

Sterling Integrator



Performance Management

Version 5.1

Sterling Integrator



Performance Management

Version 5.1

Note

Before using this information and the product it supports, read the information in "Notices" on page 211.

Copyright

This edition applies to Version 5 Release 1 of Sterling Integrator and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 2000, 2014.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Chapter 1. Overview 1

Intended Audience	1
System Components	1
Performance Concepts Overview	2
Performance Tuning Methodology	4

Chapter 2. Performance Recommendations Checklist 7

Sterling Integrator: General Recommendations Checklist	7
Sterling File Gateway: Specific Recommendations Checklist	10
EBICS Banking Server Module: Specific Recommendations	13

Chapter 3. Database Management System 17

Planning	17
Server Sizing	17
Storage and File Systems	17
Database Management for Sterling Integrator	19
Dynamically Manage JDBC Pools	23
Oracle Database Configuration and Monitoring	26
init Parameter Configuration Checklist	26
Rolling Back or Undoing Changes in Oracle Database	28
Database Storage	29
Monitoring with the AWR Report	31
Index and Table Statistics	34
Enable BLOB Caching	34
Oracle Real Application Clusters (RAC)	35
IBM DB2 for LUW Configuration and Monitoring	37
Recommended Settings for IBM DB2 Registry Variables	37
Recommended Settings for DBM CFG Parameters	40
Recommended Settings for DB CFG Parameters	41
Recommended Settings for DB2 9.7	43
Storage Subsystem	45
Log Configuration	48
DB2 Monitoring	49
DB2 Index and Table Statistics	51
Microsoft SQL Server Configuration and Monitoring	52
Recommended Instance-Specific Settings for Microsoft SQL Server	52
Microsoft SQL Server Memory with Address Windowing Extensions (AWE)	53
Recommended DB-Specific Settings for Microsoft SQL Server	53
Row Versioning-Based Isolation Level (READ_COMMITTED_SNAPSHOT) for Microsoft SQL Server	53
Recommended Settings for Microsoft SQL Server 2008	54
Storage Subsystem	56

Monitoring Microsoft SQL Server Using Microsoft SQL Server Management Studio	57
Monitoring Microsoft SQL Server Using SQL Server Profiler	57
Microsoft SQL Dynamic Management Views	58
Microsoft SQL System Monitor	58
Microsoft SQL Server Maintenance Plan	58
Microsoft SQL Server Index, Table Statistics, and Index Rebuilds	59
Windows Defragmentation	60
Microsoft SQL Server Tips	60

Chapter 4. Java Virtual Machines 63

Overview of JVM	63
View the Performance Tuning Configuration Settings for ASI Nodes	63
View JVM Settings in the tuning.properties file	64
Edit Performance Tuning Configuration Settings for ASI Node	64
Edit Performance Tuning Configuration Settings for Container JVMs	65
Garbage Collection Statistics	65
JVM Verbose Garbage Collection	66
IBM JVM Garbage Collection Example	66
Introduction to HotSpot JVM Performance and Tuning	67
HotSpot JVM Heap Memory and Garbage Collection	67
HotSpot JVM Garbage Collection Tools	68
HotSpot JVM Startup and Runtime Performance Optimization	69
HotSpot JVM Default Parameters	69
Enable JVM Verbose Garbage Collection	73
Display HotSpot Statistics for HP JVM	74
Enable Verbose Garbage Collection for Container JVM	74
Display HotSpot Statistics for Sun JVM	75
HotSpot JVM Heap Monitoring	76
HotSpot JVM Thread Monitoring	76
Performing a Thread Dump	78
HotSpot Thread Dump Analysis Tools	78
HotSpot JVM DeadLock Detection	79
HotSpot JVM Blocked Thread Detection	79
HotSpot JVM Troubleshooting Tips	79
Introduction to the IBM JVM Performance and Tuning Guidelines	84
Before You Begin Tuning Your IBM JVM	85
IBM JVM Default Parameters for Sterling Integrator	85
IBM JVM Troubleshooting Tips	87

Chapter 5. Monitoring Operations 91

Monitoring Operations	91
Managing System Logs	91
Auditing	102

Monitoring a Business Process Thread	104
Monitoring Messages	105
Reviewing System Information	107
Monitoring Node Status	115
Monitoring Deprecated Resources	117
Chapter 6. Performance Tuning.	121
Performance Tuning Utility	121
Calculating the Settings Required.	121
Viewing or Editing Performance Configuration Settings	123
Manual Performance Tuning	128
Tuning Your ASI Environment.	129
Scheduling Policy	133
Queues in an Application Server Independent Environment	133
Monitoring Queues using Queue Watcher	137
Compress Cache Contents	140
Set the Document Body Serialization Threshold Property	141
Allocate JNDI Contexts	142
Changing Persistence Levels	142
Changing System Logging	145
Enable Trusted Domains for Schemas	145
Enable the Console Listener	145
Enable the JMX Event Listener	146
Enhance EDI Performance	147
Customizing Property Files	148
Managing System Recovery	148
Change the SoftStop Time Limit	150
Changing Advanced File Transfer Settings.	150
Optimize System Performance for Sterling e-Invoicing	152
The Order of Segments in Maps and Data	153
Purge Performance Recommendations	154
Optimizing and Controlling the System Threads	155
Chapter 7. Performance Statistics	161
Manage Performance Statistics.	161

Turning On and Turning Off Performance Statistics	171
Reporting Performance Statistics	171

Chapter 8. Troubleshooting 173

Database Maintenance Check	173
Full Database Issues and Resolution.	176
Database Connection Issues	180
Database Down Check	182
Tracking JDBC Connections	183
Types of Cache Memory.	184
Symptoms and Causes of Inefficient Cache Usage	185
Resolving Inefficient Cache Usage	185
Correcting Dropped Data Problems	187
Correcting Out-Of-Memory Errors	187
Understanding Business Processes	187
Symptoms and Causes of Poor Business Process Execution Time.	188
Resolving Halted, Halting, Waiting, or Interrupted Business Processes.	189
Slow System: Symptoms, Causes, and Resolution	191
Conducting a Thread Dump	193
Scenario-Based Troubleshooting Tips and Techniques	196

Chapter 9. Appendix A 199

The Performance and Tuning Worksheet	199
--	-----

Chapter 10. Appendix B 203

workflowLauncher: Running a Business Process from a Command Line	203
--	-----

Index 207

Notices 211

Chapter 1. Overview

Intended Audience

This document is intended for, but not limited to:

- Technical Architects
- Performance Engineers
- Configuration Engineers
- Application Administrators
- Database Administrators
- System Administrators

System Components

Performance management activities affect all aspects of the system, including computer nodes, network, disks, database, and so on.

One person or role may be responsible for one, several, or all the components. Some of the roles include:

- Hardware Engineer
- System Administrator
- Network Engineer - Local Area or Wide Area
- Application Server Administrator
- Database Administrator
- Capacity Planning Engineer
- Performance Analyst
- Sterling Integrator Administrator

Performance management documentation includes the following information:

- Background information about the different performance and tuning issues when running Sterling Integrator.
- Step-by-step information that helps you:
 - Optimize the performance of Sterling Integrator.
 - Diagnose and resolve performance issues, if any, with Sterling Integrator to suit your environment.

You can work through most performance and tuning issues using the following tools:

- Properties file settings, which you can access through the properties directory in your installation directory.
- The Sterling Integrator user interface. Use the Operations option in the Administration menu to access the following tools:
 - JDBC Monitor (**Operations > JDBC Monitor**)
 - JVM Monitor (**Operations > System > Performance > JVM Monitor**)
 - Message Monitor (**Operations > Message Monitor**)
 - Performance Statistics Report (**Operations > System > Performance > Statistics**)

- Performance Tuning Wizard (**Operations > System > Performance > Tuning**)
- SWIFTNet Monitor (**Operations > SWIFTNet Monitor**)
- System Troubleshooting (**Operations > System > Troubleshooter**)
- Thread Monitor (**Operations > Thread Monitor**)

Based on the diversity of the roles and the responsibilities associated with them, the *Sterling Integrator Performance Management Guide* is divided into the following sections:

- Overview - Provides performance tuning overview, audience information, performance concepts overview, performance tuning methodologies, and tuning issues that may arise when running Sterling Integrator.
- Performance recommendations checklist - Provides general, planning, and implementation (single node and cluster) checklists. This topic caters to hardware engineers, system administrators, network engineers, capacity planners, and Sterling Integrator administrators.
- Database management system - Discusses key recommendations for Oracle®, IBM DB2®, Microsoft SQL Server™, and MySQL databases. This topic caters to database administrators.
- Java Virtual Machines - Explains configuration, recommendations, and so on. This topic caters to application server administrators.
- Monitoring operations, performance tuning, and performance statistics - Discusses monitoring operations, performance tuning (utility and manual), and management of performance statistics.
- Monitoring operations - Provides information about system logs, auditing, business process threads, messages, system information, and cluster node status.
- Performance tuning - Explains performance tuning using the performance tuning wizard, and manual performance tuning recommendations such as scheduling policies, cache contents, persistence levels, property files, and system recovery.
- Performance statistics - Provides information about managing, enabling and disabling, and reporting performance statistics.
- Troubleshooting - Provides information about resolving full database issues, Java Database Connectivity (JDBC™) connections, cache issues, memory issues, slow system issues, and improving business process execution time.
- Performance and tuning worksheet - Helps you to take an inventory of the different parts of your Sterling Integrator installation that affects performance.
- workFlow Launcher - Explains how to run business processes from the command line.

Note: It is recommended that the performance analyst, capacity planner, and Sterling Integrator administrators read all the topics.

Performance Concepts Overview

Typically, performance in general, and optimal performance in particular, are associated with the following criteria:

- Latency – The amount of time taken to generate a response to a request (speed).
- Throughput – The amount of data transferred during a specified period (volume).
- Scalability – The ability of the system to adapt to increasing workload (additional hardware).

- Load – The ability of the system to continue performing at optimal level even when the load on the system increases.

Sterling Integrator's performance can be tuned to meet various processing requirements, including higher processing speed and ability to sustain high volumes. The amount of resources given to the interactive and batch components of a mixed workload determines the trade-off between responsiveness (speed) and throughput (work completed in a certain amount of time).

When using the Sterling Integrator, if you face any performance issues, perform the applicable task or tasks from the following list:

- Change performance parameters in the properties files or through the performance tuning utility. For more information about changing performance parameters, refer to the topic Performance Tuning Utility.
- Add additional hardware.
- Tune your business processes to run more efficiently on Sterling Integrator.
- Monitor and archive the database to free up resources.
- Create Sterling Integrator cluster for load balancing and scaling.

Before You Begin Performance Tuning

Before you carry out performance tuning actions, you must consider capacity planning issues. The "Performance and Tuning Worksheet" topic provides information about how to determine your capacity requirements. This worksheet, and other capacity planning tools, also help you adjust your current workload, regardless of your future requirements.

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 that are available and can be dedicated to Sterling Integrator.
- Memory available to meet your processing requirements.
- Disk space available.
- Hard disk Redundant Array of Independent Disks (RAID) level. RAID arrays use two or more drives in combination for fault tolerance and performance. The recommended RAID level for Sterling Integrator is Level 5.
- Database size requirements.

Note: When conducting a performance tuning activity, keep the information provided in this topic at hand for easy reference and future planning.

Performance Tuning Methodology

For effective performance tuning, you must first identify the performance issues and determine the corrective actions. Following is a methodology that helps you focus your efforts and reduce redundant tasks.

Use the following methodology to identify and resolve performance issues in Sterling Integrator:

1. Fill out the Requirements Planning Worksheet to determine the hardware requirements for a given processing volume and speed.

Note: If you purchased Sterling Commerce Consulting Services, you would have received a report containing information about your hardware requirements based on your business needs.

2. Verify that your hardware and memory specifications match the hardware and memory recommendations provided during the performance planning phase.
3. Verify that you have installed the requisite patches on your system, which will help you fix your performance issues from the Sterling Commerce Customer Center Web site at <https://cn.sterlingcommerce.com/login.jsp>.
4. Verify that you have the supported Java Virtual Machine (JVM™) on the computer running Sterling Integrator, and on the DMZ computer if you are running perimeter servers. Both the JVM versions must match each other and your Sterling Integrator version requirements. 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 JDBC drivers.
7. Verify that you have created your business processes using the most current adapters and services, well-structured XPath statements, and the lowest persistence by step and business process.
8. Tune Sterling Integrator using the Performance Tuning Utility and the information that you provided in the Requirements Planning Worksheet (refer to the topic Performance and Tuning Worksheet). The Performance Tuning Utility enables you to tune the cache, memory, and other system components, but not the business processes.

For more information about this utility, refer to the topic Performance Tuning Utility.

9. Navigate to **Operations > System > Performance > Statistics**. In the **On/Off** section of the **Performance Statistics Report Manager** page, select the **Enable Performance Statistics** option. The Performance Statistics reports provide information about the time taken to perform business processes and execute activities, database connections and persistence, and business process queue performance. You can use this information to tune your system according to your volume and speed requirements.

For more information about performance statistics, refer to the topic Manage Performance Statistics.

10. Review your history of incoming and outgoing documents and enter this information in charts, showing daily, weekly, and monthly processing trends. Use these charts to determine your peak volume processing levels based on your processing pattern history.
11. Conduct initial performance testing by running your business processes with sample data that is the same size as the data that will be processed in production. In addition, run your business processes with data that

approximates your anticipated peak processing volume. This helps you tune your system as close 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 information about the areas that can be tuned.
14. Retune Sterling Integrator using the Performance Tuning Utility, based on the information you gathered from your initial performance testing.
Continue this process until your processing time and volume requirements are met.
15. Create a new Performance Statistics Report called *Benchmarksdd/mm/yy*.
16. Conduct the same test that you conducted in step 11.
17. Review the *Benchmarksdd/mm/yy* Performance Statistics Report. If the statistics in this report are not similar to your previous statistics, repeat steps 11 - 14.
18. Compare your monthly or weekly Performance Statistics Reports with 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.

Chapter 2. Performance Recommendations Checklist

Sterling Integrator: General Recommendations Checklist

The general recommendations checklist provides a list of guidelines to plan for the required elements, and to enhance the performance of Sterling Integrator.

In the following table, the Test and Production columns indicate whether the recommendations are Recommended (R), Critical (C), or Not Applicable (NA) in the test and production environments.

Note: It is recommended to setup a Sterling Integrator test environment with a sample set of data to verify the recommendations provided in this checklist.

Recommendation	Test	Production	Comments
OS version and OS kernel parameters	C	C	You should ensure that you install Sterling Integrator on certified OS versions and levels. Refer to the <i>System Requirements</i> documentation of the corresponding OS versions.
Network speed	C	C	You should ensure that your network cards are operating at the highest speeds. The network interface and the network switch can negotiate to lower speed. When that happens, performance degrades even under normal processing periods.
AIX page space allocation	C	C	The AIX default page space allocation policy does not reserve swap space when processes allocate memory. This can lead to excessive swap space, which forces AIX to kill processes when it runs out of swap space. You should ensure that you either have sufficient swap space, or set the following environment policy variables: <i>PSALLOC=EARLY</i> <i>NODISCLAIM=TRUE</i>
Monitor CPU Utilization	NA	C	You should monitor CPU utilization to ensure that there is no CPU contention.
Monitor Swap Usage	C	C	If not enough space is left on the swap device (or paging file), the operating system may prevent another process from starting, or in some cases, be forced to kill the running processes.

Recommendation	Test	Production	Comments
Monitor Paging	C	C	<p>The JVMs and database management systems rely on large memory buffers or heaps, and are sensitive to paging. Performance can noticeably degrade if enough memory is not available to keep the JVM heap in memory.</p> <p>You can monitor paging levels using standard operating system or third-party measurement tools, such as:</p> <ul style="list-style-type: none"> • UNIX/Linux – SAR • Windows – System Monitor
Monitor Heap Garbage Collection Performance	C	C	<p>Monitoring heap GC performance is critical for performance and availability. For example, if the amount of heap that is free after a GC is continually increasing, and approaching the maximum heap size, the JVM can experience OutOfMemory exceptions.</p>
Planning			
Server node sizing	NA	C	<p>You should ensure that you have sufficient computing capacity to process peak transaction volumes.</p> <p>Refer to <i>System Requirements</i> documentation to get an estimate of the processor, memory, and other requirements.</p> <p>You can also engage Sterling Commerce Professional Services to conduct a capacity plan study of your system. This study involves measuring your system, and using the measurements to forecast resource requirements at anticipated peak processing periods.</p>
Database disk sizing	NA	C	<p>You should ensure that you have sufficient disk space for the database server.</p> <p>The size of the database disk subsystem may vary from a few gigabytes to several terabytes. The size depends on the processing volume, its complexity, and the length of time you want to keep the data active in the database.</p>
JVM			
JVM version	C	C	<p>You must ensure that you install Sterling Integrator on certified JVM versions and levels.</p> <p>Run the -version command in the command prompt to ensure that you have installed the correct version.</p>
Verbose GC statistics	NA	C	<p>You can enable verbose GC statistics collection. Understanding the health of GCs for each JVM is critical for performance.</p>

Recommendation	Test	Production	Comments
Paging	C	C	The JVM heap must be resident in the memory. The performance can noticeably degrade if the operating system has to page portions of the heap out to disk.
OutOfMemory Exceptions	C	C	OutOfMemory exceptions can cause unpredictable behavior. You should ensure that you have allocated sufficient physical memory to Sterling Integrator based on your processing requirements.
JVM VM mode	C	C	For HotSpot JVM, the server mode is applicable for long-running workloads.
Heap size	C	C	<p>Correct heap size configuration is critical for both performance and availability. If the heap size is too big, the GC pauses can be long. Similarly, if the heap size is too small, it can lead to OutOfMemory exceptions. You should ensure that heap size is not set larger than physical memory to avoid thrashing.</p> <p>Sterling Integrator supports both 32-bit and 64-bit JVMs. You should allocate more heap space when you are running a 64-bit JVM.</p>
Database			
Monitor and regulate indexes	C	C	<p>Sterling Integrator comes with a default set of indexes. In some cases, the indexes may not apply to your operational environment.</p> <p>You should regularly monitor the resource cost of frequently used queries to check if additional indexes are required. Similarly, you can also monitor the database to delete indexes if they are not required.</p>
Cursor sharing	C	C	<p>If you are using Oracle database, cursor sharing enables dynamic SQL to be reusable, thereby reducing the contention on the shared pool.</p> <p>You should ensure that you set <code>cursor_sharing=EXACT</code></p>
Parameters governing DB2 locking strategy	C	C	Set <code>DB2_EVALUNCOMMITTED</code> , <code>DB2_SKIPDELETED</code> , and <code>DB2_SKIPINSERTED</code> to reduce lock contention.
Parameters governing DB2 memory	C	C	Set parameters to manage various memory structures such as <code>LOCKLIST</code> , <code>SORTHEAP</code> , and so on, to <code>AUTOMATIC</code> .
Volatile table	NA	C	Mark tables that change significantly, as volatile.

Sterling File Gateway: Specific Recommendations Checklist

Sterling File Gateway is installed on an instance of Sterling Integrator, and shares many of the resources with the latter, including:

- Communication Adapters
- Business Processes
- Security Services
- Perimeter Services
- Encryption
- Decryption
- Account Management

You should, therefore, tune your Sterling Integrator installation first, and then perform the Sterling File Gateway-specific tuning and troubleshooting tasks. Be aware that the changes you make to Sterling File Gateway can also affect the performance of Sterling Integrator.

Note: Do not edit the properties files. Make all the changes in the `customer_overrides.properties` file. For example, to change the `pgpCmdLine2svcname` property, enter the following line in the `customer_overrides.properties` file:

```
filegateway.pgpCmdLine2svcname=CUSTOM
```

In this line, replace **CUSTOM** with the name of your Command Line 2 adapter. For more information about the `customer_overrides.properties` file, refer to the *Sterling Integrator Property Files* documentation, which can be accessed from:

<http://www.sterlingcommerce.com/Documentation/SI51/PropFilesPage.htm>

The following table describes some of the key parameters that must be configured to optimize Sterling File Gateway's performance.

In the following table, the Test and Production columns indicate whether the recommendations are Recommended (R), Critical (C), or Not Applicable (NA) in the test and production environments.

Recommendation	Test	Production	Comments
Increase the value of Sterling File Gateway. Property: <code>fgRouteConcurrentSessionLimit</code>	R	R	Number of Sterling File Gateway services that can be run concurrently. The services are split into two groups, and each group has this value as the limit. Therefore, the total number of services that can run concurrently is equal to the value for this property multiplied by two. Set this to a value that is higher than the sum of business process threads in queues 4 and 6 (where Sterling File Gateway services run). Default value: 8 (Maximum: 64)

Recommendation	Test	Production	Comments
<p>If you are processing very large files, increase the probe values to avoid timeout conditions.</p> <p>Property:</p> <ul style="list-style-type: none"> • filegateway.bpCompletionProbes.2 • filegateway.bpCompletionSleepMsec.2 	R	R	<p>Timeouts and sleep intervals that control the period for which Sterling File Gateway waits for each of the sub-business process it invokes. The timeouts and sleep intervals control the timeouts when a business process is executed synchronously during routing. The types of business processes that run during routing are consumer identification and PGP processing. Setting the values for these properties also enables one set of relatively quick probes, followed by a second set of slower probes. The first set will be reactive, but consumes more processor capacity. The second set will be activated for longer-running processes and will consume less processor capacity.</p> <p>First, probe 120 times, with 100 milliseconds between each probe, for a total of 12 seconds.</p> <p>Default value:</p> <ul style="list-style-type: none"> • bpCompletionProbes.1=120 • bpCompletionSleepMsec.1=100 <p>Then, probe 600 times with 2000 milliseconds between each probe, for a total of 1200 seconds (20 minutes).</p> <p>Default value:</p> <ul style="list-style-type: none"> • bpCompletionProbes.2=600 • bpCompletionSleepMsec.2=2000
<p>If you have a high volume of PGP traffic, you can improve your performance by specifying a group for the file gateway.</p> <p>Property: pgpCmdline2svcname</p>	R	R	<p>The name of the Command Line 2 adapter to be used for PGP packaging and unpackaging. You can override this property in the <code>customer_overrides.properties</code> file if a custom Command Line 2 adapter is used for PGP operations. You can also specify an adapter group name to balance the outbound PGP sessions load across multiple adapter instances.</p> <p>Default value: pgpCmdline2svcname=PGPCmdlineService</p>
<p>If you have very large files that will be processed by PGP, increase the value of the file gateway.</p> <p>Property: fgRoutePGPCmdLineSocketTimeout</p>	R	R	<p>Timeout value, in milliseconds, for PGP package and unpackage operations invoked by Sterling File Gateway.</p> <p>Default value: 240000 milliseconds (4 minutes)</p>

Recommendation	Test	Production	Comments
<p>If you have high volumes of FTP traffic, you can improve your performance by specifying a group.</p> <p>Property: filegateway.ftpClientAdapterName</p>	R	R	<p>The FTP Client Adapter instance or service group that the FileGatewayDeliverFTP business process will use. You can override this property in the customer_overrides.properties file to use a custom FTP Client Adapter instance to contact trading partners. You can also specify an adapter group name to balance the outbound FTP sessions load across multiple adapter instances.</p> <p>Default value: ftpClientAdapterName=FTPClientAdapter</p>
<p>Decrease the value of evaluation frequency.</p> <p>Property:</p> <ul style="list-style-type: none"> • MailboxEvaluateAllAutomaticRules <p>or</p> <ul style="list-style-type: none"> • MailboxEvaluateAllAutomaticRulesSubMin 	R	R	<p>You can enable either MailboxEvaluateAllAutomaticRules or MailboxEvaluateAllAutomaticRulesSubMin.</p> <p>MailboxEvaluateAllAutomaticRulesSubMin verifies the presence of routable messages once every 10 seconds, and can be edited for other intervals of less than one minute by modifying the MailboxEvaluateAllAutomaticRulesSubMin business process.</p>
<p>Suppress Duplicate Messages</p> <p>Property: mailbox.disallowDuplicateMessages=true</p>	R	R	<p>Prevents duplicate messages from using system resources.</p>
<p>Increase the number of steps a business process must complete prior to returning to the queue.</p> <p>Property: noapp.AE_ExecuteCycle.#</p>	R	R	<p>Number of steps involved in the completion of a business process before the business process returns to the queue. Higher values will accelerate individual business process execution, while lower values will provide smoother multitasking capabilities. Interactive use favors a lower number of steps, while batch processing favors a higher number of steps. The value of noapp.AE_ExecuteCycle.# can be different for each queue. # indicates the queue number.</p> <p>When a business process has one service to begin a protocol session and another service to use the protocol session, a very low AE_ExecuteCycle may lead many business processes to be in the queue, with only the first service running. This may result in many protocol sessions accumulating in an open state, and session limits being met sooner than is necessary.</p>

Recommendation	Test	Production	Comments
Increase the time period that a business process can use a thread, before releasing it to be used for another business process. Property: noapp.AE_ExecuteCycleTime.#	R	R	Maximum time period, in milliseconds, for which a business process can use a thread before releasing it for use by another business process. This value will override the value set for AE_ExecuteCycle. Tuning the value for this property ensures that a series of unusually slow steps will not tie up a thread completely. This value can be different for each queue. # indicates the queue number. A value that is too low may result in the accumulation of more sessions than are recommended.
Increase the number of concurrent threads. Property: noapp.MaxThreads	R	R	Total number of concurrent threads that Sterling File Gateway is allowed to use. This value is the total number of threads available to a workflow engine to execute business process steps. Other, non-workflow engine threads do not come under the purview of this limit. For example, the threads set in fgRouteConcurrentSessionLimit do not come under the purview of this limit.
Set storage type. Property: filegateway_ui.storageType=file	R	R	File System is more efficient. Default value: database

EBICS Banking Server Module: Specific Recommendations

The EBICS Banking Server Module is installed on an instance of Sterling Integrator, and shares many of the resources with the latter.

