commands:
 - ◆ `./workflowLauncher.sh -n BPname [option(s)] (UNIX)`

◆ `workflowLauncher.cmd -n BPname [option(s)]` (Windows)

BPname is the name of the business process you want to launch. The available options are listed in the following table:

Option	Description
-n <i>BPname</i>	Specifies the name of the business process definition to launch. Mandatory. Example: <code>workflowLauncher.sh -n testBP</code> An error will be generated if the business process name is not provided.
-s	Specifies Silent mode. Optional. In Silent mode, less information about the business process is output to the standard output. Default is Verbose mode.
-u <i>username</i>	Specifies the username to include the userToken from. Optional. Example: <code>-u username</code> This parameter is used to specify the user token of the business process.
-d <i>directory</i>	Specifies the output directory. Optional. Example: <code>-d absolutePath/directory</code> Used in conjunction with the <code>-t</code> parameter when you generate step XML trace files. This is used to set the directory output of the trace files.
-t	Turns on step XML trace generation. Optional. Used in conjunction with the <code>-d</code> parameter when you want to generate XML trace files. By default, the trace file name will be prefixed with the business process ID. To turn this feature off, use the <code>-o</code> parameter.
-o	Used to prevent prefixing the trace file name with the business process ID when generating step XML trace files. Optional.
-v <i>versionNumber</i>	Specifies a specific version of the business process definition to launch. Optional. Example: <code>-v business_process_version_number</code> The version number is <i>not</i> the user description string, it is the internal running number (a positive integer) that is automatically incremented each time you check in a new version of the same business process name. If there is only one version, the version number will be 1. An error will be generated if the version number is not found. You can view the version number by resting your mouse on, or clicking the information icon ⓘ for the business process in the Change History section of the Gentran Integration Suite Business Process Source Manager.
-f <i>inputFile</i>	Specifies a file to use as the primary document when launching the business process. Optional. Example <code>-f path/filename</code> Only one file is allowed. An absolute path or a relative path to the file may be used. To specify the encoding of the file, use the <code>-e</code> parameter. An error will be generated if the file is not found.
-e <i>encoding</i>	Specifies the encoding of the primary document. Optional. Example <code>-e encodingName</code> Used in conjunction with the <code>-f</code> parameter. To see what encoding names are supported, see the <code>encodings.properties</code> file in the <code>install_dir/properties</code> directory. The full listing is shown in the <code>encodings_large.properties</code> file.

Option	Description
<code>-w count</code>	Specifies the number of times to sleep. Default is 10. Example: <code>-w count</code> Used in conjunction with the <code>-i</code> parameter to define the total timeout value. For an example showing the total timeout value, see the <code>-i</code> parameter.
<code>-i time</code>	Specifies the length of time, in milliseconds, to sleep. Default is 3000 (3 seconds). Example: <code>-i time</code> Used in conjunction with the <code>-w</code> parameter to define the total timeout value. For example, the default value for <code>-i time</code> is 3000 milliseconds and the default value for <code>-w count</code> is 10. Therefore, the default timeout value is 30,000 milliseconds (3000 milliseconds x 10), or 30 seconds.
<code>-x</code>	Specifies not to return an exit code. Optional. By default, the utility returns an exit code.
<code>-h -?</code>	Displays the help screen.

Exit Codes

The workflowLauncher utility returns the following exit codes:

Exit Code	Description
0	The business process completed with success status.
1	The business process completed with error status.
99	The workflowLauncher utility reached timeout while waiting for status from the business process.

Resolving Cache Issues

You may find that your processing slows if you have not tuned your cache settings appropriately, resulting in calls to the database for data where having the data available for use in the appropriate cache would be more efficient.

This section covers the following topics:

- ◆ *Types of Memory Cache* on page 108
- ◆ *Symptoms and Causes of Inefficient Use of Cache* on page 109
- ◆ *Resolutions for Inefficient Use of Cache* on page 110

Types of Memory Cache

Gentran Integration Suite utilizes the following types of memory cache to execute processes:

Memory Types	Description
Soft reference cache	When objects are removed from memory cache to keep memory cache size constant, they are moved to a soft reference cache which can grow or shrink based on available memory. If the JVM needs to reclaim memory space, it takes it from the soft reference cache.
Memory cache	Uses the amount of memory necessary to hold the objects in memory at all times (user configurable using <code>noapp.properties</code>).
Disk cache	Objects can be read more quickly from disk than database. When objects are no longer in the soft reference cache because they were garbage collected, the disk cache provides a faster access mechanism than object retrieval from database (user configurable using <code>noapp.properties</code>).