You should, therefore, tune your Sterling Integrator installation first, and then perform the EBICS Banking Server Module-specific tuning. Be aware that the changes you make to the EBICS Banking Server Module can also affect the performance of Sterling Integrator.

The following table describes some of the key parameters that must be configured to optimize the EBICS Banking Server Module's performance:

Recommendation	Comments
Allocation of additional threads to the HTTP Server Adapter	<p>You can allocate additional threads to the HTTP Server adapter when the rate of concurrent requests to the EBICS Banking Server Module is moderately high. Complete the following steps to allocate additional threads to the HTTP Server adapter:</p> <ol style="list-style-type: none"> 1. Modify the <code>numOfmaxThread</code> parameter in the <code>http.properties.in</code> file by setting <code>numOfmaxThread</code> to four times the number of processor cores on the system the HTTP Server Adapter is running on. For more information about calculating the settings to tune the performance of your system, refer to the topic “Calculating the Settings Required” on page 121. 2. Run the <code>setupfiles.sh</code> (<code>setupfiles.cmd</code> for Windows) script to apply the changes.
Reduction of the disk I/O amount	<p>In order to reduce the amount of disk I/O on the system, change the persistence of the following EBICS business processes to Error Only:</p> <ul style="list-style-type: none"> • <code>handleEBICSRequest</code> • <code>EBICSOrderAuthorisationProcessing</code> • <code>EBICSOrderProcessing</code>
Storage size	<p>The database space required to store the files processed by the EBICS Banking Server Module is approximately 2.5 times the size of the transacted files. For more information about managing databases in Sterling Integrator, refer to “Database Management for Sterling Integrator” on page 19.</p>

Recommendation	Comments
Purge documents	<p>The document lifespan in an EBICS transaction is set to ten years. As a result, the business processes associated with the documents remain in the live system databases, and may occupy a large memory and slow down the performance of the system.</p> <p>The Index Business Process service scans the live systems, and flags the records that have reached their purge eligibility date and time. To reset the document lifespan, schedule the Index Business Process to run every ten minutes. The Index Business Process resets the lifespan after you delete the messages from the mailbox.</p> <p>For more information about implementing and configuring the Index Business Process service and scheduled purge, refer to Index Business Process Service and Purge Service documentation respectively, which can be accessed from http://www.sterlingcommerce.com/Documentation/SI51/HomePage.htm.</p>

Chapter 3. Database Management System

Planning

This section provides information about the tasks elements that have to be completed prior to the implementation phase. The key planning tasks include, but are not limited to:

- Selecting a certified database management server software and version. Refer to the *Sterling Integrator System Requirements* for a list of supported database configurations.
- Determining the size and configuration of the database server node
- Determining the size and configuration of the database disk subsystem
- Determining the disk technology

Monitoring Disk Space

At the operating system level, it is important to monitor disk space utilization, as this is one of the most common causes of database freeze up. On various UNIX/Linux-based platforms, run the **df** command as a line command or in a shell script. Various third-party tools are also available.

Server Sizing

At appropriate times in the project lifecycle, you can request a Server Sizing study from your Sterling Commerce Professional Services Project Manager or a Sterling Commerce Sales Executive. This study starts with the Sterling Integrator Server Sizing Questionnaire. Sterling Commerce Performance Engineering creates a sizing document that provides information about the estimated processor, memory, and network requirements for the standard/baseline Sterling Integrator. On your part, you must consider additional requirements such as other workloads on the same node (for example, additional third-party software, customization, performance monitors, and so on).

Storage and File Systems

This section covers the following topics:

- Capacity Requirements
- Use of Native File Systems
- Monitoring Disk Space

Capacity Requirements

Your disk capacity requirement is a very important input to the disk configuration planning process. This involves many considerations to ensure survivability, manageability, scalability, and performance.

The following table provides information about the tasks that you must perform to ensure that the required behavior is achieved.

Goal	Strategy
Survivability	<ul style="list-style-type: none"> • Configure disks with the ability to survive single or multiple disk failures (for example, RAID-1, RAID-5, or RAID-10). • Configure the disk array with multiple I/O paths to the server to survive I/O path failures. • Configure disks to be accessible from multiple server nodes to tolerate single-node failures.
Manageability	If you have limited-time windows to back up the database, select techniques such as array snapshots or SAN-based (storage area network) mirroring that allow logical backups.
Scalability and Performance	<ul style="list-style-type: none"> • Configure the disk array with many small disks instead of a few large disks to increase the number of I/O paths. • Configure the disk array with large NVRAM cache to improve read and write performance. • Configure the disks with striping (for example, RAID-0 or RAID-10). • In some circumstances, you may also want to consider using solid-state disks.

Capacity Requirements: An Example

Let us consider the following example to understand the concept of capacity requirements better. Assume you require 900 GB of storage and you have disk arrays or SANs that are made up of 93 GB disks. Let us further assume that the database is implemented over ninety 10 GB data files.

The following table summarizes the choices for the common disk organizations and the trade-offs involved.

Technology	Scalability	Survivability	Maintainability	Num Disks
JBOD	Poor. Subject to throughput of individual disks	Poor. Single-disk failure creates outage and requires database recovery	Poor. High disk utilization skew.	10
RAID-0	Excellent. Striping n disks provides read/write throughput at n times a single disk.	Poor. Single-disk failure creates outage and requires database recovery.	Excellent. Expect near-uniform disk utilization within a logical unit. Potential LUN utilization skew.	10
RAID-1	Poor. Similar performance to that of JBOD.	Better. Can survive multiple disk failures, assuming that these occur in different mirrored sets.	Poor. High disk utilization skew.	20

Technology	Scalability	Survivability	Maintainability	Num Disks
RAID-5	Excellent for read (similar to RAID-0). Potentially poor for write performance.	Better. Able to survive a single-disk failure. Multiple disk failures create an outage and require database recovery.	Excellent. Low disk utilization skew. Possible LUN utilization skew.	11
RAID-01	Excellent read/write performance.	Can tolerate up to two disk failures as long as both failures are not in the same mirrored set.	Excellent. Low disk utilization skew. Possible LUN utilization skew.	20
RAID-10	Excellent read/write performance.	Can tolerate up to n disk failures as long as two failures do not occur in the same mirrored set.	Excellent. Low disk utilization skew. Possible LUN utilization skew.	20

Use of Native File Systems

Using file systems for storage can simplify administration tasks, although potentially at some loss of some performance over using raw disks. It is recommended that you consult with your storage administrator and storage vendor.

Operating System	Native File System
Windows	NTFS
Linux	EXT3
Solaris	UFS
AIX	JFS2
HP-UX	VxFS

Database Management for Sterling Integrator

Sterling Integrator uses a database server as a repository for transactional, reference, and history data that it generates and uses. Refer to the *Sterling Integrator Installation Guide* for a list of supported database configurations.

This section covers the following concepts:

- JDBC Connection Pool Overview
- Configuring Parameters for a Database Pool
- Schema Objects

JDBC Connection Pool Overview

Sterling Integrator uses internal connection pooling in an effort to minimize delays in the creation of new connections to the database server. When a connection is requested, an existing connection is used whenever possible. This improves performance by removing the overhead of going to the driver and creating and destroying a new connection each time one is needed.

Sterling Integrator internal connection pooling mechanism implements a connection pool in every JVM started for the product. In a cluster with multiple nodes, the number of connections the database has to establish is the total of all connections for all the connection pools.

For example, if you have four JVM instances and each connection pool averages around 30 connections, your database server will have to establish 120 database connections.

With this implementation, the database will report any idle connections to the system. This is the expected behavior. JDBC properties can be tuned in the `jdbc.properties` file to destroy idle connections and minimize usage.

Note: Override the JDBC properties in the `customer_overrides.properties` file. Do not modify the `jdbc.properties` file directly. For more information about the `jdbc.properties` file, refer to the inline comments in the `jdbc.properties` file.

How Connection Pooling is Implemented

Each database pool is created with an initial number of connections, as defined by the `initsize` property in the `jdbc.properties` file. As more connections are required, Sterling Integrator requests additional connections up to the maximum number defined by the `maxsize` property. When Sterling Integrator finishes using a connection, it returns it to the pool of connections for later use.

If Sterling Integrator requires connections beyond the maximum size limit for each database pool, and every pool is in use, Sterling Integrator can request new connections up to the value defined by the `buffer_size` property. Connections in this “buffer range” (that is, beyond the number of connections defined by the `maxsize` property) are created as needed, and destroyed when the calling process is completed. To detect this, search for warning messages such as connections exceeded in the `noapp.log`. Performance will drop if Sterling Integrator runs in this range for a long period of time.

The `buffer_max` value is a hard maximum. When the number of connections defined by `buffer_max` (`maxsize + buffersize`) is reached, additional requests for connections are denied. An error is written to the logs indicating that the pool has been exhausted and the attempt to create a new connection has failed.

Summary

Each JVM can create connections for each of the pools defined in the `jdbc.properties` file and configured in the `customer_overrides.properties` file. Each pool grabs the initial number of connections (`initsize`) configured.

Because a pool requires additional connections beyond the initial size, it creates new connections up to the maxsize limit for that pool. Connections that are created with fewer than the maxsize are returned to that pool for reuse, thus improving performance.

If a pool is using all the connections up to the maxsize limit, it creates a new connection as needed within the buffersize limit. Connections beyond maxsize and below the buffersize are not reused. Each connection is created when needed, and destroyed when no longer needed by the calling process. This is expensive and will hit performance if Sterling Integrator runs for continued amounts of time in this state.

When the number of connections (maxsize + buffersize) is reached, new requests for connections are refused.

You can now manage JDBC pools dynamically. For more information about dynamically managing JDBC pools, refer to the topic “Dynamically Manage JDBC Pools” on page 23.

Configuring Parameters for a Database Pool

The following table describes the primary parameters of a database pool and the recommended value for each:

Parameter	Description and Recommended Value
initsize	<p>When a JVM is created and connection pool objects are instantiated for each of the pools, a call is made to the JDBC driver for each pool and an initial number of connections are created for each. Connections created from the initsize are part of the connections that will exist in the pool for reuse, as needed.</p> <p>Since various JVMs may load the initial number of connections, but may not really need them, it is recommended that you do not set the initsize to a large value.</p> <p>Default value: 0 or 1</p> <p>Recommended value: 1</p>
maxsize	Maximum size of the connection pool pertaining to a database pool. Once a connection is created up to this value, it is not destroyed. It remains idle until needed.
Buffersize	<p>Maximum number of connections that can be created. Connections created beyond the maxsize value are created and destroyed as needed.</p> <p>Once all the connections are used (maxsize + buffersize), and a new request for a connection comes in, the new request is denied and an error is written stating that a connection cannot be created and that the connection pool has been exhausted.</p>
maxConn	This parameter is not used in Sterling Integrator.

Frequently Asked Questions

- **Q:** How many connections will I need?

A: You can start with the recommended settings provided by the tuning wizard and then monitor noapp.log to ensure that you are not close to the maximum size of the connection pool (maxsize). You should monitor the usage over a period of time and observe if the usage is increasing. If the usage limits exceeds 25% of the settings recommended by the tuning wizard and the demand for connections is increasing, contact Sterling Commerce Professional Services onsite consultant or Sterling Commerce Customer Support.

- **Q:** Why are connections shown as being idle?
A: Getting a new connection from the database is expensive. Connections are held in a pool in order to be available when the system needs them. This means that connections will often be shown as being idle. This is a performance trade-off that enables connections to be available when needed.
- **Q:** Can I kill idle connections?
A: Configure the pool to keep fewer connections, and release connections after a specified amount of time.

Schema Objects

This topic covers the following concepts:

- Placement of Schema Table Objects
- Placement of Indexes
- Sterling Integrator Database Tables

Placement of Schema Table Objects

The *Sterling Integrator Installation Guide* provides directions about how to create the database, including the necessary changes to some critical parameters. The DDL statements allow you to create a simple schema that is suitable for general use. You will need to review, and possibly modify, these statements prior to production.

Placement of Indexes

The DDL statements create a minimal set of indexes for general use. You may need to create additional indexes to suit your business needs. You should contact Sterling Commerce Professional Services or Sterling Commerce Customer Support to create additional indexes. You should also maintain a list of the indexes added and ensure that they are present after you perform an upgrade and monitor its usage. These indexes may or may not be added to the Sterling Integrator and is specific to your operating environment.

You may want to create additional tablespaces for storing indexes in addition to those for the base tables. Tablespaces should have multiple data files for I/O load balancing.

Important: Indexes can also be separated out into different storage areas. This should be done in coordination with your database administrator. However, if you are moving tables to different storage areas, do so only in coordination with Sterling Commerce Professional Services to prevent upgrade issue in the future.

Sterling Integrator Database Tables

The following table lists some of the tables defined in Sterling Integrator experience substantial input, output, and rapid growth.

Table Name	Table Function
DATA_TABLE and TRANS_TABLE	These tables hold the binary payload of documents used for processing within Sterling Integrator.
DOCUMENT, CORRELATION_SET	Contains document metadata used for searchability and enablement of various document processing tasks.
WORKFLOW_CONTEXT	Contains step status and business process flow information.

Table Name	Table Function
ARCHIVE_INFO	Holds lifespan information pertaining to all business processes and document storage in the system.

These tables are candidates for moving to a tablespace or segment that is separate from the rest of the database.

Important: Move these tables only in coordination with Sterling Commerce Professional services. When patches or upgrades that need to recreate a table are installed, the installation will attempt to create them in the default tablespace.

Dynamically Manage JDBC Pools

Sterling Integrator now has the ability to dynamically add, modify, remove JDBC Pools and manage JDBC pools with effective dates for passwords.

Adding JDBC Pools

About this task

Adding a New Database Pool to `jdbc.properties`

Complete these steps to add a new database pool to `jdbc.properties`:

Procedure

1. From the Operations menu, select JDBC Monitor. The JDBC Monitor page appears. If you want to verify that the database you want to add does not already exist, click the link next to View JDBC Report.
2. In the `customer_overrides.properties` file, create a new database connection pool. For additional information, see *Adding New Database Pools in the Lightweight Java Database Connectivity (JDBC) Adapter*.
3. After adding the pool properties in `customer_overrides.properties`, go to JDBC monitor page and click the **Refresh JDBC Pools** button, or run the REFRESHJDBC OPS command from the install root directory:

```
./bin/opscmd.sh -cREFRESHJDBC -nnode1
```

Modifying a Database Pool in `jdbc.properties`

About this task

Complete these steps to update a Database Pool in `jdbc.properties`:

Procedure

1. In the `install_dir/install/properties` directory, locate the `customer_overrides.properties` file.
2. Open the `customer_overrides.properties` file in a text editor.
3. Modify the properties you want to change in the customer pools list of properties.

Note: You can modify any properties for user added pools. For system pools, you cannot change the database type (for example, Oracle to MSSQL), but you can change the database type for customer pools.

4. Save the `customer_overrides.properties` file.

5. After modifying the pool properties in `customer_overrides.properties`, go to JDBC monitor page and click the **Refresh JDBC Pools** button, or run the REFRESHJDBC OPS command from the install root directory: `./bin/opscmd.sh -cREFRESHJDBC -nnode1`

Removing Pool from `jdbc.properties` About this task

Complete these steps to remove a customer pool from `jdbc.properties`:

Procedure

1. In the `install_dir/install/properties` directory, locate the `customer_overrides.properties` file.
2. In the `customer_overrides.properties` file, delete the pool you want to remove.

Note: Verify that all the pool properties are removed for the pool you want to delete, including, `jdbc.properties_*_ext`, `jdbc_customer.properties` and `customer_overrides.properties` files.

3. Save the `customer_overrides.properties` file.
4. After removing the pool properties in `customer_overrides.properties`, go to JDBC monitor page and click the **Refresh JDBC Pools** button, or run the REFRESHJDBC OPS command from the install root directory: `./bin/opscmd.sh -cREFRESHJDBC -nnode1`

Controlling User and Password Credentials with Effective Dates About this task

You can now change database passwords on a scheduled basis in Sterling Integrator. When you add or modify a pool, you now can control user and password credentials with effective dates. Multiple user and password credentials are associated with a pool. A date/time entry indicates to Sterling Integrator when to start using that credential for new connections. This applies primarily to external pools, although Sterling Integrator database pools will also work.

You can use the following variables for the date format:

- 15:00:00 3/16/09
- 3/16/09 15:00:00
- 3/16/2009 15:00:00
- Sat, 12 Aug 1995 13:30:00 GMT
- Sat, 12 Aug 1995 13:30:00 GMT+0430

Note: Other formats may be used as long as they follow the Internet Engineering Task Force (IETF) standard date syntax. For additional information, see <http://www.ietf.org/rfc/rfc3339.txt>.

Pool Property	Description
---------------	-------------

newDBPool.password.1=<new password>	<p>You can specify alphabets and combination of alphabets and numbers for the password. You can use numbers for newDBPool.password.1 or newDBPool.password.2 as well as following examples:</p> <ul style="list-style-type: none"> • newDBPool.password.a=password_a • newDBPool.effective.a=1/01/2005 09:35:00 • newDBPool.password.b=password_b • newDBPool.effective.b=02/01/2009 09:35:00 • newDBPool.password.c=password_c • newDBPool.effective.c=06/18/2009 11:07:00
newDBPool.effective.1=<The date for the new password starts to take affect>	<p>You can specify alphabets and combination of alphabets and numbers for the password. You can use numbers for newDBPool.password.1 or newDBPool.password.2 as well as following examples:</p> <ul style="list-style-type: none"> • newDBPool.password.a=password_a • newDBPool.effective.a=1/01/2005 09:35:00 • newDBPool.password.b=password_b • newDBPool.effective.b=02/01/2009 09:35:00 • newDBPool.password.c=password_c • newDBPool.effective.c=06/18/2009 11:07:00

System Logs and Error Logs

About this task

System Logs

When applicable, the following items are logged in system logs:

- Logging the switch from one credential to the next, as well as the initialization of the pool dates and user IDs being used (not the passwords).
- Logging if the connection is expired when it returns to the pool.
- Logging if two passwords have the same effective dates. In this case, the system randomly selects a password and log that two passwords had the same effective dates. Additional logs on passwords and effective dates may be added.
- Logging when pool properties are changed. If you changed the pool related property like maxSize, or lifespan the following message appears in the system log: "for pool name ***** <PROPERTY> is changed".

Error Logs

The following list provides descriptions of the different types of errors that can be logged:

- Failed to add the pool <pool name>
- Failed to delete the pool <pool name>
- Failed to modify the pool <pool name>

- Failed to create the connections from the pool <pool name>

Oracle Database Configuration and Monitoring

This topic provides information about configuring and monitoring Oracle database. It also provides information about concepts such as init parameter configuration, rolling back or undoing changes to the database, database storage, and monitoring with Automatic Workload Repository (AWR) reports.

init Parameter Configuration Checklist

The following table lists the settings that have been found useful or critical to Sterling Integrator performance on an Oracle database.

Parameter	Mandatory or Recommended	Description and Recommend Value
processes	Recommended	Must be greater than the number of connections required by Sterling Integrator (sum of transactional or local and NoTrans pools in jdbc.properties), and operational management tools.
sga_max_size sga_target pga_aggregate_target	Recommended	1 GB to <i>n</i> GB, depending on the amount of physical memory on your database server. If the server is running only this database, up to 80% of physical memory. To size SGA (Shared Global Area) pools automatically, refer to the topic <i>SGA Pool Sizing</i> .
cursor_sharing	Mandatory	Recommended value: Exact
timed_statistics	Recommended	Recommended value: True
optimizer_mode	Recommended	Recommended value: All_rows
open_cursors	Recommended	Recommended value: 2000 or higher if prepared statement caching is to be used.
db_cache_size	Recommended	Recommended value: Greater than or equal to 19200.
Character set	Mandatory	Recommended value: AL32UTF8

init Parameter Descriptions

Parameter	Description
Number of processes	<p>Maximum number of processes that the Oracle server can create. Each inbound connection from a client requires an available process on the Oracle server, as well as internal processes that execute in the Oracle server itself.</p> <p>This setting needs to be set high enough to accommodate the expected peak connections from Sterling Integrator as well as additional processes for internal server processes plus the possible usage of buffer connections from the Sterling Integrator pools. Note that the expected peak connections from Sterling Integrator are per node, and this number should be multiplied by the number of nodes in a cluster.</p> <p>Important: Exceeding the allocated processes can destabilize the entire Oracle server and cause unpredictable issues.</p> <p>Normally, the Sterling Integrator logs a report if this has occurred, by logging an "ORA-00020 maximum number of processes exceeded" error.</p>

Parameter	Description
open_cursors	<p>Number of cursors that the Oracle server can create. Each process that is servicing a connection will normally spawn two or more cursors, plus additional cursors that are needed for internal server processes.</p> <p>As with the processes, this needs to be set high enough to accommodate the expected peak connections, multiplied by two, plus an allocation for possible buffer connections. This should also be multiplied by the number of nodes in the cluster.</p> <p>A simple rule of thumb is to set this to four times the number of processes, four times the number of nodes of Sterling Integrator, plus an additional 10%.</p> <p>If the number of cursors is exceeded, Sterling Integrator logs "ORA-01000: maximum open cursors exceeded."</p>
Character Set	<p>Controls the storage of character data in the Oracle database. The UTF8 set most closely matches the default character set for Java and will prevent any conversion of character data.</p> <p>AL32UTF8 is the preferred setting for Oracle database.</p>
cursor_sharing	<p>Controls whether SQL sent from a client is reparsed each time, or if the server can reuse ("share") the compiled plan.</p> <p>Sterling Integrator requires the setting to be "exact" for both performance and data integrity.</p>
sga_max_size	<p>Maximum size of the memory allocated to the System Global Area (SGA). This controls the maximum memory that can be allocated to the system area of the Oracle instance. The size of the SGA should never exceed the size of the actual memory (RAM) installed on the server.</p> <p>A rule of thumb is that the SGA maximum size should be allotted as much memory as possible, but should not exceed 80% of the total memory of the server.</p>
sga_max_target	<p>Target value that Oracle uses for automatic allocation of SGA pool resources.</p> <p>Recommended setting is equal to the sga_max_size. The exception is for extremely large databases, where this may need to be adjusted manually.</p>
pga_aggregate_target	<p>Specifies the target aggregate Program Global Area memory available to all server processes attached to the Oracle instance. These are normally internal Oracle processes, and are not used by clients connecting to the instance.</p> <p>Recommended setting is a non-zero value. This allows the Oracle instance to size the SQL working areas automatically as needed.</p>
timed_statistics	<p>Controls whether database statistics for particular times are logged by Oracle. This information may be useful to monitor system or application performance.</p> <p>Setting timed_statistics to TRUE may incur a slight overall load increase on the Oracle instance. However, the information it creates is valuable for diagnosing issues with performance of the instance.</p>
optimizer_mode	<p>Controls the mode that the optimizer uses to select the proper execution path for fulfilling query results, among other functions.</p> <p>Setting this to an incorrect mode can dramatically affect the overall performance of the server.</p> <p>Recommended setting for Sterling Integrator: all_rows</p>

Parameter	Description
db_cache_size	<p>Controls the size of the memory data buffer cache region. This setting works in conjunction with the <code>sga_target</code> to control memory allocation to the data buffer in the Oracle instance. Having sufficient cache available allows the Oracle instance to load the required data blocks into memory, and increase the speed query results and reduce the number of I/O operations required to do so.</p> <p>Recommended setting: 80% of total SGA, with minimum of 19200</p>

SGA Pool Sizing

As an alternative to sizing the main SGA pools in Oracle database manually, it is recommended that you use ASMM (Automatic Shared Memory Management). This feature controls pool size dynamically at run time, readjusting the sizes of the main pools (`db_cache_size`, `shared_pool_size`, `large_pool_size`, `java_pool_size`) based on existing workloads.

Set the following values to enable Oracle ASMM:

- Use an spfile for the `init.ora` values
- Set `sga_target` to a non-zero value
- Set `statistics_level` to `TYPICAL` (the default) or `ALL`
- Set `shared_pool_size` to a non-zero value

Rolling Back or Undoing Changes in Oracle Database

Oracle database supports `AUTO UNDO` management. Its use is recommended to avoid manual monitoring of UNDO segments. Set the `UNDO_MANAGEMENT=AUTO` parameter in `init<SID>.ora`. Your database administrator needs to determine the `UNDO_RETENTION` setting.

Ensure that the file system where `UNDOTBS1` tablespace is located has enough space to use the `AUTOGROW` setting. As a starting point, size the undo tablespace at 20% of the total database size.

Redo Logs

Redo logs are critical for database and instance recovery. Correct redo log configuration is also critical for performance. We have seen best performance while implementing redo logs on ASM. The following is recommended:

- Consider implementing redo logs on dedicated disk devices, preferably RAID 10.
- Consider implementing redo log group log files on alternating disks.

Redo File Size

Your choice of redo file size depends on your trade-off between performance and availability, specifically the time required to recover the Oracle instance in the event of a failure. For performance, some installations opt to create large redo logs in order to reduce the frequency of log switches. However, this means potentially more transactions in the redo logs that must be replayed during recovery.

The general rule for sizing redo log files is to consider the time it takes to switch log files. Generally, do not switch logs more than once every 20 to 30 minutes.

Issue the following query to see how often the redo log files are changing:

```
select * from v$loghist order by first_time desc
```

Following is an example of the output:

```
THREAD# SEQUENCE# FIRST_CHANGE# FIRST_TIME SWITCH_CHANGE#
1 97 7132082 10/20/2008 11:47:53 PM 7155874
1 96 7086715 10/20/2008 11:32:04 PM 7132082
1 95 7043684 10/20/2008 11:15:07 PM 7086715
1 94 6998984 10/20/2008 11:00:57 PM 7043684
1 93 6950799 10/20/2008 10:48:03 PM 6998984
```

In this example, the logs are switched every 15 minutes.

Database Storage

This topic provides information about database storage.

ASM

It is recommended that you use Oracle Automatic Storage Management (ASM) to manage database storage. The benefits of using ASM include:

- Improved I/O performance and scalability
- Simplified database administration
- Automatic I/O tuning
- Reduction in number of objects to manage

For more information, consult these sources:

- ASM documents available on the Oracle Web site.
- Storage vendor best practice papers about how to configure ASM for their storage products.

Raw Disks

For performance-critical systems, importing only raw devices into ASM is recommended. For high-volume processing environments, ensure that ASM is configured with disk devices from high-performance disk storage arrays. Some characteristics that you should look for include large NVRAM caches in order to buffer the disk reads and writes, and for efficient RAID implementation.

Configure ASM with External Redundancy to ensure that redundancy is provided by your storage array.

Raw devices (raw partitions or raw volumes) can be used directly under Oracle on UNIX-based or Linux-based systems and can help performance. However, raw disks have the following potential disadvantages:

- Raw devices may not solve problems with file size writing limits.
To display current UNIX file size limits, run the following command:

```
ulimit -a
```

- It may not be possible to move files to other disk drives if you are using raw devices. If a particular disk drive has intense I/O activity, and performance will benefit from the movement of an Oracle data file to another drive, it is likely that no acceptably-sized partition or volume exists on a drive with less I/O activity.
- Raw devices are more difficult to administer than data files stored on a file system or in an Automatic Storage Management disk group.

When deciding whether to use raw devices, consider the following issues:

- Raw disk partition availability

Use raw partitions for Oracle files only if you have at least as many raw disk partitions as Oracle data files. If disk space is a consideration, and the raw disk partitions are already created, match the data file size to the partition size as closely as possible to avoid wasting space.

You must also consider the performance implications of using all the disk space on a few disks as opposed to using less space on more disks.
- Logical volume manager

Logical volume managers manage disk space at a logical level and hide some of the complexity of raw devices. With logical volumes, you can create logical disks based on raw partition availability.

The logical volume manager controls fixed-disk resources by:

 - Mapping data between logical and physical storage
 - Enabling data to span multiple disks and to be discontinuous, replicated, and dynamically expanded

For RAC, you can use logical volumes for drives associated with a single system, as well as those that can be shared with more than one system of a cluster. Shared drives enables all the files associated with a RAC database to be placed on these shared logical volumes.
- Dynamic performance tuning

To optimize disk performance, you can move files from disk drives with high activity to disk drives with low activity. Most hardware vendors who provide the logical disk facility also provide a graphical user interface (GUI) that you can use for tuning.

Refer to your operating system documentation on how to use raw devices with your OS.

Tablespaces (Single or Multiple)

Prior to production, you should plan the overall storage strategy. The DDLs to create temporary tablespaces and data tablespaces are left to the discretion of the customer. General recommendations include:

- Implement these tablespaces as locally managed tablespaces (LMTs) by specifying extent management as local when creating the tablespace.
- Implement tablespaces with Automatic Space Management by specifying segment space management as auto.
- With LMTs, you may want to consider creating tablespaces that store small reference tables with the autoallocate extent allocation model.
- Consider putting large tables into their own tablespace and using the uniform extent allocation model. Use the default extent size 1 MB.

- Create your temporary tablespace as a temporary data file (temp files). Temp files are used to store intermediate results, such as from large sort operation. Changes to temp files are not recorded in the redo logs.

Monitoring with the AWR Report

Important: The information in an Automatic Workload Repository (AWR) report needs to be considered in relation to Sterling Integrator performance. This information is not meaningful in isolation. You should monitor AWR over a period of time to establish your steady state performance baseline. You should monitor AWR even if Sterling Integrator is performing to your satisfaction; doing so could uncover issues before they become noticeable.

The AWR is used to collect performance statistics, including:

- Wait events used to identify performance problems
- Time model statistics indicating the amount of DB time associated with a process from the V\$SESS_TIME_MODEL and V\$SYS_TIME_MODEL views
- Active Session History (ASH) statistics from the V\$ACTIVE_SESSION_HISTORY view
- Selected system and session statistics from the V\$SYSSTAT and V\$SESSTAT views
- Object usage statistics
- Resource-intensive SQL statements

The repository is a source of information for several other Oracle database features, including:

- Automatic Database Diagnostic Monitor
- SQL Tuning Advisor
- Undo Advisor
- Segment Advisor

By default, AWR snapshots of the relevant data are taken every hour and retained for seven days. It is recommended to set shorter intervals, for example, 30 minutes, because at longer intervals, the issues are less noticeable. Default values for these settings can be altered with statements such as:

```
BEGIN

DBMS_WORKLOAD_REPOSITORY.modify_snapshot_settings(

retention => 43200, -- Minutes (= 30 Days).

Current value retained if NULL.

interval => 30); -- Minutes. Current value retained if NULL.

END;

/
```

A typical Sterling Integrator deployment contains characteristics of both interactive and batch data processing systems, including a high volume of small transactions that rapidly update the database and periodic long-running transactions that touch many records.

The key areas to inspect in an AWR report are:

- Top SQL (by CPU or I/O)
- Top wait events
- Buffer cache hit ratio (refer to the topic "Instance Efficiency Percentages")
- I/O contention (hot spots) on disks (refer to the topic " Tablespace IO Stats and File IO Stats")

Instance Efficiency Percentages

When viewing the Instance Efficiency Percentages area, focus on Buffer Hit % and Buffer Nowait % as shown in the following figure. These are the key performance indicators on the Oracle database server. The goal is to keep these values at 95% or greater, which can typically increase efficiency by increasing buffer cache available to the DB server.

Note: When Sterling Integrator performs internal maintenance, these values may be affected.

Instance Efficiency Percentages (Target 100%)

Buffer Nowait %:	100.00	Redo NoWait %:	100.00
Buffer Hit %:	99.80	In-memory Sort %:	100.00
Library Hit %:	99.96	Soft Parse %:	99.79
Execute to Parse %:	61.03	Latch Hit %:	99.98
Parse CPU to Parse Elapsed %:	233.33	% Non-Parse CPU:	90.82

The following figure shows the buffer pool statistics.

Buffer Pool Statistics

- Standard block size Pools D: default, K: keep, R: recycle
- Default Pools for other block sizes: 2k, 4k, 8k, 16k, 32k

P	Number of Buffers	Pool Hit%	Buffer Gets	Physical Reads	Physical Writes	Free Buff Wait	Writ Comp Wait	Buffer Busy Waits
D	962,240	100	206,779	428	1,391	0	0	0

Tablespace IO Stats and File IO Stats

Review the IO Stats area to obtain information about hotspots or drive contention. Analysis of information such as Buffer Wait and Average Buffer Waits can help you determine whether you are getting the best performance from your drives, and identify potential bottlenecks. Separation of objects into their own unique tablespaces can contribute to performance tuning of specific areas.

Wait Events

The AWR report for Sterling Integrator shows typical wait events as displayed in the following figure. To see if the wait events are problematic, compare it to a baseline in your environment. The following figures are an example of what you can see normally for Sterling Integrator.

Wait Events

- s - second
- cs - centisecond - 100th of a second
- ms - millisecond - 1000th of a second
- us - microsecond - 1000000th of a second
- ordered by wait time desc, waits desc (idle events last)

Event	Waits	%Time -outs	Total Wait Time (s)	Avg wait (ms)	Waits /txn
db file sequential read	424	0.00	0	1	2.90
os thread startup	9	0.00	0	38	0.06
log file parallel write	353	0.00	0	1	2.42
control file parallel write	204	0.00	0	0	1.40
db file scattered read	13	0.00	0	6	0.09

Event	Waits	%Time -outs	Total Wait Time (s)	wait (ms)	Waits /txn
log file sync	11,748,794	.0	68,378	6	1.0
db file sequential read	16,471,961	.0	26,834	2	1.4
log file parallel write	8,371,728	.0	15,820	2	0.7
enq: TX - index contention	665,441	.0	7,239	11	0.1
log file sequential read	187,827	.0	2,100	11	0.0
FX Deq Credit: send blkl	42,843	.0	740	17	0.0
Log archive I/O	204,864	.0	681	3	0.0

The SQL statements that should be inspected are SQL ordered by Gets and SQL ordered by Reads as shown in the following figure. Again, compare these with the baseline of the system to determine if any SQL statement is behaving oddly. If yes, refer to the query plans to determine if the query plan is appropriate.

SQL Statistics

- SQL ordered by Elapsed Time
- SQL ordered by CPU Time
- SQL ordered by Gets
- SQL ordered by Reads
- SQL ordered by Executions
- SQL ordered by Parallel Calls
- SQL ordered by Shared Memory
- SQL ordered by Version Count
- Complete List of SQL Text

[Back to Top](#)

SQL ordered by Elapsed Time

- Resources reported for PL/SQL code includes the resources used by all SQL statements called by the code
- % Total DB Time is the Elapsed Time of the SQL statement divided into the Total Database Time multiplied by 100

Elapsed Time (s)	CPU Time (s)	Executions	Elap per Exec (s)	% Total DB Time	SQL ID	SQL Module	SQL Text
1	1	1	1.37	152.88	6wssq0f56m1f		insert into wfs_opp_jgpl_ag...
1	0	1	0.78	80.80	qf1vzb28h772h		select rownum, obj, descr...
0	0	381	0.00	15.36	ab75wz2m1m0x	DEV_SystemPool	begin MGMT_JOB_HOME_get_ache...
0	0	10	0.01	12.48	8qezbtou8ka2q		DECLARE job_BINARY_INTEGER = ...
0	0	8	0.01	6.18	1787yfm1m6p1t	DEV_SoundedPool	INSERT INTO MGMT_SEVERITYACTL...
0	0	10	0.00	5.50	2b88484skwv1y	DEV_SystemPool	RETURN END_NOTIFICATION_QUEUE...
0	0	1	0.05	5.32	34y4bbsr3k1h0		insert into wfs_jobh (snoo...

Occasionally, you may observe full table scans on some tables. It is not a cause for concern unless the overall system performance is slow or business level SLAs are not met.

Index and Table Statistics

Database optimizers rely on relatively up-to-date table and index statistics to generate optimal access plans. Oracle does not require statistics to be absolutely correct or current, just relatively correct and representative. You should refresh statistics periodically (several times in a day or at least once every day). You can refresh statistics by using the default job provided in Oracle or by creating the custom jobs. Ensure that you update statistics when there is a significant change in data (20% or more).

Oracle database gathers statistics automatically during its maintenance window (10 p.m. to 2 a.m. on week nights and all day on weekends) for tables that have undergone sufficient changes. Oracle bypasses statistics generation for tables that have not changed significantly.

To manually refresh statistics, use the DBMS_STATS package. This sample invocation refreshes the statistics for the entire database:

```
EXECUTE DBMS_STATS.GATHER_SCHEMA_STATS(<schema owner>, CASCADE => TRUE,  
OPTIONS =>'GATHER AUTO')
```

where **<schema owner>** = Sterling Integrator schema owner.

One way to determine the update frequency for statistics is when the data for one customer changes by approximately 20%. This is governed by the document's lifespan.

Speak with the Sterling Integrator Administrator to find out more about document lifespans. Generally, updating statistics once every 24 hours is sufficient.

Index Rebuilds

Index rebuilds are also recommended for the Sterling Integrator schema when the data in the tables change by approximately 20%. However, this must be validated by looking at the system performance as a whole. Online index rebuilds, which cause minimal impact to the system, is possible.

The Sterling Integrator system administrator should coordinate with the DBA to identify the indexes that should be rebuilt. You should always rebuild indexes at schema level.

Sterling Integrator is a high-volume OLTP application, and rarely, there may be times when you may need to shrink the tables or rebuild indexes. The following Oracle documentation provides details about these:

- Shrinking Database Segments Online –
http://download.oracle.com/docs/cd/E11882_01/server.112/e17120/schema003.htm#ADMIN10161
- Rebuilding an Existing Index –
http://download.oracle.com/docs/cd/E11882_01/server.112/e17120/indexes004.htm#ADMIN11734

Enable BLOB Caching

By default, Oracle database configures itself with the BLOB caching disabled when you choose Oracle as your database during Sterling Integrator installation.

You can significantly improve performance by enabling the cache on the BLOB data object. Run the following commands while connected to the database:

```
ALTER TABLE DATA_TABLE MODIFY LOB DATA_OBJECT CACHE
```


```
ALTER TABLE TRANS_DATA MODIFY LOB DATA_OBJECT CACHE
```

For regular tuning and best practices pertaining to Oracle, refer to the corresponding Oracle documentation.

Oracle Real Application Clusters (RAC)

Oracle Real Application Clusters (RAC) ensures high availability. Oracle RAC supports deployment of a database across multiple machines, thereby reducing the downtime due to hardware failures or outages, planned or unexpected.

The following diagram shows how all transactions should be sent to one Oracle RAC node at a time. If that database node goes down, all Sterling B2B Integrator nodes reconnect to the second available database node. This scenario requires that load balancing is turned off, and provides better overall system performance.



The following table lists the features and benefits of Oracle RAC when used with Sterling B2B Integrator:

Features	Benefits
Load balancing	Oracle RAC supports initial connection time load balancing and runtime connection load balancing. However, running Sterling B2B Integrator with load balancing turned on is likely to result in performance issues. On the client side (jdbc url), you should set the parameter LOAD_BALANCE to OFF. On the server side, turn off load balancing on the database service.

Features	Benefits
High availability (Failover)	<p>Oracle RAC supports high availability by eliminating the dependence on a single server, where the chances of database downtime are more due to hardware failure.</p> <ul style="list-style-type: none"> Initial connection time failover <p>In initial connection time failover, the database node availability is determined while establishing the connection. The connection is established with the available database node. However, if a database node goes down while a transaction is in progress, the transaction will fail and throw a SQLException error, and a new connection will be established with the available node.</p> <p>Initial connection time failover can be configured on the JDBC multi-instance URL on the client side by setting this parameter:</p> <p>FAILOVER=ON</p> <p>Sterling B2B Integrator supports two-node Oracle RAC setup.</p> <p>To establish initial time connection failover in Sterling B2B Integrator, perform the following tasks:</p> <ol style="list-style-type: none"> Configure the following properties in ASI node and in customer_overrides.properties.in file in containers. You can set values for these properties (in milliseconds) that are suitable for your operating environment. <p>Setting a value too low would result in genuine SQL queries getting terminated before they have a chance to be processed and return the dataset. A value too high means that the application threads are holding connections which are waiting on a query timeout for that duration and are not available for other tasks. This could result in lower throughput. An optimum setting depends on the customer's environment.</p> <pre> jdbcService.oraclePool.prop_oracle.jdbc. ReadTimeout=<milliseconds> jdbcService.oraclePool_local.prop_oracle.jdbc. ReadTimeout=<milliseconds> jdbcService.oraclePool_NoTrans.prop_oracle.jdbc. ReadTimeout=<milliseconds> jdbcService.oracleUIPool.prop_oracle.jdbc. ReadTimeout=<milliseconds> jdbcService.oracleArchivePool.prop_oracle.jdbc. ReadTimeout=<milliseconds> </pre> From the <i>install_dir/install/bin</i> (<i>install_dir\install\bin</i> for Windows) directory, run (UNIX or Linux) setupfiles.sh or (Windows) setupfiles.cmd. Restart the ASI and container JVMs to apply the changes.

Features	Benefits
High availability (Failover), continued	<ul style="list-style-type: none"> Runtime connection failover When one Oracle RAC node goes down, the active connections from the node that went down will failover to the active node. When the node that went down is up and running again, the connections are not redistributed. You can redistribute the connection pools by setting lifespan parameter to a value greater than 0, which ensures that connection pools are terminated after the lifespan value is met and load balancing occurs when new connections are established. However, configuring a low value lifespan may lead to connections being terminated quickly thereby losing the benefit of connection pool. To override all the Oracle pools lifespan to a value greater than 0, set the following value in the customer_overrides.properties file: jdbcService.<oraclepool>. lifespan=180000
Visibility Event Queues	<p>You can increase the number of visibility queues and event input queue capacity in customer_overrides.properties to ensure that adapters are responsive.</p> <p>To increase the default queue size and the thread counts, perform the following tasks:</p> <ol style="list-style-type: none"> Set the following values in the ASI node and in customer_overrides.properties.in file in containers: dmvisibility.number_visibility_queues=6 dmvisibility.event_input_queue_capacity=3072 From the <i>install_dir/install/bin</i> (<i>install_dir\install\bin</i> for Windows) directory, run (UNIX or Linux) setupfiles.sh or (Windows) setupfiles.cmd. Restart the ASI and container JVMs to apply the changes.
Fast Connection Failover (FCF)	Sterling B2B Integrator does not support Fast Connection Failover (FCF).
Transparent Application Failover (TAF)	Sterling B2B Integrator does not support Transparent Application Failover (TAF).

IBM DB2 for LUW Configuration and Monitoring

This topic provides information about configuring and monitoring IBM DB2 for LUW (Linux, UNIX, and Windows). It describes the recommended registry variables and DBM CFG parameters. It also provides details about the storage subsystem, log configuration, database monitoring, and index and table statistics.