A *soft* cache is an in-memory cache. The difference between a soft cache and a traditional cache is that the memory used to cache the objects is reclaimed as soon as memory constraints become limited. When the system is low on memory, the cached objects are automatically removed from memory. This is possible because the objects themselves are backed by persistent data in the database or disk and can be re-read from there.

This behavior is very similar to an operating system disk cache. The disk cache retains recently read files, or portions of files, in memory to speed up subsequent access to that data. However, if the system needs memory for applications to run, the cache size is reduced proportionally. You can see this in Windows by watching the size of the disk cache, in the performance monitoring tools, fall as you open a large file.

Gentran Integration Suite follows this same practice with soft caches. Most soft caches in Gentran Integration Suite are implemented by Java Soft References. For more information about Java Soft References, see <http://java.sun.com/developer/technicalArticles/ALT/RefObj/>

The JVM short-lived min/max memory values are used for New Age Generations for Garbage Collections. References here:

http://publib.boulder.ibm.com/infocenter/javasdk/v5r0/index.jsp?topic=/com.ibm.java.doc.diagnostics.50/html/mm_gc_generational.html

<http://java.sun.com/docs/hotspot/gc1.4.2/faq.html>

Based on the amount of memory allocated to Gentran Integration Suite and the number of CPUs, using the tuning wizard Gentran Integration Suite will allocate disk cache and in-memory cache which is a combination of soft and memory cache displayed on the cache usage screen. Objects that are used most often, such as using various functions on the UI screen are stored in softcache. If a business process that is executing needs more memory allocation, the JVM handles the Garbage Collection by flushing out expired objects from MEMORY.

Symptoms and Causes of Inefficient Use of Cache

Symptoms of inefficient use of cache include:

- ◆ Slow running business processes
- ◆ Longer completion times for business processes
- ◆ Increasing hit numbers on the Cache Usage report on the **Operations > System > Troubleshooter** page.

Determining the cause of the inefficient use of cache may require you to investigate many areas of your system and how you are implementing Gentran Integration Suite.

Causes of inefficient use of cache may include:

- ◆ Improperly tuned performance properties
- ◆ High cache usage of less frequently used large objects
- ◆ Low cache usage of frequently used small objects

To determine the cause of inefficient use of cache, review the Cache Usage report on the **Operations > System > Troubleshooter** page for the number of counts, requests, and hits for each cache in Gentran Integration Suite.

An increase in the number of requests and a decrease in the number of hits indicates that you may need to increase the cache size for the impacted cache. Each request that is not met with a hit, requires a call to the database for the data, which slows processing times and uses more resources that could be used by other components of Gentran Integration Suite.

For example, if the SchemaCache on the Cache Usage report had 5 counts, 5 requests and 0 hits, this shows that the schema cache may be too small for the number of requests from the cache or too small for the size of the documents used. The 0 hits means that the business processes requesting the schema from the cache did not find the data requested and called the database to receive the schema data, which slowed the processing times.

Resolutions for Inefficient Use of Cache

The following are resolutions for the inefficient use of cache, depending on the identified cause:

Improperly Tuned Cache Performance Properties

If you review the Cache Usage report and notice the number of requests increasing and the number of hits decreasing for the same cache, increase the cache values using the Performance Tuning Utility or manually increase the values in the `install_dir/properties/noapp.properties` file.

For more information about increasing the cache values in the Performance Tuning Utility, see *Viewing or Editing Performance Configuration Settings* on page 46.

For more information about increasing the cache values manually, see *Manual Performance Tuning* on page 53.

Trade offs in increasing cache sizes include:

- ◆ Increasing the cache sizes too much can make the system no more effective than reading from the disk. Depending on your system and your configuration, you may need to adjust the cache settings to attain peak performance without losing the benefit of cache over disk.
- ◆ Increasing the cache sizes for items that are not used frequently may degrade performance because more resources are allocated to the caches, but are not being used.