Recommended Settings for IBM DB2 Registry Variables

The following table describes the IBM DB2 registry values that are either recommended or are critical for IBM DB2 performance with Sterling Integrator:

Variable	Mandatory or Recommended	Recommended Value
DB2_USE_ALTERNATE_PAGE_CLEANING	Recommended	Recommended value: ON

Variable	Mandatory or Recommended	Recommended Value
DB2_EVALUNCOMMITTED	Recommended	<p>Enabling this variable can reduce the amount of unnecessary lock contention from Read Share and Next Key Share. By default, DB2 requests share locks on the index or record before verifying if the record satisfies the query predicate. Queries that scan a set of records in tables with high frequency of inserts or updates can unnecessarily block records that do not belong to its result set.</p> <p>When you set DB2_EVALUNCOMMITTED to ON, DB2 performs an uncommitted read on the record to perform the predicate verification. If the record satisfies the predicate, DB2 requests a share lock on that record.</p> <p>Recommended value: ON</p>
DB2_SKIPDELETED	Mandatory	<p>Allows index-range queries or table-scan queries to skip records that are in an uncommitted delete state. This reduces the amount of lock contention from Read Share and Next Key Share locks from range queries in tables with a high frequency of deletes.</p> <p>When enabled, DB2_SKIPDELETED allows, where possible, table or index access scans to defer or avoid row locking until a data record is known to satisfy predicate evaluation. This allows predicate evaluation to occur on uncommitted data.</p> <p>This variable is applicable only to statements using either Cursor Stability or Read Stability isolation levels. For index scans, the index must be a type-2 index. Deleted rows are skipped unconditionally on table scan access while deleted keys are not skipped for type-2 index scans unless DB2_SKIPDELETED is also set.</p> <p>Recommended value: ON</p>
DB2_SKIPINSERTED	Mandatory	<p>Allows SELECTs with Cursor Stability or Read Stability isolation levels to skip uncommitted inserted rows. This reduces record lock contention on tables with heavy insert rates.</p> <p>Recommended value: ON</p>

Variable	Mandatory or Recommended	Recommended Value
DB2_PARALLEL_IO	Recommended	<p>Changes the way in which DB2 calculates I/O parallelism to the tablespace. By default, DB2 sets I/O parallelism to a tablespace equal to the number of containers in that tablespace. For example, if the tablespace has four containers, prefetches to that tablespace are performed as four extent-sized prefetch requests.</p> <p>Set the DB2_PARALLEL_IO variable if you have implemented containers on striped devices (for example, RAID-5, RAID-10, or RAID-01). If you set DB2_PARALLEL_IO=ON or DB2_PARALLEL_IO=*, DB2 assumes that containers are implemented on a RAID 5 (6+1) configuration: six data disks plus one parity disk.</p> <p>In this example, prefetches to the four-container tablespace mentioned above are performed in 24 extent-sized prefetch requests.</p> <p>To assess the effectiveness of your prefetch parallel I/O settings, monitor the unread_prefetch_pages and prefetch_wait_time monitor element with the snapshot_database monitor. The unread_prefetch_pages monitor element tracks the number of prefetch pages that were evicted from the buffer pool before it was used. A continually growing number indicates that the prefetch requests are too large, either because the prefetch size is larger than the pages needed or the prefetch activities are bringing in too many pages for the capacity of the buffer pool. In either case, you may want to consider reducing the prefetch size.</p> <p>If you have high prefetch_wait_time values, the application might be waiting for pages.</p>
DB2_NUM_CKPW_DAEMONS	Recommended	<p>Set this on IBM AIX 5.3 only if you observe a memory leak during connect authentication within DB2's authentication daemons (db2ckpwd processes) as a result of calling the AIX loginsuccess() API.</p> <p>Symptoms may include excessive memory usage, an instance crash due to a trap in the db2ckpwd process, or general authentication failures. Verify by monitoring the SZ and RSS values for db2ckpwd processes. Use the <code>ps aux grep db2ckpwd</code> command and look at the fifth and sixth columns of output.</p> <p>DB2 workaround exists for this problem. Set the following registry variable, and recycle this instance:</p> <pre>db2set DB2_NUM_CKPW_DAEMONS=0</pre> <p>Also see IBM APAR IY78341.</p> <p>Recommended value: 0</p>
DB2LOCK_TO_RB	Recommended	Recommended value: STATEMENT

Recommended Settings for DBM CFG Parameters

Let DB2 automatically manage the following parameters for DB2 9.x by accepting the default values:

- INSTANCE_MEMORY
- FCM_NUM_BUFFERS
- FCM_NUM_CHANNELS

Enabling the following monitor switches is recommended in production:

Parameter	Mandatory or Recommended	Recommended Value
DFT_MON_BUFPOOL	Recommended	ON
DFT_MON_LOCK	Recommended	ON
DFT_MON_SORT	Recommended	ON
DFT_MON_STMT	Recommended	ON
DFT_MON_TABLE	Recommended	ON
DFT_MON_TIMESTAMP	Recommended	ON
DFT_MON_UOW	Recommended	ON
MON_HEAP_SZ	Recommended	AUTOMATIC
MAXAGENTS	Recommended	<p>Limits the number of database manager agents (both coordinator or subagents) that can run concurrently. Pick a high enough number so that the combined connection requirements from the application servers, agents, monitoring tools, and so on do not exceed the MAXAGENTS limit during peak processing periods. If you do, you must restart the DB2 instance to increase this limit.</p> <p>Recommend value: Must be greater than the number of connections needed by Sterling Integrator (that is, the sum of transactional (local) plus NoTrans pools in jdbc.properties file) plus the connections needed by operational management tools.</p> <p>With the use of connection pooling in the NoApp Server, the number of database connections is less than the number of users who are logged in. Depending on your anticipated peak workload traffic, this parameter may range from a small number such as 25 to a large number in the thousands.</p> <p>Monitor the number of concurrent connections in production (especially during peak periods) to ensure that it does not reach the maximum. When the MAXAGENTS limit is reached, DB2 refuses to establish new connection requests.</p>

Estimation Guidelines for Number of Connections

You can roughly estimate the number of concurrent user connections required for Sterling Integrator with this formula:

$$\text{Concurrent (DB2) connections} = A + B + C + D$$

Variable	Description
A	Maximum number of UI-based connections
B	Maximum NoApp Server connection pool max size times the number of NoApp Server instances (as in a cluster)
C	Any additional connections that are opened by customized code or user exits that do not go through the NoApp Server connection pools. This connection requirement is specific to your implementation.
D	Number of connections required by the containers.

Benchmarking

Benchmarking your system to validate assumptions and estimates prior to a production implementation is strongly recommended. During the test, monitor the connection pool usage levels in each of the application server instances, the number of agents to run to meet your processing and service levels, and the actual DB2 database connections established.

Recommended Settings for DB CFG Parameters

For optimal performance, the parameters and values described in the following table are recommended for DB2 9.x.

Parameter	Mandatory or Recommended	Description and Recommended Value
SELF_TUNING_MEM	Recommended	Enables the DB2 self-tuning memory manager (STMM) to automatically and dynamically set memory allocations to memory consumers such as buffer pools, lock lists, package cache, and sort heap. Recommended value: ON
DATABASE_MEMORY	Recommended	Allows DB2 to adjust the amount of database memory depending on load, memory pressures, and other factors. Recommended values: AUTOMATIC (for Windows and AIX) COMPUTED (for Linux, HP-UX, and Solaris)
LOCKLIST	Recommended	Allows STMM to dynamically manage memory allocations. Recommended value: AUTOMATIC
MAXLOCKS	Recommended	Allows STMM to dynamically manage memory allocations. Recommended values: AUTOMATIC If AUTOMATIC is not supported: 100
PCKCACHESZ	Recommended	Allows STMM to dynamically manage memory allocations. Recommended value: AUTOMATIC
SHEAPTHRES_SHR	Recommended	
SORTHEAP	Recommended	
NUM_IOCLEANERS	Recommended	

Parameter	Mandatory or Recommended	Description and Recommended Value
NUM_IOSERVERS	Recommended	Recommended value: AUTOMATIC
DFT_PREFETCH_SZ	Recommended	Recommended value: AUTOMATIC
MAXAPPLS	Recommended	Recommended value: AUTOMATIC
APPLHEAPSZ	Recommended	Recommended value: AUTOMATIC
APPL_MEMORY	Recommended	Recommended value: AUTOMATIC
APP_CTL_HEAP_SZ	Recommended	512 Note: APP_CTL_HEAP_SZ is deprecated in DB2 9.5.
DBHEAP	Recommended	Amount required depends on the amount of memory available and the traffic volume. Recommended value: AUTOMATIC Note: The default DBHEAP value is inadequate.
LOGFILSIZ	Recommended	Refer to the topic Log Configuration. Recommended value: 65536 if configuring 20 transaction logs of 256 MB (65536 4 K pages)
LOGPRIMARY	Recommended	Number of primary transaction logs. Refer to the topic Log Configuration. Recommended value: 40 or more
LOGSECOND	Recommended	Number of secondary transaction logs. These are allocated by DB2 when it cannot reuse any of the primary logs due to active transactions. Refer to the topic Log Configuration. Recommended value: 12
NUM_LOG_SPAN	Recommended	Refer to the topic Log Configuration. Recommended value: LOGPRIMARY - Safety buffer
DFT_DEGREE	Recommended	Sets the default degree of parallelism for intrapartition parallelism. In general, online transactional applications such as Sterling Integrator typically experience a high volume of short queries that do not benefit from parallel queries. Recommended value: 1 – Disable intrapartition parallelism
DB2LOCK_TO_RB	Recommended	Recommended value: STATEMENT
Database Code Set	Mandatory	Recommended value: UTF-8

Recommended Settings for DB2 9.7

The following table lists the recommended settings specific to DB2 9.7:

Feature	Description	Value
Currently Committed Semantics	<p>The DB2 database used the default Cursor Stability (CS) isolation level in all versions until DB2 9.7. If an application changed a row and another application tried to read that row before the first application committed the changes, the second application waited until the commit.</p> <p>You can now set the currently committed semantics of the CS level, which informs DB2 that when the second application tries to read a locked row, it will get an image of what the row looked like before committing the change.</p> <p>In the enhanced currently committed semantics, only committed data is returned, as it used to be earlier. However, now read operation does not wait for the write operation to release the row locks. Instead, the read operation returns the data prior to the start of the write operation.</p> <p>The currently committed semantics is turned on by default in the new DB 9.7 database. The new database configuration parameter, cur_commit is used to override this behavior.</p> <p>Currently committed semantics requires more log space for write operations. In addition, extra space is required for logging the first update of a data row during a transaction, which can have an insignificant or measurable impact on the total log space used.</p>	Recommended value: ON

Feature	Description	Value
	<p>Limitations</p> <p>The following list provides limitations that apply to the currently committed semantics:</p> <ul style="list-style-type: none"> • The target table in sections that is used for data update or delete operation does not use currently committed semantics. • An uncommitted modification to a row forces the currently committed read operation to access appropriate log records and determine the currently committed version of the row. Although log records that are no longer present in the log buffer can be physically read, currently committed semantics does not support the retrieval of log files from the log archive. • The following scans do not use currently committed semantics: <ul style="list-style-type: none"> – Catalog table scans – Referential integrity constraint enforcement scans – LONG VARCHAR or LONG VARCHAR column reference scans – Range-clustered table (RCT) scans – Spatial or extended index scans 	
Statement concentrator	<p>DB2 uses less server resources while processing queries that have parameters instead of queries that have literal values in them. DB2 will compile an SQL statement once and will cache it. It presents the same query execution plan from the cache the next time for the same query thereby utilizing fewer resources to compile/prepare the same statement.</p> <p>However, it becomes a difficult task when SQL statements use literal values instead of parameters matching incoming statements to what is already present in the statement. To prevail over this situation, statement concentrator modifies dynamic SQL statements at the database server so that similar, but not identical, SQL statements can share the same access plan.</p>	<p>You can enable statement concentrator by running the following SQL statement:</p> <pre>db2 update db cfg for <db-alias> using stmt_conc literals</pre>

Feature	Description	Value
Inline LOBs	<p>Sterling Integrator extensively uses large objects (LOBs). These LOBs are usually a few kilobytes in size.</p> <p>The LOB data access performance can be improved by placing the LOB data within the rows on data pages instead of the LOB storage object.</p> <p>Typically, LOBs are stored in a separate storage object that is outside the base table. LOBs can be stored in the formatted rows of the base table if they are sufficiently sized. Depending on the page size, the maximum length of a LOB that can qualify for in table in-lining is 32660 bytes. Such LOBs are commonly known as inline LOBs.</p> <p>Earlier, the processing of such LOBs created bottlenecks for application. However, now LOBs improve the performance of queries that access the LOB data as there is no additional I/Os required for fetching, inserting, or updating the data.</p> <p>LBO inlining is enabled for all LOB columns in the new or existing tables for DB2 9.7 and for all existing LOB columns in an upgrade.</p>	Enable LOB inlining through the <code>INLINE LENGTH</code> option on the <code>CREATE TABLE</code> statement or the <code>ALTER TABLE</code> statement.

Storage Subsystem

Prior to production, plan the overall storage strategy. The DDLs to create temporary tablespaces and data tablespaces are left to the discretion of the customer. This topic discusses some general recommendations.

The Sterling Integrator installation creates tables and indexes. The tables use different page sizes: 4K, 8K, and 16K. You should have a tablespace to accommodate tables with these page sizes. DB2 automatically places tables and indexes in the available tablespaces using its internal logic.

SMS Tablespaces in DB2

In an SMS (System Managed Space) table space, the operating system's file system manager allocates and manages the space where the table is stored. The storage model typically consists of many files, representing table objects, stored in the file system space. The user decides on the location of the files, DB2 controls their names, and the file system is responsible for managing them. By controlling the amount of data written to each file, the database manager distributes the data evenly across the table space containers. By default, the initial table spaces created at database creation time are SMS.

SMS table spaces are defined using the `MANAGED BY SYSTEM` option in the `CREATE DATABASE` command, or in the `CREATE TABLESPACE` statement.

Consider two key factors when you design your SMS table spaces:

- Containers for the table space. You must specify the number of containers that you want to use for your table space. It is crucial to identify all the containers you want to use, because you cannot add or delete containers after an SMS table space is created.
- In a partitioned database environment, when a new partition is added to the database partition group for an SMS table space, the ALTER TABLESPACE statement can be used to add containers for the new partition.

Each container used for an SMS table space identifies an absolute or relative directory name. Each of these directories can be located on a different file system (or physical disk). Estimate the maximum size of the table space with:

Number of containers * maximum file system size supported by the operating system

This formula assumes that there is a distinct file system mapped to each container, and that each file system has the maximum amount of space available. In practice, this may not be the case, and the maximum table space size may be much smaller. There are also SQL limits on the size of database objects, which may affect the maximum size of a table space.

- Extent size for the table space
The extent size can only be specified when the table space is created. Because it cannot be changed later, it is important to select an appropriate value for the extent size. If you do not specify the extent size when creating a table space, the database manager will create the table space using the default extent size as defined by the dft_extent_sz database configuration parameter. This configuration parameter is initially set based on the information provided when the database is created. If the dft_extent_sz parameter is not specified in the CREATE DATABASE command, the default extent size is set to 32.

DMS Tablespaces in IBM DB2

In a DMS (Database Managed Space) table space, the database manager controls the storage space. The storage model consists of a limited number of devices or files whose space is managed by DB2. The database administrator decides which devices and files to use, and DB2 manages the space on those devices and files. The table space is essentially an implementation of a special-purpose file system designed to meet the needs of the database manager.

A DMS table space containing user-defined tables and data can be defined as:

- A regular table space to store any table data, and optionally, index data
- A large table space to store long field or LOB data, or index data

When designing your DMS table spaces and containers, you should consider the following:

- The database manager uses striping to ensure an even distribution of data across all containers.
- The maximum size of regular table spaces is 64 GB for 4 KB pages, 128 GB for 8 KB pages, 256 GB for 16 KB pages, and 512 GB for 32 KB pages. The maximum size of large table spaces is 2 TB.

Unlike SMS table spaces, the containers that make up a DMS table space do not have to be the same size. However, the use of unequal container sizes is not usually recommended because it results in uneven striping across the containers, and results in suboptimal performance. If any container is full, DMS table spaces use the available free space from other containers.

- Because space is preallocated, it must be available before the table space can be created. When using device containers, the device must also exist with enough space for the definition of the container. Each device can have only one container defined on it.

To avoid wasted space, the size of the device and the size of the container should be equivalent. If, for example, the device is allocated with 5000 pages, and the device container is defined to allocate 3000 pages, 2000 pages on the device will not be usable.

By default, one extent in every container is reserved for overhead. Only full extents are used. For optimal space management, use the following formula to determine an appropriate size when allocating a container:

$$\text{extent_size} * (n + 1)$$

In this formula:

- extent_size is the size of each extent in the table space
 - n is the number of extents that you want to store in the container
- Device containers must use logical volumes with a "character-special interface," and not physical volumes.

You can use files instead of devices with DMS table spaces. No operational difference exists between a file and a device; however, a file can be less efficient because of the run-time overheads associated with the file system. Files are useful when devices are not directly supported, a device is not available, maximum performance is not required, or you do not want to set up devices.

If your workload involves LOBs or LONG VARCHAR data, you can derive performance benefits from file system caching.

Automatic Storage Management (ASM)

Automatic storage grows the size of your database across disk and file systems. It removes the need to manage storage containers manually by taking advantage of the performance and flexibility of database managed storage. In DB2 9.x, automatic storage is enabled by default.

A database needs to be enabled for automatic storage when it is created. DB2 9.5 and DB2 9.7 enable automatic storage by default when you create new databases. You cannot enable automatic storage for a database if it was not created to accommodate for it. Similarly, you cannot disable automatic storage for a database that was originally created to use it.

With automatic storage, you no longer need to worry about tasks such as creating additional table spaces for capacity, adding containers, and monitoring container growth. When you want to restore a database backup, on a different system (using different directory or path structures), you can redefine the storage paths, such that the new paths are used instead of the ones stored in the backup.

The following examples illustrate automatic storage usage on UNIX and Linux systems.

When a database is created, you can specify the storage pool for use by the database. If no storage paths are specified, the default database path (dftdbpath) is used.

CREATE DATABASE test on /data/path1, /data/path2

You can add additional storage paths to the pool:

```
ALTER DATABASE ADD STORAGE /data/path3, /data/path4
```

Earlier, when you created tablespaces, you had to specify containers for them. You can now specify that they automatically use the database storage pool:

```
CREATE TABLESPACE ts1 MANAGED BY AUTOMATIC STORAGE
```

You can also define policies for storage growth and limits:

```
CREATE TABLESPACE ts2
```

```
INITIAL SIZE 500K
```

```
INCREASE SIZE 100K
```

```
MAXSIZE 100M
```

Note: For performance and simplicity, Automatic Storage Management is recommended. If this is not possible, define SYSCATSPACE and SYSTOOLSPACE, temporary tablespaces as SMS, and other tablespaces holding application data as DMS.

Log Configuration

This topic provides information about the following logs:

- LOGFILSIZ, LOGPRIMARY, LOGSECOND
- NUM_LOG_SPAN

LOGFILSIZ, LOGPRIMARY, LOGSECOND

DB2At a minimum, configure 20 transaction logs (LOGPRIMARY=20) of 256 MB (LOGFILSIZ=65536 4K-pages) for Sterling Integrator.

As an additional precaution, configure at least 12 secondary transaction logs (LOGSECOND=12). DB2 allocates secondary logs when it cannot reuse any of the primary logs due to active transactions.

Track the following monitor elements to assess the effectiveness of these settings, and adjust as needed:

- total_log_used and tot_log_used_top to see how much of the logs are used
- Which workloads are consuming or holding the transaction logs when LOGPRIMARY approaches the total primary log capacity. If needed, raise the setting for LOGPRIMARY.
- sec_log_used_top and sec_logs_allocated to see if secondary transaction logs are used. Investigate how often logging spills over to the secondary logs and what workloads are running during the spill. If needed, increase LOGPRIMARY to prevent log spills.

NUM_LOG_SPAN

This parameter limits the number of logs a transaction can span, which prevents situations where DB2 cannot switch transaction logs because all the transaction logs are active. For example:

- Someone may have updated a record in Control Center, but may have forgotten to commit the change.

- Updates to one or more database records might not get committed due to a software bug.

Set NUM_LOG_SPAN to a minimum of 12 so that valid long-running transactions are not prematurely forced, and a maximum of LOGPRIMARY minus a safety buffer. For example, if LOGPRIMARY=20 and you decide upon a safety buffer of 4, the maximum NUM_LOG_SPAN=16.

DB2 Monitoring

DB2 includes many facilities for tracing system activity at just about any level of detail. This section covers the following topics:

- Snapshot Monitors
- Snapshot Monitor Commands
- DB2 Event Monitors
- Use an Event Monitor to Diagnose and Flush Deadlocks
- DB2 Performance Expert

Snapshot Monitors

Information about snapshot monitors is available by searching at <http://www.ibm.com/developerworks>.

Snapshot monitors collect information about the state of a DB2 instance and any databases it controls at a specific point in time. Snapshots are useful for determining the status of a database system. When taken at regular intervals, they can provide information that lets you observe trends and identify potential problem areas. Snapshot monitoring is performed by executing the **GET SNAPSHOT** command from the DB2 Command-Line Processor (CLP).

Snapshot Monitor Commands

To use snapshot monitors, turn on the monitor switches and view data. The following table describes the command and syntax for specific tasks:

Task	Command and Syntax
Turn on monitor switches	db2 UPDATE MONITOR SWITCHES USING options Options are: lock on sort on statement on table on timestamp on uow on Note: Monitor switches can also be controlled in the database manager configuration.
Check status of monitor switches	GET MONITOR SWITCHES

Task	Command and Syntax
View snapshot data	<p>Switches must be turned on to view data.</p> <p>GET SNAPSHOT FOR options</p> <p>Options:</p> <p>bufferpools on database_name</p> <p>locks on database_name</p> <p>dynamic sql on database_name</p> <p>tables on database_name</p> <p>applications on database_name</p> <p>tablespaces on database_name</p> <p>database on database_name</p>

DB2 Event Monitors

Information about snapshot monitors is available by searching at <http://www.ibm.com/developerworks>.

Event monitors are used to monitor events in an area over time, such as by producing event records at the start and end of multiple events. These records are useful for resource planning, trend analysis, and so on.

The most common uses for event monitors are for connections, locks, and statements. Output can be written to files, named pipes (useful if writing programs to retrieve monitor data), or to tables.

Event monitor output can be either blocked or nonblocked. Blocked output ensures against data loss, but should generally be avoided for systems that produce a large volume of records and for statement event monitors to avoid crashes.

You can limit the granularity of an event monitor to a single user or application.

Connection events can be used to track system usage on a user basis or application basis. This data allows you to identify poorly performing programs, the heaviest users, and usage trends. Daily review facilitates discussions with users about the activity or about adjusting the DB2 physical design.

Use an Event Monitor to Diagnose and Flush Deadlocks

Sterling Integrator, along with the necessary DB2 registry variables, is designed to minimize the occurrence of deadlocks. However, deadlocks may still occur in some situations, including the following:

- Custom BP code may obtain records in a different order.
- DB2 may choose an access plan that retrieves records in a different order.

The following table describes the actions to be performed for specific tasks:

Task	What to Do
Set an event monitor to help diagnose deadlocks	<p>Run this command:</p> <pre>db2 -v create event monitor \$MON for deadlocks with details write to file \$OUTDIR buffersize 64 nonblocked</pre> <pre>db2 -v set event monitor \$MON state = 1</pre> <p>Options:</p> <p>MON – Monitor name (for example, DLMON)</p> <p>OUTDIR – directory to store deadlock information</p>
Determine that a deadlock has occurred	View data in the db2diag.log or in the Sterling Integrator logs
Release a deadlock by flushing buffers	<p>Run the following command:</p> <pre>db2 flush event monitor \$MON</pre> <pre>db2evmon -path \$OUTDIR</pre> <p>The flush ensures that the deadlock records in the buffers are written out. The db2evmon command formats the deadlock information.</p>

Optim Performance Manager Extended Edition

Optim Performance Manager is a performance analysis and tuning tool for managing a mix of DB2 systems. It can be used to identify, diagnose, solve, and prevent performance problems.

For more information about the Optim Performance Manager Extended Edition, refer to:

<http://www-01.ibm.com/software/data/optim/performance-manager-extended-edition>

DB2 Index and Table Statistics

DB2 uses catalog statistics to determine the most efficient access plan for a query. Out-of-date or incomplete statistics for a table or an index may slowdown query execution.

Manual statistics management can be time-consuming. Automatic statistics collection can run the RUNSTATS utility in the background to collect and maintain the statistics you require.

To configure your database for automatic statistics collection, set each of the following parameters to ON:

- AUTO_MAINT
- AUTO_TBL_MAINT
- AUTO_RUNSTATS

Volatile Tables

In some cases, the content of the WORKFLOW_CONTEXT, TRANS_DATA, and other tables can fluctuate significantly during the day. The resulting statistics, which represent a table at a single point in time, can be misleading. In such situations, mark the table as volatile with the following command:

```
alter table <table name> volatile cardinality
```

For information about regular tuning and best practices for DB2, refer to the corresponding IBM documentation.

Update Table Statistics Manually

In rare cases you may need to update statistics for a table manually. Run the following command for the table in the Sterling Integrator schema:

```
db2 runstats on table <table name> on key columns with distribution on key columns and sampled detailed indexes all allow read access
```

Microsoft SQL Server Configuration and Monitoring

This topic provides information about configuring and monitoring Microsoft® SQL Server® 2005 and Microsoft® SQL Server® 2008. It describes the recommended instance-specific settings, database-specific settings, maintenance plan, and system monitoring. It also provides information about Address Windowing Extensions (AWE), storage subsystem, dynamic management views, and index and table statistics.

Recommended Instance-Specific Settings for Microsoft SQL Server

The default collation of Microsoft® SQL Server® should match the collation for the Sterling Integrator database to prevent collation conversions. A tempdb must be created with the same collation as that of Microsoft SQL Server's default collation. The Microsoft SQL Server uses tempdb for results that are too large to fit in memory. If the collation is different between the tempdb and the Sterling Integrator database, the engine must convert from the Sterling Integrator collation to the tempdb collation, and then back again before sending the results to the Sterling Integrator server. This may lead to severe performance issues. In order to avoid such a situation, the server's default collation and the Sterling Integrator database's collation should be set to SQL_Latin1_General_CP850_Bin.

The parameter settings described in the following table are recommended:

Parameter	Value	Mandatory/ Recommended
Server Collation	SQL_Latin1_General_CP850_Bin	Mandatory
Max server memory (MB)	500 MB to x MB depending on the amount of physical memory available on your database server. if the server is running only this SQL Server instance; x can be up-to 80% of the physical memory (RAM).	Recommended
Min server memory (MB)	0	Recommended

Additionally, you must also perform these tasks:

- Allow Microsoft SQL Server 2005 or Microsoft SQL Server 2008 to manage memory dynamically (default).
- Disable any antivirus software running on the Microsoft SQL Server data, transaction log, and binaries directory.