High Cache Allocation for Less Frequently Used Large Objects

If you review the Cache Usage report and notice the number of requests to a cache are low, you may be able to reduce the size of the cache values using the Performance Tuning Utility, or manually decrease the values in the `install_dir/properties/noapp.properties` file. The low request number indicates that the objects in the cache are not used frequently with your business processes. Review the value for the cache property and reduce the cache size if the cache size is large.

For more information about increasing the cache values in the Performance Tuning Utility, see *Viewing or Editing Performance Configuration Settings* on page 46.

For more information about increasing the cache values manually, see *Manual Performance Tuning* on page 53.

The trade off in decreasing cache sizes is that decreasing the cache sizes too much may cause a reduced number of hits to the caches, which causes a call to the database for the data and increases the processing times. Depending on your system and your configuration, you may need to adjust the cache settings to attain peak performance without losing the benefit of cache over disk.

Low Cache Allocation for More Frequently Used Small Objects

If you review the Cache Usage report and notice the number of requests to a cache are high and the number of hits to the cache are low, you may be able to increase the size of the cache values using the Performance Tuning Utility, or manually increase the values in the *install_dir/properties/noapp.properties* file. The high request number indicates that the objects in the cache are used frequently with your business processes. Review the value for the cache property and increase the cache size if the cache size is small. This is especially important for the smaller objects that are used frequently and are static in value. Cache retrieval is faster than disk or database retrieval.

For more information about increasing the cache values in the Performance Tuning Utility, see *Viewing or Editing Performance Configuration Settings* on page 46.

For more information about increasing the cache values manually, see *Manual Performance Tuning* on page 53.

The trade off in increasing cache sizes for smaller objects is that increasing the cache sizes too much may cause a reduced number of hits to the caches for other caches for larger objects. This causes a call to the database for the data and increases the processing times. Depending on your system and your configuration, you may need to adjust the cache settings to attain peak performance without losing the benefit of cache over disk.

Improving Business Process Execution Times

You may find that your business process execution times increase after you run Gentran Integration Suite. An increase in volume or improperly written business processes may cause longer than expected business process execution times.

This section covers the following topics:

- ◆ *Symptoms and Causes of Poor Business Process Execution Times* on page 112
- ◆ *Resolving Business Process Execution Times* on page 113

Symptoms and Causes of Poor Business Process Execution Times

Symptoms of increasing business process execution times may include:

- ◆ Performance Statistics report showing Min (ms), Max (ms), and Average (ms) times increasing for business processes, activities, and internal system activities.
- ◆ Increasing business process execution times on the wf.log in the *install_dir/logs* directory.

Determining the cause of increasing business process execution times may require you to investigate many areas of the system and how you are implementing Gentran Integration Suite.

Typically, the cause of increasing business process execution times is either the database being full or receiving too many requests, or improperly designed business processes.

Improperly designed business processes can involve the following issues:

- ◆ Business process persistence levels set too high, which causes too much data to be stored to the database.
For more information about persistence levels, see *Changing Persistence Levels* on page 68.
- ◆ Unnecessary service invocations, which uses resources that could be used by other business processes or other components of Gentran Integration Suite.

- ◆ Improperly tuned cache properties, which cause the business process to query the database or disk for information that could be provided in the queue.

For more information about caches, see *Resolving Cache Issues* on page 108.

- ◆ Database full or database maintenance issues, which cause longer query times and increase execution times.

For more information about caches, see *Resolving Database Issues* on page 89.

- ◆ Unnecessary loops, which cause the use of resources that could be used by other business processes or other components of Gentrans Integration Suite.
- ◆ Improperly written XPath statements, which use resources that could be used by other business processes or other components of Gentrans Integration Suite.

Use relative paths in writing XPath statements. For example, write `PurchaseOrder/text` instead of `/ProcessData/PurchaseOrder/text()`. In addition, do not use `//` at the beginning of an XPath statement; this causes the entire process data to be searched and slows down the search time significantly.

- ◆ Running services in modes other than literal mode. Literal mode is significantly faster than the other modes available.
- ◆ Using older services and adapters instead of newer more efficient ones, which may combine the functionality of more than one of the older services or adapters into one service.

Note: Adapters and services that are retiring are placed on the Retiring stencil in the Graphical Process Modeler. Adapters and services on the Retiring stencil are available for use, however, if you are designing new business processes, use the newer adapters and services in your business processes.