Microsoft SQL Server Memory with Address Windowing Extensions (AWE)

You can use the Microsoft SQL Server memory (with AWE) when running Microsoft SQL Server (32 bit) on Windows Server (32 bit). This is helpful when your server has more than 4 GB of RAM and you want Microsoft SQL Server to be able to use that. For using Microsoft SQL Server with AWE, use the recommended values for Min and Max server memory for instance level settings. For more information on Enabling AWE Memory for Microsoft SQL Server Memory, refer to the Microsoft Developer Network Web site, which can be accessed from: [http://msdn.microsoft.com/en-us/library/ms190673\(SQL.90\).aspx](http://msdn.microsoft.com/en-us/library/ms190673(SQL.90).aspx).

Recommended DB-Specific Settings for Microsoft SQL Server

The collation required for the Sterling Integrator database is a collation that most closely matches the character set used by Java[®]. By using this collation, you can remove the necessity to perform character data conversions before they are stored in the database tables. Use the parameters described in the following table when configuring the DB-specific settings:

Parameter	Value	Mandatory / Recommended	Notes
Collation	SQL_Latin1_General_CP850_Bin	Mandatory	
IsAutoCreateStatistics	True	Recommended	This parameter can be set to False if you have a maintenance plan.
IsAutoUpdateStatistics	True	Recommended	This parameter can be set to False if you have a maintenance plan.
Page_verify_option	Checksum (Microsoft SQL Server default)	Recommended	
READ_COMMITTED_SNAPSHOT	On	Recommended	

Row Versioning-Based Isolation Level (READ_COMMITTED_SNAPSHOT) for Microsoft SQL Server

This new feature, which is available in Microsoft SQL Server 2005 and later versions, can help in the following ways:

- Resolve concurrency issues such as excessive blocking
- Reduce deadlocks

The following T-SQL statement enables the READ_COMMITTED_SNAPSHOT for a database:

```
ALTER DATABASE <DB NAME> SET READ_COMMITTED_SNAPSHOT ON;
```

This snapshot option increases the number of I/Os as well as the size of tempdb. It is important to have tempdb on fast disks as well as to have it sized according to your workload. For more information about Using and Understanding Snapshot Isolation and Row Versioning, refer to the Microsoft Developer Network Web site, which can be accessed from: [http://msdn.microsoft.com/en-us/library/tcbchxcb\(VS.80\).aspx](http://msdn.microsoft.com/en-us/library/tcbchxcb(VS.80).aspx).

Recommended Settings for Microsoft SQL Server 2008

The following table lists the recommended settings specific to Microsoft SQL Server 2008:

Feature	Description
Activity Monitor	During troubleshooting, a database administrator (DBA) executes several scripts or verifies number of sources to collect general information about the processes being executed and to find out the source of the problem. Microsoft SQL Server 2008 consolidates this information in detail graphically by running the recently executed processes.

Feature	Description
Data Compression	<p>The following list describes the two type of data compression supported by Microsoft SQL Server 2008:</p> <ul style="list-style-type: none"> • Row compression Row compression compresses the individual columns of a table. Row compression results in lower overhead on the application and utilizes more space. • Page compression Page compression compresses the data pages using row, prefix, and dictionary compression. Page compression affects application throughput and processor utilization, but requires less space. Page compression is a superset of row compression, which implies that an object or a partition of an object that is compressed using page compression is compressed at the row level too. <p>The amount of compression achieved is dependant on the data types and the data contained in the database.</p> <p>Compression, row or page, can be applied to a table or an index in an online mode without interrupting the availability of Sterling Integrator. The hybrid approach, where only the largest tables that are few in number, are compressed, results in best performance in saving significant disk space and resulting in minimal negative impact on performance. Disk space requirements should be considered before implementing compression. Compressing the smallest objects first minimizes the additional disk space requirements.</p> <p>Run the following SQL query to determine how compressing an object may affect its size:</p> <p>sp_estimate_data_compression_savings</p> <p>The following Sterling Integrator tables may be the most likely candidates for compression:</p> <ul style="list-style-type: none"> • DATA_TABLE • TRANS_DATA • CORRELATION_SET • WORKFLOW_CONTEXT • ARCHIVE_INFO <p>For more information about implementing row and page compression, refer to the Microsoft Developer Network Web site, which can be accessed from:</p> <ul style="list-style-type: none"> • http://msdn.microsoft.com/en-us/library/cc280576.aspx • http://msdn.microsoft.com/en-us/library/cc280464.aspx

Feature	Description
Hot Add CPU and Hot Add Memory	<p>Hot Add CPU enables you to add CPUs dynamically to the servers without shutting down the server or limiting client connections.</p> <p>Hot Add Memory enables you to add physical memory dynamically without restarting the server.</p> <p>For more information about dynamically adding CPU and physical memory, refer to the Microsoft Developer Network Web site, which can be accessed from:</p> <ul style="list-style-type: none"> • http://msdn.microsoft.com/en-us/library/bb964703.aspx • http://msdn.microsoft.com/en-us/library/ms175490.aspx
Extended Events	<p>The extended events infrastructure enables administrators to investigate and address complex problems such as excessive CPU usage, deadlocks, application timeouts, and so on. Extended events can be correlated with Windows events to obtain more information of the problem.</p> <p>For more information about extended events, refer to the Microsoft Developer Network Web site, which can be accessed from: http://msdn.microsoft.com/en-us/library/bb630354.aspx.</p>

Storage Subsystem

Prior to production, you should plan the overall storage strategy. You must have DDLs to create and place tempdb and the user database for Sterling Integrator. A typical subsystem configuration would be:

- OS and SQL Server binaries on a RAID 1 disk set
- SQL Server data files on one or more RAID 5 disk sets
- SQL Server transaction logs on a RAID 10 disk set

Consider and ensure the following when planning a storage subsystem:

- Place SQL Server binaries on a separate set of physical disks other than the database data and log files.
- Place the log files on physical disk arrays other than those with the data files. This is important because logging is more write-intensive, and the disk arrays containing the SQL Server log files require sufficient disk I/O to ensure that performance is not impacted.
- Set a reasonable size for your database. Estimate how big your database will be. This should be done as part of presales exercise working with Sterling Commerce Professional Services.
- Set a reasonable size for the transaction log. The transaction log's size should be 20–25 per cent of the database size.
- Leave the Autogrow feature on for the data files and the log files. This helps the SQL Server to automatically increase allocated resources when necessary.
- Set a reasonable size for the Autogrow increment. Setting the database to automatically grow results in some performance degradation. Therefore you should set a reasonable size for the Autogrow increment to prevent the database from growing automatically often.

- Set the maximum size for the data files and log files in order to prevent the disk drives from running out of space.
- If you have several physical disk arrays, try to create at least as many files as there are physical disk arrays so that you have one file per disk array. This improves performance because when a table is accessed sequentially, a separate thread is created for each file on each disk array in order to read the table's data in parallel.
- Place the heavily accessed tables in one file group and place the tables' indexes in a different file group on a different physical disk arrays. This improves performance, because separate threads will be created to access the tables and indexes. For more information about Sterling Integrator tables, refer to the "Schema Objects" and "Sterling Integrator Database Tables" sections in the topic "Database Management for Sterling Integrator" on page 19.

Monitoring Microsoft SQL Server Using Microsoft SQL Server Management Studio

Microsoft SQL Server Management Studio is a tool kit for configuring, managing, and administering all components of Microsoft SQL Server. Microsoft SQL Server Management Studio combines the features of Enterprise Manager, Query Analyzer, and Analysis Manager. For more information about Microsoft SQL Server Management Studio, refer to the Microsoft Developer Network Web site, which can be accessed from: [http://msdn.microsoft.com/en-us/library/ms174173\(SQL.90\).aspx](http://msdn.microsoft.com/en-us/library/ms174173(SQL.90).aspx).

Monitoring Microsoft SQL Server Using SQL Server Profiler

SQL Server Profiler is a graphical tool used to monitor an instance of Microsoft SQL Server. This tool is a good troubleshooting tool, but should not be enabled for day-to-day operations because there is an inherent overhead in capturing this data daily. The data about each event can be captured to a file or a table for analysis at a later date.

The SQL Server Profiler can be used to:

- Monitor the performance of an instance of the SQL Server Database Engine
- Identify procedures and queries that are executing slowly
- Replay traces of one or more users
- Perform query analysis
- Troubleshoot problems
- Audit and review activity
- Correlate performance counters
- Configure trace problems

For more information about Using the SQL Server Profiler, refer to the Microsoft Developer Network Web site, which can be accessed from: [http://msdn.microsoft.com/en-us/library/ms187929\(SQL.90\).aspx](http://msdn.microsoft.com/en-us/library/ms187929(SQL.90).aspx).

Enable the following events in SQL Server Profiler to capture deadlock-related information:

- Deadlock graph
- Lock: Deadlock
- Lock: Deadlock Chain
- RPC:Completed

- SP:StmtCompleted
- SQL:BatchCompleted
- SQL:BatchStarting

For more information about Analyzing Deadlocks with SQL Server Profiler, refer to the Microsoft Developer Network Web site, which can be accessed from: [http://msdn.microsoft.com/en-us/library/ms188246\(SQL.90\).aspx](http://msdn.microsoft.com/en-us/library/ms188246(SQL.90).aspx).

For more information about Troubleshooting Deadlocks in Microsoft SQL Server 2005, refer to the Microsoft Developer Network Web site, which can be accessed from: <http://blogs.msdn.com/sqlserverstorageengine/archive/2006/06/05/617960.aspx>.

Microsoft SQL Dynamic Management Views

The Dynamic Management Views (DMVs) introduced in Microsoft SQL Server 2005 provide DBA information about the current state of the SQL Server machine. These values help an administrator diagnose problems and tune the server for optimal performance. For more information about dynamic management views and functions, refer to the Microsoft Developer Network Web site, which can be accessed from:

[http://msdn.microsoft.com/en-us/library/ms188754\(SQL.90\).aspx](http://msdn.microsoft.com/en-us/library/ms188754(SQL.90).aspx)

Microsoft SQL System Monitor

The performance monitor (Perfmon) or system monitor is a utility used to track a range of processes and provide a real-time graphical display of the results. It can also be used to measure SQL Server performance. You can view SQL Server objects, performance counters, and the behavior of other objects, such as processors, memory, cache, threads, and processes. For more information about Monitoring Resource Usage (System Monitor), refer to the Microsoft TechNet Web site, which can be accessed from: [http://technet.microsoft.com/en-us/library/ms191246\(SQL.90\).aspx](http://technet.microsoft.com/en-us/library/ms191246(SQL.90).aspx).

Microsoft SQL Server Maintenance Plan

A maintenance plan is a set of measures (workflows) taken to ensure that a database is properly maintained and routine backups are scheduled and handled. Microsoft SQL Server maintenance plans can be configured by a wizard in Microsoft SQL Server Management Studio, which can help alleviate some of the burden involved in creating the plan. In Microsoft SQL Server Database Engine, maintenance plans create an Integration Services package, which is run by an SQL Server Agent job. The goal of a maintenance plan is to:

- Back up the Sterling Integrator database regularly using either the simple model or the full recovery model
- Update the statistics on all Sterling Integrator tables and associated indexes
- Rebuild or reorganize indexes on Sterling Integrator tables
- Run database consistency checks

For more information about MS SQL Maintenance Plans, see the Microsoft Developer Network Web site: [http://msdn.microsoft.com/en-us/library/ms187658\(SQL.90\).aspx](http://msdn.microsoft.com/en-us/library/ms187658(SQL.90).aspx)

You can use either the simple recovery model or the full recovery model with the database. If you are unsure about which model to use, use the full recovery model.

Consider the simple recovery model if you are agreeable to the following:

- Point of failure recovery is not necessary. If the database is lost or damaged, you are willing to lose all the updates between a failure and the previous backup.
- You are willing to risk the loss of some data in the log.
- You do not want to back up and restore the transaction log, preferring to rely exclusively on full and differential backups.
- You are willing to perform a differential database backup every day and a full database backup over the weekends (during low activity period).

For more information about Backup Under the Simple Recovery Model, see the Microsoft Developer Network Web site:

[http://msdn.microsoft.com/en-us/library/ms191164\(SQL.90\).aspx](http://msdn.microsoft.com/en-us/library/ms191164(SQL.90).aspx)

Consider the full recovery model if you are agreeable to the following:

- You want to recover all the data.
- You want to recover to the point of failure.
- You want to be able to restore individual pages.
- You are willing to incur the cost of additional disk space for transaction log backups.
- Performance of transaction log backups every 30-45 minutes.

For more information about Backup Under the Full Recovery Model, refer to the Microsoft Developer Network Web site, which can be accessed from:

[http://msdn.microsoft.com/en-us/library/ms190217\(SQL.90\).aspx](http://msdn.microsoft.com/en-us/library/ms190217(SQL.90).aspx).

For more information about Transaction Log Backups, refer to the Microsoft Developer Network Web site, which can be accessed from: [http://msdn.microsoft.com/en-us/library/ms190440\(SQL.90\).aspx](http://msdn.microsoft.com/en-us/library/ms190440(SQL.90).aspx).

Microsoft SQL Server Index, Table Statistics, and Index Rebuilds

Your maintenance plan should include updation of statistics and rebuilding or reorganizing indexes. If you have a maintenance plan for updating statistics, set `IsAutoCreateStatistics` and `IsAutoUpdateStatistics`, at the database level, to `False`. This helps control when the maintenance plan runs, which should be at low load periods. If you have not included this in your maintenance plan, then set `IsAutoCreateStatistics` and `IsAutoUpdateStatistics`, at the database level to `True`.

Following is an example of a T-SQL statement for updating the statistics on a table, including the statistics pertaining to all the indexes on the table:

```
UPDATE STATISTICS <TABLE NAME>
```

When the data in the tables have changed by approximately 20 per cent, index rebuilds are recommended for the Sterling Integrator database. This must be validated by looking at system performance as a whole. Online index rebuilds, which cause minimal impact to the system, are possible. To find out more about document life-spans, and when the data in the tables change by approximately 20 per cent, speak to the Sterling Integrator system administrator.

You can rebuild indexes either online or offline. Online indexes can be rebuilt, with the following exceptions:

- Clustered indexes if the underlying table contains LOB data types
- Nonclustered indexes that are defined with LOB data type columns. Nonclustered indexes can be rebuilt online if the table contains LOB data types, but none of these columns are used in the index definition as either key or nonkey columns.

For ease of maintenance, it is easier to either build all the indexes offline because offline rebuilding does not have the restrictions listed previously, or reorganize the index.

Following is an example of a T-SQL statement for rebuilding indexes offline:

```
ALTER INDEX ALL ON <TABLE NAME> REBUILD
```

Following is an example of a T-SQL statement for reorganizing indexes offline:

```
ALTER INDEX ALL ON <TABLE NAME> REORGANIZE
```

For more information about Reorganizing and Rebuilding Indexes, refer to the Microsoft Developer Network Web site, which can be accessed from: [http://msdn.microsoft.com/en-us/library/ms189858\(SQL.90\).aspx](http://msdn.microsoft.com/en-us/library/ms189858(SQL.90).aspx).

For more information about Alter index (Transact-SQL), refer to the Microsoft Developer Network Web site, which can be accessed from: [http://msdn.microsoft.com/en-us/library/ms188388\(SQL.90\).aspx](http://msdn.microsoft.com/en-us/library/ms188388(SQL.90).aspx).

Windows Defragmentation

You must run Windows defragmentation on disks with SQL Server data files and transaction logs once a month. This helps reduce the fragmentation in the SQL Server files at the file system level. You can create a schedule for this using Windows Task Scheduler. For more information about the Disk Defragmenter Tools and Settings, refer to the Microsoft TechNet Web site, which can be accessed from: [http://technet.microsoft.com/en-us/library/cc784391\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/cc784391(WS.10).aspx).

Microsoft SQL Server Tips

Following are some additional tips pertaining to Microsoft SQL Server in the context of Sterling Integrator:

- When using Sterling Integrator with Microsoft SQL Server, Windows Integrated authentication is not supported.
- Ensure that network components such as routers, firewalls, and so on, do not drop the idle connections between Sterling Integrator and Microsoft SQL Server where they are on separate physical servers. Sterling Integrator uses JDBC connection pool, and idle connections are typical.
- It is recommended that you run Microsoft SQL Server and Sterling Integrator on separate physical servers because this helps improve performance, ease of maintenance, and recoverability.
- It is important to understand the difference between simple blocking and deadlocks:
 - Blocking is an unavoidable characteristic of Microsoft SQL Server because it uses lock-based concurrency. Blocking occurs when one session holds a lock on a specific resource, and a second session attempts to acquire a conflicting

lock type on the same resource. Typically, the time frame for which the first session locks the resource is small. When it releases the lock, the second session is free to acquire its own lock and continue processing. This is the normal behavior of Microsoft SQL Server with Sterling Integrator, and is generally not a cause for concern. It is a cause for concern only when sessions are getting blocked for a long time.

Deadlocks are much worse than simple blocking. A deadlock typically occurs when a session locks the resources that another session has to modify, and the second session locks the resources that the first session intends to modify. Microsoft SQL Server has a built-in algorithm for resolving deadlocks. It will select one of the deadlock participants and roll back its transaction. This session becomes the deadlock victim. Microsoft SQL Server has two trace flags that can be set to capture deadlock-related information. The flags are Trace Flag 1204 and Trace Flag 1222. These trace flags can be used as an alternative to using SQL Server Profiler.

For more information about Detecting and Ending Deadlocks, refer to the Microsoft Developer Network Web site, which can be accessed from: [http://msdn.microsoft.com/en-us/library/ms178104\(SQL.90\).aspx](http://msdn.microsoft.com/en-us/library/ms178104(SQL.90).aspx).

Chapter 4. Java Virtual Machines

Overview of JVM

Java Virtual Machine (JVM) is a platform-independent programming language that converts Java bytecode into machine language and executes it. When you compile a Java source, you get an intermediate Java file called the Java class. The class file is made up of bytecodes representing abstract instruction codes. These codes are not directly executable by any computer processor.

To run a Java program, you start a JVM and pass the class file to the JVM. The JVM provides many services, including loading the class file and interpreting (executing) the byte codes. The JVM is the core technology that provides the runtime environment in which a Java application runs.

Each Java program or application runs in its own JVM. For example, if you configured an application server cluster with ten managed server instances that are controlled by one administrative instance, your configuration runs 11 JVM processes.

Since JVM is the underlying processing engine, it is critical that the JVMs are optimally configured and are running efficiently. Incorrect JVM settings may lead to poor application performance or JVM outages too.

Run the following command to find out the JVM version installed in your system:

```
$JAVA_HOME/bin/java -version
```

View the Performance Tuning Configuration Settings for ASI Nodes

About this task

To view the JVM performance tuning configuration settings for application server-independent (ASI) nodes:

Procedure

1. From the **Administration** menu, select **Operations > System > Performance > Tuning**.
2. In the Performance Tuning page, under **View**, next to **View Performance Configuration**, click **Go!**
3. Review the performance settings. Performance settings includes information on:
 - Number of CPUs
 - Physical memory (MB) allocated to SI
 - BP Queue
 - Database Pool
 - Memory
 - BP Execution
 - Cache
4. Click **Return** when finished.

Note: You can view container JVM parameters in `install_dir/bin/tmp.sh_platform_ifcresources_ext.in` file (for Windows - `install_dir\bin\InstallContainerWindowsService.cmd.in`).

View JVM Settings in the `tuning.properties` file

About this task

To view JVM settings in the `tuning.properties` file:

Procedure

1. Locate the `tuning.properties` file. (The file should be located in the `install_dir/properties` directory.)
2. Review the file using a file editor.
3. Look for the following properties with reference to the appropriate host OS:

Note: The 32-bit JVM attributes are prefixed with OS name, for example, `HP-UX.INIT_HEAP`. The 64-bit JVM attributes are prefixed with OS name and architecture (64-bit), for example, `HP-UX-64.INIT_HEAP`.

JVM Attribute	Attribute Description
<code>INIT_HEAP</code>	Initial amount of JVM long-lived memory reserved at startup
<code>MAX_HEAP</code>	Maximum amount of JVM long-lived memory that can be reserved
<code>INIT_AGE</code>	Initial amount of JVM short-lived memory
<code>MAX_AGE</code>	Maximum amount of JVM short-lived memory

Edit Performance Tuning Configuration Settings for ASI Node

About this task

To edit performance tuning configuration settings for application server-independent (ASI) node:

Procedure

1. From the Administration menu, select **Operations > System > Performance > Tuning**.
2. In the Performance Tuning page, under **Edit**, next to **Edit Performance Configuration**, click **Go!**
3. When prompted with the following message, click **OK**. Performance Tuning UI lock enabled message, click **OK**. The Performance Tuning UI lock releases when you click **Cancel** or **Finish** during the edit 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. In the Edit Mode page, select **Edit settings** and click **Next**.
Sterling Integrator displays the current setting for each property (system, BP queue, database pool, memory, BP execution, cache). You can either keep the current setting or enter a new value based on your performance requirements.
5. After completing the necessary modifications for each property, click **Next**.
6. In the Performance Tuning: Confirm page, click **Finish** to retain the modified settings.
7. Stop Sterling Integrator ASI node.

8. From the *install_dir/bin* directory, run one of the following commands:
 - (UNIX or Linux) **setupfiles.sh**
 - (Windows) **setupfiles.cmd**
9. Start Sterling Integrator ASI node.

Edit Performance Tuning Configuration Settings for Container JVMs

About this task

To edit performance configuration settings for container JVMs:

Procedure

1. Navigate to the *install_dir/bin* (for Windows - *install_dir\bin*) directory and modify either one of the following files depending on your operating system:
 - (For UNIX) - *tmp.sh_platform_ifcresources_ext.in*
 - (For Windows) - *InstallContainerWindowsService.cmd.in*
2. Modify the JVM parameter settings, for example,

Existing	New
<code>echo -Xms256m >> %PARAM_FILE%</code>	<code>echo -Xms1024m >> %PARAM_FILE%</code>
<code>echo -Xmx256m >> %PARAM_FILE%</code>	<code>echo -Xmx1024m >> %PARAM_FILE%</code>

Save the file.

3. Stop the container by running one of the following commands:
 - (For UNIX) - **./stopContainer.sh**
 - (For Windows) - **stopContainerWindowsService.cmd**
4. Apply the changes by running one of the following commands:
 - (For UNIX) - **setupfiles.sh**
 - (For Windows) - **setupfiles.cmd**

Run the following scripts in order:

- **uninstallContainerWindowsService.cmd** *<Container_Number>*
- **InstallContainerWindowsService.cmd** *<Container_Number>*
- **InstallACPassPhraseWindowsService.cmd**

The changes are applied to all container JVMs.

5. Start the container by running one of the following commands:
 - (For UNIX) - **./startContainer.sh**
 - (For Windows) - **startContainerWindowsService.cmd**

Garbage Collection Statistics

Garbage Collection (GC) statistics provide heap-related information such as:

- What are the sizes of the different heaps?
- How full is each section of heap?
- How fast is the heap getting full?
- What is the overall overhead of GC to clean the non-live objects?

Collecting and analyzing GC statistics help size the different sections of a heap correctly. It is recommended that you continuously collect garbage collection statistics for all the JVMs, even in production. The collection overhead is minor compared to the benefit. With these statistics, you can tell if:

- The JVM has or is about to run into a memory leak
- Garbage collection is efficient
- Your JVM heap settings are optimal

JVM Verbose Garbage Collection

JVM Verbose garbage collection (GC) statistics are critical and must always be enabled in production environments. These statistics can be used to understand the behavior of the JVM heap management and the efficiency of the JVM.

IBM JVM Garbage Collection Example

Following is an example of the GC output for JVM:

```
<af type="tenured" id="100" timestamp="Sun Nov 25 15:56:09 2007"
intervalms="120245.593">
  <minimum requested_bytes="10016" />
  <time exclusiveaccessms="0.045" />
  <tenured freebytes="2704" totalbytes="1073741824" percent="0" >
    <soa freebytes="2704" totalbytes="1073741824" percent="0" />
    <loa freebytes="0" totalbytes="0" percent="0" />
  </tenured>
  <gc type="global" id="100" totalid="100" intervalms="120245.689">
    <refs_cleared soft="0" threshold="32" weak="0" phantom="0" />
    <finalization objectsqueued="0" />
    <timesms mark="35.301" sweep="5.074" compact="0.000" total="40.426" />
    <tenured freebytes="808526296" totalbytes="1073741824" percent="75" >
      <soa freebytes="808526296" totalbytes="1073741824" percent="75" />
      <loa freebytes="0" totalbytes="0" percent="0" />
    </tenured>
  </gc>
  <tenured freebytes="808516280" totalbytes="1073741824" percent="75" >
    <soa freebytes="808516280" totalbytes="1073741824" percent="75" />
    <loa freebytes="0" totalbytes="0" percent="0" />
  </tenured>
  <time totalms="40.569" />
</af>
```

In this example, `<af type="tenured" id="100"` indicates that this is the 100th time an attempt to allocate memory has failed, and as a result, a GC was initiated. An allocation failure is not an error in the system or code. When enough free space is not available in the heap, the JVM automatically initiates a garbage collection. The last time an allocation failure occurred was 120245.593 milliseconds ago (or 120.245 seconds).

The lines starting with `<gc type="global" id="100"` provide information about the collection process. In this example, garbage collection initiated the mark phase and the sweep phase, which were completed in 35.301 and 5.074 milliseconds respectively. The JVM determined that the heap was not fragmented, and that compacting the heap was not required. At the end of the GC, the heap had 808,516,280 bytes of available memory.

It is important that the frequency of GCs be monitored. This can easily be achieved by looking at the time between the allocation failures. Typically, a healthy JVM will spend less than 0.5 seconds in each GC cycle. Also, the overall percentage of time spent on garbage collection should be less than 3 percent. To calculate the percent of time spent performing garbage collection, divide the sum of the garbage collection time over a fixed interval by the fixed interval.

IBM provides documentation pertaining to its Garbage Collector and how to interpret its GC statistics.

Refer to either the *IBM JDK 6.0: Java Diagnostics Guide* or the IBM DeveloperWorks article "Fine-tuning Java garbage collection performance", which is available at the following Web site:

<http://www.ibm.com/developerworks/ibm/library/i-gctroub/>

Introduction to HotSpot JVM Performance and Tuning

The Sun HotSpot JVM is used when you deploy Sterling Integrator on a Sun Solaris operating system running on Sun UltraSPARC processor-based servers and on a Windows operating system.

The HP HotSpot JVM is used when you deploy Sterling Integrator on a HP-UX for IA64 (Itanium) or HP-UX on any other processor.

HotSpot JVMs provide many tuning parameters. There is no golden set of JVM settings that apply to all customers and conditions. Fortunately, the HotSpot JVMs provide good measurement feedback that allows you to measure the effectiveness of the settings. The settings, especially memory settings, are highly dependent on:

- Transaction mix
- Amount of data cached
- Complexity of the transactions
- Concurrency levels


This topic describes the processes involved in planning, implementing, configuring, monitoring, and tuning the HotSpot Java Virtual Machines.

HotSpot JVM Heap Memory and Garbage Collection

The JVM run-time environment uses a large memory pool called the heap, for object allocation. The JVM automatically invokes garbage collections (GC) to clean up the heap of unreferenced or dead objects. In contrast, memory management in legacy programming languages such as C++ is left to the programmer. If the JVM heap settings are not set correctly, the garbage collection overheads can make the system appear unresponsive. In the worst case, your transactions or the JVM may abort due to `outOfMemory` exceptions.

Garbage collection techniques are constantly being improved. For example, the Sun JVM supports a "stop-the-world" garbage collector where all the transactions have to pause at a safe point for the entire duration of the garbage collection. The Sun JVM also supports a parallel concurrent collector, where transactions can continue to run during most of the collection.

The Sun heap and HP heap are organized into generations to improve the efficiency of their garbage collection, and reduce the frequency and duration of user-perceivable garbage collection pauses. The premise behind generational collection is that memory is managed in generations or in pools with different ages. The following diagram illustrates the layout of the generational heap.



At initialization, a maximum address space is virtually reserved, but not allocated, to physical memory unless it is needed. The complete address space reserved for object memory can be divided into young and tenured (old) generations.

New objects are allocated in the Eden. When the Eden fills up, the JVM issues a scavenge GC or minor collection to move the surviving objects into one of the two survivor or semi spaces. The JVM does this by first identifying and moving all the referenced objects in the Eden to one of the survivor spaces. At the end of the scavenge GC, the Eden is empty (since all the referenced objects are now in the survivor space) and ready for object allocation.

The scavenge GC's efficiency depends on the amount of referenced objects it has to move to the survivor space, and not on the size of the Eden. The higher the amount of referenced objects, the slower the scavenge GC. Studies have, however, shown that most Java objects live for a short time. Since most objects live for a short time, one can typically create large Edens.

Referenced objects in the survivor space bounce between the two survivor spaces at each scavenge GC, until it either becomes unreferenced or the number of bounces have reached the tenuring threshold. If the tenuring threshold is reached, that object is migrated up to the old heap.

When the old heap fills up, the JVM issues a Full GC or major collection. In a Full GC, the JVM has to first identify all the referenced objects. When that is done, the JVM sweeps the entire heap to reclaim all free memory (for example, because the object is now dead). Finally, the JVM then moves referenced objects to defragment the old heap. The efficiency of the Full GC is dependent on the amount of referenced objects and the size of the heap.

The HotSpot JVM sets aside an area, called permanent generation, to store the JVM's reflective data such as class and method objects.

HotSpot JVM Garbage Collection Tools

If you want to analyze the Garbage Collection (GC) logs, use some of the tools described in the following table:

Tool Name	For Additional information, Refer To
GCViewer	http://www.tagtraum.com/gcviewer.html

Tool Name	For Additional information, Refer To
IBM Pattern Modeling and Analysis Tool for Java Garbage Collector	http://www.alphaworks.ibm.com/tech/pmat
visualgc	http://java.sun.com/performance/jvmstat/visualgc.html Note: Visualgc can be tied to running JVM processes at any time. It is helpful if you have not enabled the GC flags, but want to take a look at the heaps and GC overheads, and you do not want to restart your JVM.
HPjmeter	https://h20392.www2.hp.com/portal/swdepot/displayProductInfo.do?productNumber=HPJMETER

HotSpot JVM Startup and Runtime Performance Optimization

In some cases, an application's startup performance is more important than its runtime performance. Applications that start once and run for a longer period should be optimized for runtime performance. By default, HotSpot JVMs are optimized for startup performance.

The Java Just-In-Time (JIT) compiler impacts the startup and runtime performance. The time taken to compile a class method and start the server are influenced by the initial optimization level used by the compiler. You can reduce the application startup times by reducing the initial optimization level. This degrades your runtime performance because the class methods will now compile at the lower optimization level.

It is not easy to provide a specific runtime performance impact statement, because compilers may recompile the class methods based on the impression that recompiling provides better performance. Short-running applications will have their methods recompiled more often than long-running applications.

HotSpot JVM Default Parameters

The following tuning options are configured and shipped out-of-the-box in Sterling Integrator on HotSpot JVMs in Windows, Solaris, and HP-UX.

The options, `-d32` and `-d64` are added to the Java launcher to specify if the program should be run in a 32-bit or a 64-bit environment. On Solaris, they correspond to the ILP32 and LP64 data models. Since Solaris contains both 32-bit and 64-bit J2SE implementation within the same Java installation, you can specify either version. If neither `-d32` nor `-d64` is specified, it will run in 32-bit environment by default. Java commands such as `javac`, `javadoc`, and so on may rarely need to be run in a 64-bit environment. However, it may be required to pass the `-d32` or `-d64` options to these commands and then to the Java launcher using `-J` prefix option, for example, `-J-d64`. For other platforms such as Windows and Linux, 32-bit and 64-bit installation packages are separate. If you have installed both 32-bit 64-bit packages on a system, you can select either one of them by adding the appropriate bin directory to the path.

Sterling Integrator 5.1 supports 64-bit architecture on most operating systems. The `-d64` option is the default value on Windows Server 2008, Solaris, and HP-UX for both noapp JVM and container JVM. However, in Windows Server 2003, the `-d32` option is default as it uses a 32-bit JDK.

Parameter	Description	Default Value
-server	HotSpot-based JVMs generally use low optimization levels, which takes less time to start up, but leads to low runtime performance. Normally, a simple JIT compiler is used. To increase the runtime performance for applications such as Sterling Integrator, an optimizing compiler is recommended. Using this method may, however, lead a JVM to take longer time to warm up.	For both noapp and container JVMs: <ul style="list-style-type: none"> • Solaris = -server (optimizing compiler) • HP-UX = -server (optimizing compiler)
-Xms	<p>Controls the initial size of a Java heap.</p> <p>If this is tuned correctly, it can reduce the overhead associated with garbage collection by improving the server response time and throughput.</p> <p>If you see a large number of minor garbage collections, the default value set to this parameter might be too low. Try increasing it.</p>	<p>For 32-bit noapp JVM:</p> <ul style="list-style-type: none"> • Windows = 768m • Solaris = 512m • HP-UX = 1024m <p>For 64-bit noapp JVM:</p> <ul style="list-style-type: none"> • Windows = 1280m • Solaris = 1280m • HP-UX = 1280m <p>For 64-bit container JVM:</p> <ul style="list-style-type: none"> • Windows = 64m • Solaris = 384m • HP-UX = 384m

Parameter	Description	Default Value
-Xmx	<p>Controls the maximum size of a Java heap.</p> <p>If this parameter is tuned correctly, it can:</p> <ul style="list-style-type: none"> • Reduce the overhead associated with the garbage collection and the risk of encountering an Out-Of-Memory (OOM) condition • Improve the server response time and throughput <p>If you see a large number of garbage collections, try increasing the value. You can set a maximum heap limit of 4 GB for a 32-bit JVM. However, due to various constraints such as available swap, kernel address space usage, memory fragmentation, and VM overhead, it is recommended to set a lower value. In 32-bit Windows systems, the maximum heap size can be set in the range from 1.4 GB to 1.6 GB. Similarly, on 32-bit Solaris kernels, the address space is limited to 2 GB. The maximum heap size can be higher if your 64-bit operating system is running 32-bit JVM, reaching until 4 GB on Solaris systems. Java SE 6 does not support Windows /3GB boot.ini feature. If you require a large heap setting, you should use a 64-bit JVM on an operating system supporting 64-bit applications.</p>	<p>For 32-bit noapp JVM:</p> <ul style="list-style-type: none"> • Windows = 1024m • Solaris = 1024m • HP-UX = 768m <p>For 64-bit noapp JVM:</p> <ul style="list-style-type: none"> • Windows = 1280m • Solaris = 1280m • HP-UX = 1280m <p>For 64-bit container JVM:</p> <ul style="list-style-type: none"> • Windows = 256m • Solaris = 512m • HP-UX = 512m
-XX:+DisableExplicitGC	<p>Disables the explicit garbage collection calls that are caused by invoking System.gc() inside the application.</p> <p>It is recommended that the developers avoid the System.gc() calls to cause programmer-initiated, full-compaction garbage collection cycles, because such calls can interfere with the tuning of resources and garbage collection for the entire application system. If your application pause times caused by System.gc() are more than your expected pause times, it is strongly recommended that you use this option to disable the explicit GC, so that the System.gc() calls will be ignored.</p>	<p>For both 32-bit/64-bit noapp and container JVMs:</p> <ul style="list-style-type: none"> • Windows = -XX:+DisableExplicitGC • Solaris = -XX:+DisableExplicitGC • HP-UX = -XX:+DisableExplicitGC

Parameter	Description	Default Value
-XX:NewSize= and -Xmn<Size>=	<p>-XX:NewSize controls the minimum young generation size in a heap, and the -Xmn sets the size of the young generation heap.</p> <p>It is recommended to use the following formula to compute -XX:NewSize and -Xmn values when the minimum heap size and maximum heap sizes are modified.</p> <ul style="list-style-type: none"> • -XX:NewSize = (0.33333* value of -Xms) • -Xmn = (0.33333 * value of -Xmx) 	<p>For 32-bit JVMs:</p> <ul style="list-style-type: none"> • Solaris <ul style="list-style-type: none"> -Xmn400m -XX:NewSize=300m • HP-UX <ul style="list-style-type: none"> -Xmn341m -XX:NewSize=341m <p>For 64-bit JVMs:</p> <ul style="list-style-type: none"> • Solaris <ul style="list-style-type: none"> -Xmn341m -XX:NewSize=341m • HP-UX <ul style="list-style-type: none"> -Xmn341m -XX:NewSize=341m <p>For 64-bit container JVMs:</p> <ul style="list-style-type: none"> • Solaris <ul style="list-style-type: none"> -Xmn256m -XX:NewSize=64m • HP-UX <ul style="list-style-type: none"> -Xmn256m -XX:NewSize=64m
-XX:MaxPermSize	<p>Stores all the class code and class-like data. The value of the parameter should be large enough to fit all the classes that are concurrently loaded. Sometimes, it is difficult to determine the actual value of this parameter because generally, this region is smaller and expands slowly, and the utilization is commonly observed at 99-100 percent of its current capacity. If you have not configured this region correctly, the JVM might fail with the <code>Java.lang.OutOfMemoryError: PermGen space error</code>.</p>	<p>For 32-bit/64-bit noapp JVMs:</p> <ul style="list-style-type: none"> • Windows = 256m • Solaris = 256m • HP-UX = 384m <p>For 64-bit container JVM:</p> <ul style="list-style-type: none"> • Windows = 128m

Parameter	Description	Default Value
-Xss	<p>Determines the stack size for each thread in the JVM.</p> <p>Every thread in a JVM gets a stack, and this value determines the number of threads you can start in a JVM. If this value is too large, you might run into an OOM. Each time a method is invoked, a stack frame is created and pushed into the thread stack. At a minimum, a stack frame contains a method's local variables and arguments. If a thread's actual stack size reaches beyond this limit, you will get a <code>java.lang.StackOverflowError</code> exception.</p>	<p>For 32-bit noapp JVM:</p> <ul style="list-style-type: none"> • Windows=256k • Solaris=256k • HP-UX=256k • HP-UX(IA64)=1024k <p>For 64-bit noapp JVM:</p> <ul style="list-style-type: none"> • Windows=1024k • Solaris=512k • HP-UX=256k • HP-UX(IA64)=1024k <p>For 64-bit container JVM:</p> <ul style="list-style-type: none"> • Windows=256k • Solaris=256k • HP-UX=384k • HP-UX(IA64)=384k
-XX:CompileThreshold	<p>Determines the number of method invocations and branches before compiling.</p> <p>Setting a low value for this parameter will trigger the compilations of hot methods sooner.</p>	<p>For both 32-bit/64-bit noapp and container JVMs:</p> <ul style="list-style-type: none"> • Windows = -XX:CompileThreshold=1000 • Solaris = -XX:CompileThreshold=1000
-Xnocatch	<p>Disables the Java catch-all signal handler. This option is used to get clean, native code stack traces.</p>	<p>For both noapp and container JVMs:</p> <p>HP-UX = -Xnocatch</p>

Enable JVM Verbose Garbage Collection

About this task

To enable Verbose Garbage Collection for IBM JVM and HotSpot JVM for noapp JVM:

Procedure

1. From the Administration menu, select **Operations > System > Performance > JVM Monitor**.
2. In the JVM Monitor page, select the check box next to **Enable GCOutput**. This adds the `-verbose:gc` flag to your JVM arguments, and the gc output will be logged to `install/logs/noapp.log`. On Windows, to collect GC statistics, uncomment the following lines in the `InstallNoappsWindowsService.cmd.in` by removing `rem`. You will see `verbosegc.log` under the `<install>/logs` directory if you use the following parameters:

```

echo -verbose\;gc >> %PARAM_FILE%
set out_path=&LOG_DIR_WIN;\verbosegc.log
set out_path=%out_path:\=\%
set out_path=%out_path:;=\%
echo -Xloggc\;%out_path% >> %PARAM_FILE%
echo -XX\:+PrintGCDetails >> %PARAM_FILE%
echo -XX\:+PrintGCTimeStamps >> %PARAM_FILE%

```
3. Stop Sterling Integrator.

4. If you are using Linux or UNIX, using a terminal window, navigate to the *install_dir/bin* directory and enter the following commands:

```
setupFiles.sh  
run.sh
```

Note: On UNIX and Linux platforms, to add any additional JVM parameters, edit *tmp.sh*, and add the parameters to the *SERVER_FLAGS* variable under the respective OS. However, it is recommended that you contact Sterling Customer Support prior to editing this script directly.

5. If you are using Windows, using a terminal window, navigate to the *install_dir/bin* directory and enter the following commands:

```
setupFiles.cmd  
uninstallWindowsService.cmd  
InstallWindowsService.cmd  
startWindowsService.cmd
```

Note:

On Windows, to add any additional flag, edit **InstallNoappsWindowsService.cmd.in**, and add a flag using the syntax shown in the following example:

```
"echo -XX\:+PrintGCTimeStamps >> %PARAM_FILE%"
```

However, it is recommended that you contact Sterling Customer Support prior to editing this script directly.

Display HotSpot Statistics for HP JVM

About this task

To display HotSpot Statistics for the HP JVM, run one of the following commands:

- `java -XX:+PrintGCDetails -XX:+PrintGCTimeStamps -Xloggc:<gcfilename> -verbose:<gc>`

This displays the following:

```
11.010: [GC [PSYoungGen: 196608K->20360K(229376K)] 196608K->  
>20360K(753664K), 0.0514144 secs] 13.987: [GC [PSYoungGen:  
216968K->32746K(229376K)] 216968K->48812K(753664K), 0.1052434 secs]
```

- `java -Xverbosegc [:help] | [0 | 1] [:file = [stdout | stderr | <filename>]]`

This displays the following:

```
<GC: 1 4 11.988605 1 64 7 201326592 64 201326592 0 20850824 33554432 0 0  
536870912 15563792 15563792 21757952 0.040957 0.040957 > <GC: 1 4  
13.400027 2 864 7 201326592 864 201326592 20850824 33539216 33554432 0  
16479936 536870912 17504224 17504224 21757952 0.088071 0.088071 >
```

Refer to the corresponding HP documentation , which can be accessed from the following Web site:

<http://docs.hp.com/en/5992-5899/ch06s01.html>

Enable Verbose Garbage Collection for Container JVM

About this task

To enable Verbose Garbage Collection for container JVM:

Procedure

1. Navigate to the *install_dir/bin* (for Windows - *install_dir\bin*) directory and modify either one of the following files depending on your operating system:
 - (For UNIX) - *tmp.sh_platform_ifcresources_ext.in*
 - (For Windows) - *InstallContainerWindowsService.cmd.in*
2. Modify the JVM parameter settings:
 - (For UNIX) - Add the following parameter and any other garbage collection related flags to OS and *SERVER_ACFLAGS* variables.
-verbose:gc
 - (For Windows) - Add the following lines after JVM memory arguments, for example, after **echo -Xmx<NNNN>m >> %PARAM_FILE%:**

```
rem "Verbose GC Parameters"
echo -verbose\:gc >> %PARAM_FILE%
set gc_out_path=&LOG_DIR_WIN;\&NODE_NAME;AC%1\verbosegc.log
set gc_out_path=%gc_out_path:\=\%
set gc_out_path=%gc_out_path:.\=%
echo -Xloggc\:%gc_out_path% >> %PARAM_FILE%
echo -XX\:+PrintGCDetails >> %PARAM_FILE%
echo -XX\:+PrintGCTimeStamps >> %PARAM_FILE%
```

Save the file.

3. Stop the container by running one of the following commands:
 - (For UNIX) - **./stopContainer.sh**
 - (For Windows) - **stopContainerWindowsService.cmd**
4. Apply the changes by running the following command(s):
 - (For UNIX) - **setupfiles.sh**
 - (For Windows) - **setupfiles.cmd**Run the following scripts in the same order against each container number:
 - **stopContainerWindowsService.cmd** <Container_Number>
 - **uninstallContainerWindowsService.cmd** <Container_Number>
 - **InstallContainerWindowsService.cmd** <Container_Number>
 - **InstallACPassPhraseWindowsService.cmd** <Container_Number>
 - **startContainerWindowsService.cmd** <Container_Number>

The changes are applied to all container JVMs.

5. Start the container by running one of the following commands:
 - (For UNIX) - **./startContainer.sh**
 - (For Windows) - **startContainerWindowsService.cmd**

Display HotSpot Statistics for Sun JVM

About this task

To display HotSpot Statistics for the Sun JVM enable one of the following command:

```
-XX:+PrintGCDetails -XX:+PrintGCTimeStamp -Xloggc:<gcfilename>
```

This displays the following:

```
0.000: [GC 0.001: [DefNew: 32192K->511K(33152K), 0.0383176 secs]
32192K->511K(101440K), 0.0385223 secs] 1.109: [GC 1.110: [DefNew:
32703K->198K(33152K), 0.0344874 secs] 32703K->697K(101440K), 0.0346844
secs]
```

Refer to the corresponding Sun documentation, which can be accessed from the following Web site:

http://java.sun.com/javase/technologies/hotspot/gc/gc_tuning_6.html

HotSpot JVM Heap Monitoring

You may want to monitor the following items in a healthy heap:

- During steady state, you should mostly see minor Garbage Collections (GC) and an occasional full GC caused by allocation failures.
- The sum of the GC times should not exceed 3 percent of the measurement interval. For example, in a 1-hour measurement interval, the time taken for all the GCs should not be more than 108 seconds.
- The JVM will choose to perform a full GC when it realizes that the live objects in the Eden and Survivor spaces cannot fit in the old generation. The JVM then tries to free up the space in the old generation by performing a full GC. The full GC pauses the application. The amount of pause time depends on the GC algorithm you are using (Sterling Integrator uses the default JVM GC algorithm, that is, Parallel GC algorithm on 1.6 JDK on server class machines), and the size of the heap. The JVM will choose to perform a full GC when it realizes that the live objects in the Eden and Survivor spaces will not fit into the old generation. In an attempt to free up the space in the old generation, the JVM will perform a full GC, which pauses the application. The amount of pause time depends on:
 - Type of GC algorithm currently in use. Sterling Integrator uses the default JVM GC algorithm, that is, Parallel GC algorithm, on 1.6 JDK on server class machines.
 - Size of the heap. Too many full GCs have a negative effect on performance.

If you observe many full GCs, try to determine if your old generation is sized too small to hold all the live objects collected from the Survivor and Eden spaces. Alternatively, there may be too many live objects that do not fit into the configured heap size. If it is the latter, increase the overall heap size.

- If you are monitoring the heap and notice an increase in the number of live objects, and see that the GC is not able to clear these objects from the heap, you might run into an Out-Of-Memory (OOM) condition, and there may be a possible memory leak. In such a situation, take a heap dump at various intervals and analyze the dump for the leak suspect.

HotSpot JVM Thread Monitoring

Monitoring JVM threads will help you locate thread deadlocks if there are blocked threads in a Java code or in SQL. It also helps you to understand which part of the code the threads are blocked in a hung application or running application.