To determine the cause of increasing business process execution times:

- ◆ Review the `wf.log` file in the `install_dir/logs` file for time and date stamp information showing increases in execution times, and errors for each business process.
- ◆ Review the Performance Statistics on the **Operations > System > Performance > Statistics** page. This report shows the number of invocations and the processing time for each business process and system activity.
- ◆ Review business processes that show increasing execution times for improper design, mode, or XPath statements.

For example, using a Command Line adapter to retrieve data instead of the File System adapter. The Command Line adapter uses more resources than the File System adapter, which can accomplish the same goal in some situations. Or, in some situations, you can create an XPath statement in an Assign service that replaces the need for both the Command Line adapter and File System adapter.

Also, use relative paths in XPath statements. For example, write `PurchaseOrder/text` instead of `/ProcessData/PurchaseOrder/text()`. In addition, do not use `//` at the beginning of an XPath statement; this causes the entire process data to be traversed.

Resolving Business Process Execution Times

The following list provides resolutions for increasing business process execution times, depending on the identified cause:

Database is Full or Receiving Many Requests

If you find that all of your business processes are experiencing increases in execution times, you may find that your database is full or receiving many requests, which causes all processing to slow down.

For more information about resolving database issues, see *Resolving Database Issues* on page 89.

Improperly Designed Business Processes

If you find that your business processes are not designed properly, your XPath statements are not worded correctly, or you are running services in other than literal mode, edit your business processes to improve processing times.

Resolving Memory Issues

Proper requirements planning and performance tuning should prevent most memory issues. However, because it is not possible to anticipate every situation, you may occasionally experience a memory problem. Some of the more common issues are covered in this section. Whenever adjusting memory allocations after initial performance tuning, you should normally increase the memory amount by the smallest increment that resolves the issue.

This section covers the following topics:

- ◆ *Correcting Dropped Data Problems* on page 115
- ◆ *Correcting Out of Memory Errors* on page 116

Correcting Dropped Data Problems

Because Gentran Integration Suite interacts extensively with other systems, sometimes problems with dropped data occur due to configuration settings of another system.

For example, When Gentran Integration Suite is under sufficient load and requires business processes to be placed on the queue, the HTTPSyncSend adapter attempts to respond with an MDN. If this does not complete in a timely manner, the client that initiated the session may release the connection due to a timeout. If this occurs, the HTTPSyncSend operation appears to complete successfully by handing off to perimeter servers, but the data may be dropped.

In this example, the timeout settings in the client could be adjusted to resolve the problem. Or, the amount of memory available to Gentran Integration Suite could be increased to allow the operation to be pulled from the queue and completed before the timeout occurs.

To increase the amount of memory available to Gentran Integration Suite, use the Performance Tuning Wizard and change the **Physical memory (MB) allocated to** Gentran Integration Suite field to a higher value. For information, see *Viewing or Editing Performance Configuration Settings* on page 46.

Correcting *Out of Memory* Errors

Occasionally, a specific operation may require more memory than the current configuration provides to Gentran Integration Suite. For example, with the default settings, attempting to view a very large code list (over 100,000 entries) may in an Out Of Memory error. This type of error is typically written to the noapp.log file.

To increase the amount of memory available to Gentran Integration Suite, use the Performance Tuning Wizard and change the **Physical memory (MB) allocated to** Gentran Integration Suite field to a higher value. For information, see *Viewing or Editing Performance Configuration Settings* on page 46.

Requirements Planning Worksheet

Completing the following planning worksheet prior to completing your performance tuning is optional, but the worksheet can assist you in tuning your Gentran Integration Suite for performance. Use the Requirements Planning Worksheet to help determine your hardware needs according to your processing volumes, operating system, hardware vendor, and number of CPUs you are running.

Request	Description	Your Answer
How many processing days are in a month?	Enter the number of processing days in a month. Use 30 days as an average. This gives 5 days of maintenance time throughout the year.	
How many processing hours are in a day?	Enter the number of hours you allow to process all data.	
How many translations are required for each transaction?	Enter the largest number of translations you will have, if there multiple translations required for a business process. For example, is data mapped to a common format, and is the common format used to interface with all external systems.	
Is content based routing required?	If content-based routing is required in your business processes, you need maps, XML encoder, BPML choice, assigns, adapters that are used for routing, and possibly other components.	
Does the content of a transaction need to be verified?	This would require translating data in to XML. The data would then need to be put in process data and verified using XPath.	
What percentage of the content of the transaction needs to be verified?	Typically, only a small amount of a transaction needs to be verified. However, if most of a transaction requires verification, you may need significant resources depending on the verification path.	
How many rules need to be applied to verify the content?	This is used to calculate the Service Latency.	
How many content elements need to be verified?	Enter the number of elements that need to be verified.	

Request	Description	Your Answer
What is the future system multiple?	Current daily volume increment in the future. The default for the current volume is 1. For example, 2 means 100% growth or doubling the system requirements.	
What is the business process persistence level required?	Enter the level of persistence (storage to the database) you require for your business process. Choices include: <ul style="list-style-type: none"> ◆ PERSISTENCE_FULL=2 – Saves a complete copy of process data along each step of the process. ◆ PERSISTENCE_MINIMAL=1– Saves all steps in a business process and selected copies of process data. ◆ PERSISTENCE_NONE=0 – Saves the first and last steps of a business process, any steps with an override persistence level, and selected steps in business process data. 	
What is the average business process size (KB)?	Unless you are very experienced and understand business process context, enter 3.5.	
What is the hardware vendor that hosts Gentran Integration Suite?	Enter the hardware vendor name. For example: IBM, Sun, HP (Compaq), or Dell.	
What is the speed of CPUs in the host?	Enter the speed in MHz. For example: 450, 1900 (1.9GHz).	
What is the number of CPUs in the host?	Enter the number of CPUs installed or the number you plan to install.	
What is the amount of main memory in the host (RAM)?	Enter in the amount of memory in MB. For example: 500, 2048MB.	
What is the operating system you are using to host Gentran Integration Suite?	Enter the host operating system (OS). For example, Solaris, Windows, Windows Server 2003, AIX, or Linux.	
What is the operating system release and version number?	For example, AIX use 4.3.3.10 (Release 4.3.3 Maintenance patch 10), Server Windows 2000 use 2000SP3.	
What are the Java Virtual Machine (JVM) version numbers?	For example, j2sdk-1_3_1_00-windows-i586. Your JVM version must match the requirements for your version of Gentran Integration Suite. The build date and lower release numbers must also match.	
What is the Gentran Integration Suite version and patch numbers?	Enter the version and patch number for your copy of Gentran Integration Suite.	

Request	Description	Your Answer
Are you integrating with an application server? If so, what is the vendor?	Enter the application server to integrate with. For example, JBoss or WebLogic.	
What is the application server version and patch number?	Enter the version and patch numbers for the application server.	
What is your relational database vendor?	Enter the relational database vendor name Gentran Integration Suite uses. For example, MYSQL, Oracle, or DB2.	
What is the relational database version and patch number?	Enter the version and patch number of the relational database.	
What are the other business applications that need to be integrated with?	Enter the vendor name of the other business applications that you use in your business. For example, SAP or PeopleSoft. This helps determine which adapters and services you need to complete processing.	
What are the other business applications version and patch numbers?	Enter the version and patch number of the other business applications.	
What data format do you exchange data with the other business applications?	Enter the different data formats that you use to exchange data with the other business applications. For example, XML or IDOC. This helps determine the amount of translation processing that may be needed that increases system resource requirements.	
What type of inbound transactions do you have?	Enter the types of inbound transactions that Gentran Integration Suite will process.	
What is the average size of each inbound transaction type?	Enter the average size of your inbound transactions that Gentran Integration Suite will process. Base your planning on the larger transactions at peak times for better performance.	
How many of each inbound transaction types do you receive daily?	Enter the sum total of the number for each inbound transaction type. This helps in determining the processing volumes.	
What is the sum total size of each inbound transaction type?	Enter the sum total size of the number for each inbound transaction type. This helps in determining the processing volumes.	
What is the average number of files in each inbound transaction by type?	Enter the average number of files in each inbound transaction type. This helps determine processing volumes.	