Following is an example of a thread dump:

```
"Thread-817" daemon prio=10 tid=0x00b9c800 nid=0x667 in Object.wait()
[0xa4d8f000..0xa4d8faf0]
  at java.lang.Object.wait(Native Method)
    - waiting on <0xf678a600> (a EDU.oswego.cs.dl.util.concurrent.LinkedNode)
  at EDU.oswego.cs.dl.util.concurrent.SynchronousChannel.poll(
SynchronousChannel.java:353)
    - locked <0xf678a600> (a EDU.oswego.cs.dl.util.concurrent.LinkedNode)
```

```

    at EDU.oswego.cs.dl.util.concurrent.PooledExecutor.getTask(
PooledExecutor.java:707)
    at EDU.oswego.cs.dl.util.concurrent.PooledExecutor$Worker.run(
PooledExecutor.java:731)
    at java.lang.Thread.run(Thread.java:595)
"Low Memory Detector" daemon prio=10 tid=0x001e64e8 nid=0x19 runnable
[0x00000000..0x00000000]

"CompilerThread1" daemon prio=10 tid=0x001e5388 nid=0x18 waiting on
condition [0x00000000..0xb72fed2c]

"CompilerThread0" daemon prio=10 tid=0x001e4510 nid=0x17 waiting on
condition [0x00000000..0xb73febac]

"AdapterThread" daemon prio=10 tid=0x001e3698 nid=0x16 waiting on
condition [0x00000000..0x00000000]

"Signal Dispatcher" daemon prio=10 tid=0x001e2928 nid=0x15 waiting on
condition [0x00000000..0x00000000]

"Finalizer" daemon prio=10 tid=0x001d6078 nid=0x14 in Object.wait()
[0xfdf6f000..0xfdf6fa70]
    at java.lang.Object.wait(Native Method)
    at java.lang.ref.ReferenceQueue.remove(ReferenceQueue.java:116)
    - locked <0xc9e89940> (a java.lang.ref.ReferenceQueue$Lock)
    at java.lang.ref.ReferenceQueue.remove(ReferenceQueue.java:132)
    at java.lang.ref.Finalizer$FinalizerThread.run(Finalizer.java:159)

"Reference Handler" daemon prio=10 tid=0x001d5b20 nid=0x13 in
Object.wait() [0xfe04f000..0xfe04f8f0]
    at java.lang.Object.wait(Native Method)
    at java.lang.Object.wait(Object.java:474)
    at java.lang.ref.Reference$ReferenceHandler.run(Reference.java:116)
    - locked <0xc9e90860> (a java.lang.ref.Reference$Lock)

"VM Thread" prio=10 tid=0x001d3a40 nid=0x12 runnable

"GC task thread#0 (ParallelGC)" prio=10 tid=0x000d8608 nid=0x2 runnable

```

The output consists of a header and a stack trace for each thread. Each thread is separated by an empty line. The Java threads (threads that are capable of executing Java language code) are printed first. These are followed by information on VM internal threads.

The header line contains the following information about the thread:

- Thread Name indicates if the thread is a daemon thread
- Thread Priority (Prio)
- Thread ID (TID) is the address of a thread structure in memory
- ID of the native thread (NID)
- Thread State indicates what the thread was doing at the time of the thread dump
- Address range gives an estimate of the valid stack region for the thread

The following table lists the possible thread states that can be printed:

Thread State	Definition
NEW	The thread has not yet started.
RUNNABLE	The thread is executing in the JVM.
BLOCKED	The thread is blocked, waiting for a monitor lock.

Thread State	Definition
WAITING	The thread is waiting indefinitely for another thread to perform a particular action.
TIMED_WAITING	The thread is waiting for another thread to perform an action for up to a specified waiting time.
TERMINATED	The thread has exited.

Note: The thread header is followed by the thread stack.

Performing a Thread Dump

About this task

To perform a thread dump in UNIX (Solaris and HP-UX):

Procedure

1. From the Administration menu, select **Operations > System > Performance > JVM Monitor**.
2. In the JVM Monitor page, under **Thread Dump**, next to **Take Thread Dump**, click **Go!**. The Thread Dump pop-up window is displayed.
3. To see the latest Dumps in the View Dumps list, click **Go!** in the Thread Dump pop-up window.
4. Close the Thread Dump pop-up window.
5. In the JVM Monitor page, under **View Dumps**, select the Name of the thread dump file and click **Go!**. The Download Dumps pop-up window is displayed.
6. Click the **Dumps download** link.
7. Open or save the file.
8. Close the Download Dumps pop-up window.

What to do next

For Windows, if you have started Sterling Integrator as a Windows service, use the stacktrace tool to take a thread dump. Follow the same procedure to perform a thread dump for container JVMs.

For more information, refer to the following Web site: <http://www.adaptj.com/main/stacktrace>.

You can use VisualVM in remote mode to take a thread dump on Windows when you start Sterling Integrator noapp or container in service mode.

Refer to the corresponding VisualVM documentation for information about remote monitoring of JVM, which is available in the following Web site:

<https://visualvm.dev.java.net/>

HotSpot Thread Dump Analysis Tools

Following is a list of tools that can be used to analyze thread dumps:

- Thread Dump Analyzer (TDA): (Refer to <https://tda.dev.java.net/>)

Note: In order to use TDA, you must strip off any starting "<" symbols in the thread dump if you are using Sterling Integrator UI to take thread dumps.

- You can also use VisualVM for live monitoring of threads in JVM. (Refer to <https://visualvm.dev.java.net/>)

HotSpot JVM DeadLock Detection

In addition to the thread stacks, the **Ctrl+Break** handler executes a deadlock detection algorithm. If any deadlocks are detected, it prints additional information after the thread dump on each deadlocked thread.

```
Found one Java-level deadlock:
=====
"Thread2":
  waiting to lock monitor 0x000af330 (object 0xf819a938, a java.lang.String),
  which is held by "Thread1"
"Thread1":
  waiting to lock monitor 0x000af398 (object 0xf819a970, a java.lang.String),
  which is held by "Thread2"
Java stack information for the threads listed above:
=====
"Thread2":
  at Deadlock$DeadlockMakerThread.run(Deadlock.java:32)
  - waiting to lock <0xf819a938> (a java.lang.String)
  - locked <0xf819a970> (a java.lang.String)
"Thread1":
  at Deadlock$DeadlockMakerThread.run(Deadlock.java:32)
  - waiting to lock <0xf819a970> (a java.lang.String)
  - locked <0xf819a938> (a java.lang.String)
Found 1 deadlock.
```

If the Java VM flag `-XX:+PrintConcurrentLocks` is set, **Ctrl+Break** will also print the list of concurrent locks owned by each thread.

HotSpot JVM Blocked Thread Detection

Look for threads that are blocked. Threads might be waiting on SQL or might be serialized on a synchronized block. If you see threads that are blocked they are waiting on another thread to complete, it means that you are serializing on some part of code.

HotSpot JVM Troubleshooting Tips

`java.lang.OutOfMemory` errors occur when either the Java heap or the native heap run out of space. These exceptions may indicate that the number of live objects in the JVM require more memory than what is available (which can be adjusted by tuning the heap), or that there is a memory leak (which may indicate a problem with an application component, the JVM, or the OS). There are many variants of this error, with each variant having its own cause and resolution.

Generally, `java.lang.OutOfMemory` (OOM) indicates that insufficient space has been allocated for a requested object in the heap, or in a particular area of the heap, even after a garbage collection is attempted. An `OutOfMemory` exception does not always indicate a memory leak. It could indicate that the heap parameters are not configured properly for an application or that the bug is complex, with you having to troubleshoot different layers of native code.

When an OOM error is from the native code, it is difficult to tell whether it is from Java heap exhaustion or native heap exhaustion because of low space. The first step in diagnosing an OOM is identifying whether the Java heap is full or the native heap is full.

Exception in thread “main” java.lang.OutOfMemoryError

This error condition might be due to a simple configuration issue. It might be that the `-Xms` value configured for this type of application is throwing the error. This error generally occurs when object allocation fails even after a Garbage Collection (GC).

This error condition may be because of a memory leak, in which case, the object references are held by the application even if they are no longer needed. Over a period of time, this unintentional object growth increases and causes this OOM.

This type of OOM error can also be seen in applications in which excessive finalizer calls are made. In such applications, where the finalizer call is made, the GCs cannot reclaim the object space. Instead, they are queued for finalization to occur, which might happen some time later. In a Sun implementation, the finalizer call is made by a separate daemon thread. If there are excessive finalize calls, the finalization thread cannot keep up with the load, and eventually, the heap might become full and an OOM might occur. For example, applications that create high-priority threads that cause the finalization to increase at a rate that is faster than the rate the finalization thread can process, may generate this error.

Exception in thread “main” java.lang.OutOfMemoryError: PermGen space

This error is seen when permanent generation is full. Permanent generation is the area in which the class and method objects are stored. For an application that loads a large number of classes, the value of `-XX:MaxPermSize` should be sized accordingly. Permanent generation also gets used when the `java.lang.String intern()` method is invoked on an object of class `java.lang.String`. Exhaustion of the permanent generation area may occur if an application interns a large number of strings.

Generally, `java.lang.String` maintains the string pool, and when an interned method is called on a string, it first verifies that an equal string is already present in the pool. If an equal string exists, `java.lang.String` returns the canonical representation of the string, which points to the same class instance; otherwise, `java.lang.String` adds the string to the pool.

Exception in thread “main” java.lang.OutOfMemoryError: Requested array size exceeds VM limit

This error occurs when the application requests the JVM to allocate an array that is larger than the heap size. For example, if the application requests an allocation of 512 MB array and your heap size is only 256 MB, this error will be seen. This error may occur because of a low `-Xms` value, or a bug in the application in which it is trying to create a huge array.

Exception in thread “main” java.lang.OutOfMemoryError: request <size> bytes for <reason>. Out of swap space?

Although the error is an OOM condition, the underlying cause is that the JVM failed to allocate the requested size from the native heap, and the native heap is close to exhaustion.

The `<size>` in the error message is the size that failed to allocate. The `<reason>` is the name of the source module reporting failure. In a few cases, the actual reason is printed.

In order to troubleshoot this error, use the OS utilities to find the actual cause. One possible cause is that the OS has been configured with insufficient swap space. Another possible cause is that other processes in the machine are consuming all the memory, or possibly a native leak, in which the application or the library code is continuously allocating memory, and the OS is not releasing it.

Exception in thread "main" java.lang.OutOfMemoryError: <reason> <stack trace>(Native method):

If you see this OOM, it means that the native method encountered an allocation failure. The main difference between this error and the Exception in thread "main" java.lang.OutOfMemoryError: request `<size>` bytes for `<reason>`. Out of swap space? error is that the allocation failure in this case occurred in the JNI native method rather than the JVM code. In order to troubleshoot this error, use the OS utilities to find the actual cause.

A Crash, Instead of an OutOfMemoryError

In rare cases, you may have a JVM crash instead of OOM because of the allocation from the native heap failing because the native code that does not check for errors returns memory allocation functions. For example, this may occur if the native code malloc returns NULL (no memory available), and if the native code is not checking for that error, and references the invalid memory location. If the diagnosis of this failure results in native code not checking for errors because of memory allocation failures, the reasons for memory allocation failures should be examined. The failures may be due to reasons such as insufficient swap space, some other processes consuming all the memory, or a native leak.

Hung Processes or Looping Processes

Generally, the most common reasons for a hang to occur are deadlocks in application code, API code, library code, or a bug in the HotSpot VM. However, in a few cases, the hang might be because of a JVM consuming all the available CPU cycles, most likely because of a bug, which in turn causes one or more threads to go into an infinite loop.

If a hang is seen, determine whether the JVM is idle or consuming all the CPU cycles. You can use the OS utilities to determine the CPU utilization. If you conclude that the hang process is using the entire CPU, the hang might be because of a looping thread. On the other hand, if the process is idle, it is most likely because of deadlock. On Solaris, for example, `prstat -L -p <pid>` can be used to report the statistics for all the LWPs in the target process. This will identify the threads that are consuming a lot of CPU cycles.

Diagnosing a Looping Process

If your observations on CPU utilization indicate that the process is looping, take a thread dump, and from the thread dump and stack trace, you should be able to gather information about where and why the thread is looping.

In the thread dump, look for the runnable threads. You will in all probability find the threads that are looping. In order to be certain about the threads that are

looping, take multiple thread dumps to see if the thread remains busy.

Diagnosing a Hung Process

If you have determined that a process is hanging, and not looping, the cause is likely to be an application (thread) deadlock.

Take a thread dump and analyze it to find the deadlocked threads.

When you take a thread dump on a HotSpot JVM, the deadlock detection algorithm is also executed and the deadlock information in the thread dump printed.

Following is an example of some deadlock output from a document. For more information, refer to the *Java Troubleshooting and Diagnostic Guide*.

Found one Java-level deadlock:

```
=====
"AWT-EventQueue-0": waiting to lock monitor 0x000ffbf8 (object 0xf0c30560, a
java.awt.Component$AWTTreeLock),
    which is held by "main"
"main":
    waiting to lock monitor 0x000ffe38 (object 0xf0c41ec8, a java.util.Vector),
    which is held by "AWT-EventQueue-0"
```

Java stack information for the threads listed above:

```
=====
"AWT-EventQueue-0":
at java.awt.Container.removeNotify(Container.java:2503)
- waiting to lock <0xf0c30560> (a java.awt.Component$AWTTreeLock)
at java.awt.Window$1DisposeAction.run(Window.java:604)
at java.awt.Window.doDispose(Window.java:617)
at java.awt.Dialog.doDispose(Dialog.java:625)
at java.awt.Window.dispose(Window.java:574)
at java.awt.Window.disposeImpl(Window.java:584)
at java.awt.Window$1DisposeAction.run(Window.java:598)
- locked <0xf0c41ec8> (a java.util.Vector)
at java.awt.Window.doDispose(Window.java:617)
at java.awt.Window.dispose(Window.java:574)
at
javafx.swing.SwingUtilities$SharedOwnerFrame.dispose(SwingUtilities.java:1743)
at
javafx.swing.SwingUtilities$SharedOwnerFrame.windowClosed(SwingUtilities.java:172)
at java.awt.Window.processWindowEvent(Window.java:1173)
at javafx.swing.JDialog.processWindowEvent(JDialog.java:407)
at java.awt.Window.processEvent(Window.java:1128)
102
at java.awt.Component.dispatchEventImpl(Component.java:3922)
at java.awt.Container.dispatchEventImpl(Container.java:2009)
at java.awt.Window.dispatchEventImpl(Window.java:1746)
at java.awt.Component.dispatchEvent(Component.java:3770)
at java.awt.EventQueue.dispatchEvent(EventQueue.java:463)
at
java.awt.EventDispatchThread.pumpOneEventForHierarchy(EventDispatchThread.java:214)
at
java.awt.EventDispatchThread.pumpEventsForHierarchy(EventDispatchThread.java:163)
)
at java.awt.EventDispatchThread.pumpEvents(EventDispatchThread.java:157)
at java.awt.EventDispatchThread.pumpEvents(EventDispatchThread.java:149)
at java.awt.EventDispatchThread.run(EventDispatchThread.java:110)
"main":
at java.awt.Window.getOwnedWindows(Window.java:844)
- waiting to lock <0xf0c41ec8> (a java.util.Vector)
at
```



```

javax.swing.SwingUtilities$SharedOwnerFrame.installListeners(SwingUtilities.java
:1697)
at
javax.swing.SwingUtilities$SharedOwnerFrame.addNotify(SwingUtilities.java:1690)
at java.awt.Dialog.addNotify(Dialog.java:370)
- locked <0xf0c30560> (a java.awt.Component$AWTTreeLock)
at java.awt.Dialog conditionalShow(Dialog.java:441)
- locked <0xf0c30560> (a java.awt.Component$AWTTreeLock)
at java.awt.Dialog.show(Dialog.java:499)
at java.awt.Component.show(Component.java:1287)
at java.awt.Component.setVisible(Component.java:1242)
at test01.main(test01.java:10)
Found 1 deadlock.

```

Note: In J2SE 6.0, the deadlock detection algorithm works only with the locks that are obtained using the synchronized keyword. This means that deadlocks that arise through the use of the java.util.concurrent package are not detected.

In the deadlock output, you can see that the thread main is the locking object <0xf0c30560>, and is waiting to enter <0xf0c41ec8>, which is locked by the thread “AWT-EventQueue-0”. However, the thread “AWT-EventQueue-0” is also waiting to enter <0xf0c30560>, which is in turn locked by “main”.

Stack Overflow

This error generally occurs when the stack space is exhausted in a JVM. Generally, this occurs because of:

- A deeply nested application
- An infinite loop within an application
- A problem in the Just-In-Time (JIT) compiled code

Not all instances of this error should be considered as programming errors. In the context of some applications, you may require a greater value for the stack size (-Xss), for example, applications having intensive graphics might require more stack size.

The stack overflow error can be either from the native code or because of an infinite loop in the Java program.

To determine if the error is in the native code, review the stack trace. In most cases, the information you get will be difficult to interpret. However, if the error is due to an infinite loop, you can see the stack trace of the error. Verify whether there are any recursive method calls, and whether they are deep. If it does not appear to be an infinite loop, try increasing either the Java stack or the native stack to resolve the issue.

Taking Heap Dumps and Profiling JVM Using Hprof

A JVM's HPROF can be used get information about CPU usage and heap allocation statistics, and to monitor contention profiles. You can also get complete heap dumps and the states of all the monitors and threads in the JVM.

In Sterling Integrator, HPROF can be invoked by adding the following option to SERVER_FLAGS in tmp.sh.in file for ASI JVM and SERVER_ACFLAGS in tmp.sh_platform_ifcresources_ext.in file for container JVMs. Run **setupfiles.sh** command and restart ASI and container JVMs.

-agentlib:hprof[=options]

Or

-Xrunhprof[:options]

On Windows, add “echo -agentlib\hprof[=options] >> %PARAM_FILE% “ in InstallNoappsWindowsService.cmd.in for ASI JVM and InstallContainerWindowsService.cmd.in for container JVMs.

Run **setupfiles.cmd** command and the following scripts in the same order to apply the changes:

- **stopContainerWindowsService.cmd** <Container_Number>
- **uninstallContainerWindowsService.cmd** <Container_Number>
- **InstallContainerWindowsService.cmd** <Container_Number>
- **InstallACPassPhraseWindowsService.cmd** <Container_Number>
- **startContainerWindowsService.cmd** <Container_Number>

Follow the same procedure to apply the changes for all containers.

For more information about how to use HPROF on Sun JVM, refer to the following Web site:

<http://java.sun.com/developer/technicalArticles/Programming/HPROF.html>

For more information about how to use HPROF HP JVM, refer to the following Web site:

http://docs.hp.com/en/5992-5899/ch05s03.html#using_xrunhprof

On HP JVM, you can also use -Xeprof to collect profiling data for performance tuning. For information about -Xeprof, refer to the following Web site:

http://docs.hp.com/en/5992-5899/ch05s03.html#using_xeprof

In order to analyze the data using -agentlib:hprof and -Xeprof, you can use HPJmeter. For more information about this tool, refer to the following Web site:

<http://docs.hp.com/en/5992-4687/ch01s07.html>

To enable hprof options in Sterling Integrator, on Solaris and HP-UX, edit tmp.sh.in and include the hprof flags described previously under the respective OS section.

Introduction to the IBM[®] JVM Performance and Tuning Guidelines

The IBM[®] Java[™] Virtual Machine (JVM) contains a number of private and proprietary technologies that distinguish it from other implementations of JVM. For instance, the IBM JVM uses mixed-mode interpretation (MMI). When the MMI detects that bytecodes have been interpreted multiple times, it invokes a just-in-time (JIT) compiler to compile those bytecodes to native instructions. Due to the significant performance benefits, the JIT and MMI are enabled by default. The JVM performance degrades considerably when JIT is disabled. For more information about the JIT compiler and MMI, refer to the *IBM JDK 6.0: Java Diagnostics Guide*.

Refer to the *Sterling Integrator System Requirements* documentation for the supported operating system and JVM combination.

This topic describes the processes involved in tuning, monitoring, and performing basic troubleshooting when deploying the Sterling Integrator using the IBM JVM.

Before You Begin Tuning Your IBM® JVM

Using the Performance Tuning Utility, Sterling Integrator calculates the recommended settings based on the number of cores and the amount of physical memory that is being made available to Sterling Integrator.

The resulting performance properties are stored in the `tuning.properties` file in the `install_dir/properties` directory. The formulae used to calculate these setting can be found in the `tuningFormulas.properties` file in the `install_dir/properties` directory.

The calculated values should be used as a guideline. Further tuning may be necessary to attain a well-tuned system. If you are still unable to attain a well-tuned system, you may want to engage Sterling Commerce Professional Services. Contact your Sales Representative for more information about this.

IBM® JVM Default Parameters for Sterling Integrator

The follow table provides the IBM® JVM parameters and the default values you should use when configuring the system:

Parameter	Description
-Xms	<ul style="list-style-type: none">Controls the initial size of the Java heapCorresponds to the JVM long-lived memory (min) value <p>For 32-bit noapp JVM:</p> <ul style="list-style-type: none">Default Value for Linux: 512mDefault Value for AIX: 512m <p>For 64-bit noapp JVM:</p> <ul style="list-style-type: none">Default Value for Linux: 1280mDefault Value for AIX: 1280m <p>For 64-bit container JVM:</p> <ul style="list-style-type: none">Default Value for Linux: 128mDefault Value for AIX: 128m
-Xmx	<ul style="list-style-type: none">Controls the maximum size of the Java heapCorresponds to the JVM long-lived memory (max) value <p>For 32-bit noapp JVM:</p> <ul style="list-style-type: none">Default Value for Linux: 1024mDefault Value for AIX: 1024m <p>For 64-bit noapp JVM:</p> <ul style="list-style-type: none">Default Value for Linux: 1280mDefault Value for AIX: 1280m <p>For 64-bit container JVM:</p> <ul style="list-style-type: none">Default Value for Linux: 512mDefault Value for AIX: 512m

Parameter	Description
-Xmns	<ul style="list-style-type: none"> Controls the initial size of the new area to the specified value when using -Xgcpolicy:gencon. Corresponds to the JVM short-lived memory (min) value <p>For 32-bit noapp JVM:</p> <ul style="list-style-type: none"> Default Value for Linux: 64m Default Value for AIX: 64m <p>For 64-bit noapp JVM:</p> <ul style="list-style-type: none"> Default Value for Linux: 256m Default Value for AIX: 256m <p>For 64-bit container JVM:</p> <ul style="list-style-type: none"> Default Value for Linux: 32m Default Value for AIX: 32m <p>It is recommended to use the following formula to compute -Xmns value when the minimum heap size is modified.</p> $-Xmns = (0.33333 * \text{value of } -Xms)$
-Xmnx	<ul style="list-style-type: none"> Controls the maximum size of the new area to the specified value when using -Xgcpolicy:gencon Corresponds to the JVM short-lived memory (max) value <p>For 32-bit noapp JVM:</p> <ul style="list-style-type: none"> Default Value for Linux: 128m Default Value for AIX: 128m <p>For 64-bit noapp JVM:</p> <ul style="list-style-type: none"> Default Value for Linux: 512m Default Value for AIX: 512m <p>For 64-bit container JVM:</p> <ul style="list-style-type: none"> Default Value for Linux: 128m Default Value for AIX: 128m <p>It is recommended to use the following formula to compute -Xmnx value when the maximum heap size is modified.</p> $-Xmnx = (0.33333 * \text{value of } -Xmx)$
-Xjit:count	<ul style="list-style-type: none"> Controls the compilation threshold of the JIT compiler Value given to count causes the Java method to be compiled after n runs <p>For noapp (32-bit/64-bit) JVM and container (64-bit) JVM:</p> <ul style="list-style-type: none"> Default Value for Linux: 1000 Default Value For AIX: 1000
-Xgcpolicy:gencon	<p>Controls the behavior of the Garbage Collector. They make trade-offs between the throughput of the application and the overall system, and the pause times that are caused by garbage collection. By specifying the gencon value, the GC policy requests the combined use of concurrent and generational GC to help minimize the time that is spent in any garbage collection pause.</p>

Parameter	Description
-Xss	<p>Controls the maximum Java stack size for any thread</p> <p>For 32-bit noapp JVM:</p> <ul style="list-style-type: none"> • Default Value for Linux: 256k <p>For 64-bit noapp JVM:</p> <ul style="list-style-type: none"> • Default Value for Linux: 512k <p>For 64-bit container JVM:</p> <ul style="list-style-type: none"> • Default Value for Linux: 256k • Default Value for AIX: 256k

IBM® JVM Troubleshooting Tips

This topic describes the various issues that may arise when using the IBM® JVM and the troubleshooting tip pertaining to each of these errors.