Request	Description	Your Answer
What type of outbound transactions do you have?	Enter the types of outbound transactions that Gentran Integration Suite will process.	
What is the average size of each outbound transaction type?	Enter the average size of your outbound transactions that Gentran Integration Suite will process. Base your planning on the larger transactions at peak times for better performance.	
How many of each outbound transaction types do you receive daily?	Enter the sum total of the number for each outbound transaction type. This helps in determining the processing volumes.	
What is the sum total size of each outbound transaction type?	Enter the sum total size of the number for each outbound transaction type. This helps in determining the processing volumes.	
What is the average number of files in each outbound transaction by type?	Enter the average number of files in each outbound transaction type. This helps determine processing volumes.	

A

- active controllers 19
- active threads 19, 23
- activities
 - monitoring 15
 - stopping 16
- AE_ExecuteCycle 54, 61
- AE_ExecuteCycleTime 54, 61
- alerterlogger.log 12
- AllowStealing 56, 63
- analyzing log files 7
- Archive Manager 16
- archive.log 12
- archiving 91
- ASI
 - performance tuning 54
 - queue tuning 60
 - queues 59
- Authentication.log 14

B

- Backuplogs.log 15
- BasicScheduling policy 58
- BLOB data 79
- BPREcovery business process 92
- business process
 - queue usage 19
 - usage 19
- business process execution times
 - causes, of increasing 112
 - determining cause 113
 - improving 112
 - resolutions, of increasing 113
 - symptoms, of increasing 112

C

- cache
 - determining cause of inefficiency 109
 - resolutions to inefficiency 110
 - resolving issues 108
 - symptoms of inefficiency 109
- cache contents, compressing 66
- cache usage 23
- CacheThreshold 55, 61
- capacity planning 117
- causes
 - database issues 90
 - halted, waiting, or interrupted business processes 104
 - increasing execution times 112
 - slow system 80
- cdinterop.log 12
- cdinterop_cdjava.log 12
- ceuinterop.log 12
- changing persistence level 68, 91
- classpath 19
- cluster environment and JNDI 23
- compressing cache contents 66
- connection, database 92
- console listener, enabling 71, 72
- consume business process 16
- controllers 19, 23
- corbaadapter.log 12
- creating performance statistics report 43

D

- data bursts 93
- data retention, optimizing 92
- database

- causes 90
- determining cause 90, 93
- increasing pool size 93
- optimizing 92
- resolutions 91, 93
- symptoms 89, 93

database connections, failing 92

database usage 19

delete.log 13

deleting performance statistics report 44

desired global threads 48

determining cause

- cache, inefficient 109
- database issues 90, 93
- halted, waiting, or interrupted business processes 105
- increasing execution times 113
- slow system 81

directory naming convention 6

disabling perimeter server 25

disk cache size 48, 54

DiskCachePath 54

document body serialization threshold 67

DTDs

- repositories 71
- trusted domains 71

E

ebXML.log 13

EDI 51

EDIINT.log 13

EnableDeadlines 55, 62

enabling

- console listener 71, 72
- perimeter server 25

envelopes 51

enveloping.properties.in file 73

environment troubleshooter page 19

event.log 13

execution times, improving 112

execution, business process 50

F

FairShareScheduling policy 58

ftp.log 13

ftpclient.log 13

G

global thread limit 59

H

halted, waiting, or interrupted business processes

- determining cause 105
- symptoms 104

http.log 13

httpclient.log 13

I

improving execution times 112

indexing 91

InitialCycles 56, 63

installation path 23

internal system statistics 38

interrupting threads 16

J

JavaPriority 55, 62

jdbc.properties.in file 67, 94

JNDI 22

- contexts, performance tuning 68
- tree structure 19

JNDI and cluster environment 23

JVM memory 49

L

large BLOB data 79

latency 29

lifecycle.log 13

listenerStartup.properties.in file 71

- load 29
- Lock Manager 16
- log file
 - naming convention 6
 - types 11
- logs
 - analyzing 7
 - changing system 70
 - naming conventions 6
 - overview 5
 - types 11
 - viewing contents 6

M

- mailbox.log 13
- managing
 - performance statistics 33
 - system logs 5
- maps 51
- MaxPoolSize 55, 59, 61
- MaxThreads 54, 60
- MaxWaitTime 56, 62
- MemCacheSize 54
- MemCacheThreshold 54
- memory in use 23
- messages
 - monitoring 16
 - searching 17
- MinPoolSize 55, 59, 62
- modifying log levels 7
- monitoring
 - activities 15
 - messages 16
 - node status 26
 - operations 5
 - threads 15

N

- naming conventions, logs 6
- noapp.log 14, 72
- noapp.properties 60, 63, 69

- noapp.properties.in file 71
- node monitoring status 26
- non-transactional pool 49

O

- oftp.log 13
- ops.log 12
- ops_exe.log 14
- opsSecurity.log 12
- opsServer.log 12
- optimizing
 - data retention 92
 - database 92

P

- performance concepts 29
- performance statistics
 - creating report 43
 - deleting report 44
 - internal system 38
 - managing 33
 - reading report 34, 98
 - statistic names 35
 - turning on/off 43
 - viewing 33, 98
- performance tuning
 - ASI 54
 - calculated recommendations 45
 - document body serialization threshold 67
 - editing, recommended settings 46
 - JNDI contexts 68
 - manual 53
 - methodology 30
 - persistence level 68
 - prerequisites 30
 - queues, ASI 60
 - system logging 70
 - utility 45
 - viewing configuration 46
- perimeter server
 - area on troubleshooting page 25
 - disabling 25
 - enabling 25

- viewing settings 25
- Perimeter.log 13
- persistence level, changing 68, 91
- persistence_level 57
- pipeline.log 13
- planning worksheet 117
- pools
 - increasing size 93
 - non-transactional 49
- produce business process 16
- property files 74
- PSLogger 15
- purging 91

Q

- QueueDepth 54, 61
- queues, ASI 59

R

- report.log 13
- requirements planning worksheet 30, 117
- rescheduling 55, 62
- ReschedulingInterval 56, 62
- resolutions
 - cache issues 108
 - cache, inefficient 110
 - database issues 91, 93
 - halted, waiting, or interrupted business processes 103
 - increasing execution times 113
 - slow system 81
- ResourceAllocation 56, 63
- RMIAPrime 25
- RMIAPrime services 20
- mif.log 13
- run.sh 51, 67, 68

S

- sap.log 14

- scalability 29
- ScheduleBackup.log 14
- ScheduleStopOps.log 15
- scheduling policy
 - changing 58
 - overview 58
- SchedulingPolicyName 54, 58
- schemas
 - repositories 71
 - trusted domains 71
- searching messages 17
- security.log 14
- serialization, document body threshold 67
- servers 23
- servicesctl.log 12
- setupfiles.cmd 51, 63, 67, 68, 69, 71, 72, 73
- setupfiles.sh 51, 63, 67, 68, 69, 71, 72, 73
- si_exe.log 14
- slow system
 - causes 80
 - determining cause 81
 - resolutions 81
 - resolving 80, 115
 - symptoms 80
- sql.log 14
- start time 19, 20
- starting system logs 7
- StartupPS.log 15
- startWindowsService.cmd 52, 67, 68
- state 23
- status viewing 5
- StealThreads 56, 63
- stopping
 - activities 16
 - system logs 7
 - threads 16
- symptoms
 - cache, inefficient 109
 - database issues 89, 93
 - halted, waiting, or interrupted business processes 104

- increasing execution times 112
 - slow system 80
- SyncEngine.log 14
- system classpath 22
- system information 18
- system logs
 - managing 5
 - modifying log level 7
 - turning on and off 7
- system logs, starting 7
- system logs, stopping 7
- system status 19
- system troubleshooting 18
- system troubleshooting page 18
- system.log 13, 14

T

- thread dump
 - conducting 82
 - reading 87
- Thread Monitor 15
- threads
 - interrupting 16
 - monitoring 15
 - stopping 16
 - viewing 15
- throughput 29
- tracking.log 13
- transactional pool 49
- translation maps 51
- troubleshooting
 - methodology 32
 - overview 5
 - system 18
- trusted domains
 - DTDs 71
 - schemas 71
- tuning, ASI queue 60
- tuning.properties file 64, 68
- txtrace.log 14

U

- ui.log 13
- up time 19, 20

V

- viewing
 - database statistics 98
 - log files 6
 - performance statistics 33, 98
 - performance tuning configuration 46
 - perimeter server settings 25
 - server status 5
 - service status 5
 - threads 15

W

- webx.log 14
- wf.log 12
- wfexception.log 12
- wfstatistics.log 12
- workflows.properties file 66

X

- XAPI output 79