OutOfMemoryError Exceptions and Memory Leaks

OutOfMemoryError exceptions occur when either the Java heap or the native heap run out of space. These exceptions indicate that there is either a memory leak or that the number of live objects in the JVM require more memory than is available.

The first step to troubleshooting an OutOfMemoryError exception is to determine whether the error is caused because of lack of space in either the Java heap or the native heap. When the OutOfMemoryError is caused because of lack of space in the native heap, an error message is displayed with an explanation about the allocation failure. If an error is not present, the exception is likely to have occurred because of lack of space in the Java heap. In the latter scenario monitor the Java heap by using the Verbose GC output.

The Java heap is consumed when the Garbage Collector is unable to compact or free the objects being referenced. Objects that are no longer referenced by other objects, or are referenced from the thread stacks, can be freed by performing the garbage collection task. As the number of referenced objects increases, garbage collection will take place more frequently. With each garbage collection instance, less memory will be freed. If this trend continues, the Garbage Collector will not be able to free enough objects to allocate new objects. When this happens, the heap will be near 100% utilized and the JVM will fail with an OutOfMemoryError exception.

Increasing the size of the Java heap may resolve an OutOfMemoryError exception. However, if the exception is due to a memory leak, increasing the heap size will not resolve this issue. In this case, further troubleshooting, including analyzing the contents of the heap using heap dump analysis or a tool such as JProbe® or OptimizeIt™ is necessary. Refer to the *IBM 6.0: Java Diagnostics Guide* for more information about heap dump analysis.

The Application Hangs

The application hangs when either a deadlock occurs or a loop is encountered. A potential deadlock scenario is one in which multiple threads in the JVM are blocked and are waiting on the same object resource. Another situation that may result in a deadlock is when there is a missed notification between threads because

of a timing error. Similarly, a loop may be encountered if there is a missed flag, which in turn may terminate the loop. A loop can also be encountered if the wrong limit has been set for the loop iterator. In either of these cases, the thread will fail to exit in a timely manner.

The recommended approach is to trigger a thread dump and interpret it. The thread dump provides all the information pertaining to the object resources in the JVM. A thread dump can be triggered to take a snapshot of all the information related to the JVM and a Java application at a particular point during execution. The information that is captured includes the OS level, hardware architecture, threads, stacks, locks, monitors, and memory.

By default, thread dumps are enabled. A thread dump is triggered when the JVM is terminated unexpectedly, an `OutOfMemoryError` is encountered, or when a user sends specific signals to the JVM. Sterling Integrator also enables users to trigger a thread dump through the Sterling Integrator UI.

Trigger Thread Dump

To trigger a thread dump for the ASI JVM:

1. From the Administration menu, select **Operations > System > Performance > JVM monitor**.
2. In the JVM Monitor page, under **Thread Dump**, next to **Take Thread Dump**, click **Go!**.

The **Thread Dump** pop-up window is displayed. It may take some time to generate the thread dump. During this time you may not be able to access other pages. If heap dumps are enabled, the delay will be longer before the process is completed.

3. Click **Go!** in the **Thread Dump** pop-up window to have the thread dumps populated on the JVM Monitor page.
4. Close the **Thread Dump Taken** pop-up window.
5. To view the dump, in the JVM Monitor page, under **View Dumps**, select the dump file, and click **Go!**.
6. In the **Download Dumps** pop-up window, click the **Download** link to download the file to the local machine.

The thread dump can also be found on the host machine in the `install_dir/noapp/bin` directory. After the download is complete, close the **Download Dumps** pop-up window.

To trigger a thread dump for the container JVM:

1. Change your working directory to `install_dir`.
2. In the command line, enter `ps -ef | grep <container_name>`.

This lists the container Java process id.

3. Enter `kill -3 <pid>`.

The thread dump is placed in the `install_dir/noapp/bin` directory for your analysis.

Because thread dumps are created in a text format, they do not require any software to make them human readable. A thread dump is broken into sections. Following is a brief description of each thread dump tag:

Thread Dump Tag	Description
TITLE	Basic information about the event that caused the thread dump, along with the timestamp and the generated name.
GPINFO	Contains general information about the operating system. General Protection Fault (GPF) information is included in this section if the failure was caused by a GPF.
ENVINFO	Contains the JRE level and details about the environment and command-line arguments that were used to launch the JVM process.
MEMINFO	Contains information about the Memory Manager, and free space, current size of the heap and the garbage collection history data.
LOCKS	Contains information about the locks and monitors being held by each Java object.
THREADS	Contains a complete list of threads that have not been stopped.
CLASSES	Contains the class loader summaries, including the class loaders and their relationships.

For more information about interpreting thread dump, refer to the *IBM JDK 6.0: Java Diagnostics Guide*.

Heapdumps

Heapdumps are useful for troubleshooting memory-related issues since they contain all the live objects used by the Java application. With this information, the objects that are using large amounts of memory can be identified. The contents of a heapdump can also help a user understand why objects cannot be freed by the Garbage Collector.

By default, a heapdump is generated in a compressed binary format know as Portable Heap Dump (PHD). Several tools are available to help analyze the dump. IBM recommends the use of the Memory Dump Diagnostic for Java (MDD4J) when performing the heapdump analysis. This tool can be downloaded from IBM Support Assistant (<http://www-01.ibm.com/software/support/isa/>).

Heapdumps are generated when the Java heap is exhausted by default. They can also be configured so that they are generated when a user sends a specific signal to the JVM.

For more information about heapdumps, refer to the *IBM JDK 6.0: Java Diagnostics Guide*.

Enable a Heapdump

To enable a heapdump to respond to user signals through the Sterling Integrator UI for noapp JVM:

1. From the Administration menu, select **Operations > System > Performance > JVM monitor**.
2. In the JVM Monitor page, select the checkbox next to **Enable Heapdump**.
3. When prompted with the steps to complete the process, click **OK**.
4. In a terminal window, change to the *install_dir/bin* directory and execute the *setupfiles.sh* script.
5. Restart the Sterling Integrator instance.

CAUTION:

It is recommended to contact Sterling Commerce Customer Support before modifying the files manually.

To enable a heapdump for container JVM:

1. Navigate to *install_dir*/bin directory.
2. Edit *tmp.sh_platform_ifcresources_ext.in* file to add the following environment variables to the respective operating system. Add the variables between **SERVER_ACFLAGS** and **;;** strings:

```
IBM_HEAP_DUMP=true
export IBM_HEAP_DUMP
IBM_HEAPDUMP=true
export IBM_HEAPDUMP
IBM_HEAPDUMP_OUTOFMEMORY=true
export IBM_HEAPDUMP_OUTOFMEMORY
```

3. Run **./setupfiles.sh** command.
4. Restart containers.
5. In the command line, enter **ps -ef | grep <container_name>**.

This lists the container pid.

6. Enter **kill -3 <pid>**.

The heapdump is placed in the *install_dir*/noapp/bin directory for your analysis.

The heapdumps will also be generated if the JVM encounters an OutOfMemory error.

Chapter 5. Monitoring Operations

Monitoring Operations

The Operations functions enable you to monitor the operational status of Sterling Integrator, its components and services, current threads and system logging, and to troubleshoot system problems.

Managing System Logs

Sterling Integrator comprises multiple components, including software applications such as archive tools, Web servers, and database servers. To monitor the activities of each component, Sterling Integrator generates log files based on the system's monitoring activity. These log files are one of the tools that enable you to monitor the way the system operates.

Each operations server on a host has its own operations log file. Log files are created in the logs directory of the installation directory. To prevent the system log files from taking up excessive storage space and main memory, Sterling Integrator generates a log file only when a component runs. This in turn improves the performance of Sterling Integrator.

Each open log file is closed once every 24 hours (usually at midnight), and a new file is created. When moving into or away from Daylight Savings Time, you must stop Sterling Integrator and restart it in order to ensure that files are created with the correct timestamp.

Log files are allowed to grow only up to a maximum size. If a log file reaches its maximum size, it is closed and another file is created. There is a limit to the number of log files that can exist simultaneously. If this limit is exceeded, the old logs files are automatically deleted. If the Sterling Integrator Dashboard interface links to a deleted log file (which will display a blank page), click the link to a newer log file.

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

Naming Conventions

This section provides information about the naming conventions to be used in Sterling Integrator.

Use the following naming convention for a directory:

- UNIX - *install_dir*/logs/directory
- Windows - *install-dir*\logs\directory

In this convention:

- *install_dir* or *install-dir* refers to the name of the installation directory.
- *logs* refers to the primary log directory.
- *directory* refers to the subdirectory created when you start Sterling Integrator.

All the old log files are moved to this subdirectory for archiving. The naming convention to be used for the old log directory is *logs_mmddyy_hhmmss*.

Use the following naming convention for a log file:

name.Log.Dyyyymmdd.Thhmmss_#

In this convention:

- *name* identifies the type of the log file.
- *Log* refers to the file name extension, which indicates the file type.
- *Dyyyymmdd* refers to the date in the year, month, and day format. The *D* at the beginning refers to Date.
- *Thhmmss* refers to the time in hours, minutes, and seconds format. The *T* at the beginning refers to Time.
- *_#* is the increment of the log file. If you attempt to write a log file that already exists, *_#* is appended to the log file name, allowing you to write a new file and save the integrity of the existing file.

For example, if *mylog.D20041101.T092022* exists, and you try to save a new log file with the same name, the new file becomes *mylog.D20041101.T092022_2*, where *_2* indicates that it is the second log in a sequence using the same file name.

Note: The date and time components in a naming convention may or may not be present, depending on the type of the log. For example, the *noapp.log* does not include date and time information, but *ui.log.Dyyyymmdd.Thhmmss* includes date and time information.

Each time Sterling Integrator is started, the log files created since the last time it was started are archived in a time-stamped subdirectory. New log files, those created subsequent to the most recent execution of Sterling Integrator, are written to the *install_dir/logs* directory.

Viewing Log File Contents

Users of Sterling Integrator can view the contents of both current log files and old log files.

Note:

- If you are working in a clustered environment, the log information that is displayed is determined by the node you select from the **Select Node** list.
- The Sterling Integrator 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 on which Sterling Integrator is installed. Open the log file in read-only mode using a text editor.

To view the current log file contents in Sterling Integrator:

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

To view the old log file contents in Sterling Integrator:

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

Analyzing Log File Contents

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

```
[2008-04-22 09:02:43.404]
ERROR 000310160001 UTIL.FRAME_POOL.ERR_Pool [Pool]
Could not create the initial objects for the pool gentranTPPool
[2008-04-22 09:02:43.405]
ERROR 000000000000 GLOBAL_SCOPE [1208869363405]
The driver manager could not obtain a database connection.
```

In this convention:

- YYYY-MM-DD refers to the date in year, month, day format.
- HH:MM:SS.ss refers to the time in hour, minutes, seconds, and hundredths of a second format.
- loglevel indicates how much information is logged and the relative importance of the message. A subsystem may log only a subset of these messages (as defined in the *subsystem.loglevel* property of *log.properties*), discarding those that have a severity that is lower (less severe) than the current log level set for that subsystem.

Log Level	Description
FATAL	Collects fatal and critical error information.
ERRORDTL	Collects only error conditions, with a detailed description of the error.
ERROR	Collects only error conditions such as exceptions and error messages (including errors from the user interface).
WARN	Collects non-fatal configuration messages.
SQLDEBUG	Collects SQL statements that are being executed.
INFO	Collects basic operational information.
TIMER	Collects timing information.
COMMTRACE	Collects communication trace information.
DEBUG	Collects basic debugging statements including system state and code paths.
VERBOSE	Collects extra debugging statements (like XML information) that describe and explain what is happening in the system.
ALL	Collects information about all the 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 or exception, 2 for debug messages, 3 for warnings, and 4 for information or 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 Sterling Integrator (such as WORKFLOW, OPS, or UTIL), the Sterling Integrator subsystem (such as FRAME_POOL, NOAPP, or SERVER), and what occurred (such as ERR_Pool or INFO_NamingException1).
- information string is a brief description of the activity that occurred.

Following 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
```

This indicates that at 11:06:55.661 a.m. on May 30, 2006, Sterling Integrator was started, and attempted to load the HTTP Send adapter (HTTP_SEND_ADAPTER). The message also provides information about:

- Scope (Services)
- The affected part of Sterling Integrator (Services Controller)
- What occurred (INFO_sdi_getName)
- Error code (0297)

Changing Log Settings

Sterling Integrator enables you to change the log settings globally and locally.

Changing Log Settings Globally

You can change log settings globally using the `customer_overrides.properties` file, which prevents customized property file changes from being overridden by updates or patches. You can change global settings in the `log.properties` file. For more information about the `customer_overrides.properties` file, refer to *Sterling Integrator Property Files* documentation.

Note: The customer override property file is not a part of the initial Sterling Integrator installation. It must be created and named `customer_overrides.properties`.

To change the property file settings using the `customer_overrides.properties` file, perform the following tasks:

1. In the `install_dir/properties` directory, either create or locate the `customer_overrides.properties` file.
2. Open the `customer_overrides.properties` file using a text editor.
3. Specify the settings for the global log properties described in the following table. These properties are displayed in the following format:
`logService.Property=Value`
 - `logService` identifies the `log.properties` file in the `customer_overrides.properties` file.
 - `Property` is the global property of the `log.properties` file that you want to set. See the following table for a list of properties.
 - `Value` is the property setting of the `log.properties` file.

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 Example: <code>logService.newloggers=true</code>
defaultlog	Specifies the name of the default log. Default is <code>systemlogger</code> . Example: <code>logService.defaultlog=logtype</code>

Property	Description
<code>logtype.maxnumlogs</code>	<p>Specifies the maximum number of logs to retain before deleting the old logs. Default is 10.</p> <p>Examples:</p> <p><code>logService.defaultlog.maxnumlogs=15</code> sets the maximum number of a log type specified as the default log (systemlogger, by default) to 15.</p> <p><code>logService.uilogger.maxnumlogs=20</code> sets the maximum number of UI type logs to 20.</p>

The following table provides the log type 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 to find the log type name.

File Name of Log (As Shown on the System Logs Screen)	Log Type (For the <code>customer_overrides.properties</code> file)
<code>alerterlogger.log</code>	<code>alerterlogger</code>
<code>archive.log</code>	<code>archivelogger</code>
<code>Authentication.log</code>	<code>AuthenticationLogger</code>
<code>cdinterop.log</code>	<code>cdinteroplogger</code>
<code>cdinterop_cdjava.log</code>	<code>cdinteropcjavalogger</code>
<code>ceuinterop.log</code>	<code>ceulogger</code>
<code>common3splogger.log</code>	<code>common3splogger</code>
<code>delete.log</code>	<code>deletelogger</code>
<code>ebXML.log</code>	<code>ebXMLlogger</code>
<code>EDIINT.log</code>	<code>EDIINTLogger</code>
<code>event.log</code>	<code>event</code>
<code>ftp.log</code>	<code>ftplogger</code>
<code>ftpclient.log</code>	<code>psftpclientlogger</code>
<code>http.log</code>	<code>httplogger</code>
<code>httpclient.log</code>	<code>httpclientlogger</code>
<code>jetty.log</code>	<code>jettylogger</code>
<code>lifecycle.log</code>	<code>lifecycleLogger</code>
<code>mailbox.log</code>	<code>mailboxlogger</code>
<code>mgmtdash.log</code>	<code>neo</code>
<code>noapp.log</code>	<code>noapplogger</code>
<code>oftp.log</code>	<code>oftplogger</code>
<code>ops_exe.log</code>	<code>opslogger</code>
<code>Perimeter.log</code>	<code>PSLogger</code>
<code>pipeline.log</code>	<code>pipelinelogger</code>
<code>report.log</code>	<code>reportlogger</code>
<code>resourcemonitor.log</code>	<code>resourcemonitorlogger</code>
<code>mif.log</code>	<code>rnlogger</code>
<code>sap.log</code>	<code>saplogger</code>
<code>schedule.log</code>	<code>schedulelogger</code>

File Name of Log (As Shown on the System Logs Screen)	Log Type (For the customer_overrides.properties file)
Security.log	SecurityLogger
servicesctl.log	sclogger
sftpclient.log	sftpclientlogger
sftpserver.log	sftpserverlogger
si_exe.log	silogger
sql.log	sqllogger
system.log	systemlogger
system.log	purgelogger
test.log	testlogger
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 the Sterling Integrator and restart it to use the new values.

Changing Log Settings for an Individual Log Type

For each log, you can specify the following information:

- Location of the log file
- The number of lines pertaining to the log file to be saved
- The amount of details to log

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

To change the log settings for an individual log type:

1. From the Administration menu, select **Operations > System > Logs**.
2. Click the icon next to the log type whose log settings you want to change.
3. In the Log Settings page, specify the settings for the options described 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, the rollover interval 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), the rollover interval is the maximum number of lines allowed in the log file before a new file is created. Select one of the following rollover interval values:</p> <ul style="list-style-type: none"> • 50000 lines • 100000 lines • 150000 lines • 200000 lines <p>Note: Here, the term “lines” refers to logical lines (entries) and not physical lines. For example, 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 details to log. Select 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 may generate an excessive amount of debugging information. You should lower the logging level after you have retrieved the debugging information you require.</p>

4. Click **Save**.

Changing Log File Location

Sterling Integrator enables you to modify the location of the log files. Perform the following tasks:

1. Modify the LOG_DIR setting in `sandbox.cfg`.
2. Run the `setupfiles.sh` script to apply the changes.

However, the following log files created by the Sterling Integrator dashboard interface are not moved to the changed location indicated in `sandbox.cfg`:

- jetspeed.log
- jetspeedservices.log
- torque.log
- turbine.log
- access.log

To modify the location of these log files, run the `deployer.sh` script. This script rebuilds and redeploys the Web ARchive (WAR) files of the dashboard interface.

The locations of the log files created by the service configuration are hard coded in the configuration. As a result, the location of the `einvoicing.log` files cannot be changed by modifying the setting in `sandbox.cfg`. To modify the location of these log files, you must modify the settings at the service configuration level.

Log File Types

The log files in Sterling Integrator can be classified under various types. The following table lists these types along with a description of the same.

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

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 the log messages the operations servers generate during startup.
Gentran:Server for UNIX		
Data Adapter		Logs activities of the Gentran:Server for UNIX adapter. You cannot turn logging on or off for Gentran:Server for UNIX data adapter. Note: The Gentran:Server for UNIX logs are displayed only if you have Sterling Integrator configured for Gentran:Server for UNIX.
Application Logs		
Adapter Server	servicesctl.log	Used by the service controller component.
AFT Routing	aftrouting.log	

Log Type	Log Name	Description
Alerter	alerterlogger.log	Logs notification failures and Alert own errors in the Alert Service. When debug is turned on, the alerter log type also logs all the alert information, such as defined alerter and filter 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 that is specific to running a business process.
Business Process Policy Statistics	wfstatistics.log	Contains workflow policy statistics generated by the workflow scheduling policy. Although the actual content depends on the scheduling policy in place, the business process policy statistics log type basically contains XML timestamps followed by XML records.
Connect:Direct Secure Perimeter Adapter	cdsp.log	
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 and related services
Connect:Enterprise Server Adapter and Services	ceuinterop.log	Used by the Connect:Enterprise Server adapter.
Crypto	crypto.log	
CSP2 FTP Adapter	cspftp.log	
CSP2 Http Adapter	csphttp.log	
Dashboard and Community Management	mgmtdash.log	Used by the Dashboard and Community Manager components.
Delete Resources	delete.log	Logs information about resources that have been deleted from Sterling Integrator.
Document Tracking	tracking.log	Logs document tracking activities.
ebXML Business Process Execution	ebXML.log	Logs ebXML business process execution activities.
EDI Log	edi.log	

Log Type	Log Name	Description
EDIINT AS1 and AS2	EDIINT.log	Used by the EDIINT components.
Embedded Engine	embeddedEngine.log	
Event Framework	event.log	Logs event framework activities for events completed in Sterling Integrator.
FTP Client Adapter and Services	ftpclient.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 Sterling Integrator interface.
Jetty HTTP Server	jetty.log	Used by the Jetty HTTP Server.
JMX Agent	jmx.log	
Log.ResourceMonitorLog	resourcemonitor.log	Used by the Resource Monitor.
Mailboxing Subsystem	mailbox.log	Used by the mailbox components in Sterling Integrator.
Odette FTP Adapter Administration	OdetteFTP.log	
OFTP Administration	oftp.log	Logs OFTP administration activities.
Perimeter Services	Perimeter.log	Used by the perimeter server components in Sterling Integrator.
Pipeline	pipeline.log	Used by the pipeline components.
Platform	platform.log	
Report	report.log	Used by the reporting components.
Reporting Services	bizIntel.log	
Resource Monitor	resourcemonitor.log	

Log Type	Log Name	Description
RosettaNet Business Process Execution	rnif.log	Used by the RosettaNet™ components.
SAP Adapter Administration	sap.log	Used by the SAP® components.
SAP XI Adapter Administration	sapxi.log	
Schedule	schedule.log	Logs scheduling activities.
Schedule Monitor	schedulemonitor.log	
Secure Proxy	secureproxy.log	
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.
SWIFTNet	swiftnet.log	
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 Log	tx.log	
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 matched each block of data against the map.
User Authentication	Authentication.log	Logs user authentication attempts and activities.
Visibility	visibility.log	
Web Extension	webx.log	Used by the Web Extensions components in Sterling Integrator.
Web Services Security	wssec.log	
WebDAV Server	webdav.log	

Log Type	Log Name	Description
WebSphereMQ Suite	WebSphereMQSuite.log	Used by the WebSphereMQ Suite Async Receiver adapter and related services.
Windows Service GI	si_exe.log	Log file created by the Sterling Integrator Windows service.
Windows Service Ops log	ops_exe.log	Log file created by the Opserver Windows service.
WorkFlow Deadline	bpdeadline.log	
WS-Reliability Routing	wstrm.log	

The following table describes the log files pertaining to the Sterling Integrator Windows service:

Log Name	Description
ScheduleBackup.log	Temporary file that is created when Sterling Integrator Windows service stops.
Backuplogs.log	Temporary file that is created when Sterling Integrator Windows service stops.
ScheduleStopOps.log	Temporary log file that can be ignored.

The following table describes the log files pertaining to the DMZ perimeter server:

Log Name	Description
PSLogger. <i>Dyyyymmdd.Thhmmss</i>	Logs perimeter server information for the DMZ perimeter server.
StartupPS.log	Logs startup activities for the DMZ perimeter server.

Auditing

In Sterling Integrator, you can find information about the creation, modification, and deletion of a Sterling Integrator resource, using the AUDIT_ADMIN table. Resources include business processes, certificates (CA, trusted, system), maps, and schemas. You can access the AUDIT_ADMIN table through a simple database query.

The AUDIT_ADMIN table contains the following information:

- The date and time of a resource operation.
- The resource that was created, modified, or deleted.
- The resource operation (creation, modification, or deletion).
- The User ID of the user who performed the operation.

The AUDIT_ADMIN table also contains information about when a user obtains or releases a lock on a resource using the Lock Manager, if the type of modification cannot be determined in the Lock Manager.

You can generate reports (by Resource Type or User ID) from the AUDIT_ADMIN table. Use the following procedure to generate an Admin-Audit Report:

1. In the **Administration** Menu, select **Operations > Reports**.

2. In the **Search** section of the screen that is displayed, select the **Admin Audit** type and click **Go!**.
3. Click **source manager** icon adjacent to the report you want to generate:
 - AdminAuditByObjectType (by Resource Type)
 - AdminAuditByPrincipal (by User ID)
4. In the Report Source Manager page that is displayed, click the **execute** icon. The Admin-Audit Report is displayed.

The Admin-Audit Report (whether by Report Type or User ID), includes the following columns. The AUDIT_ADMIN table field name is displayed within parentheses after the column name.

- Action Type (ACTION_TYPE)
Example: Modified
- Action Value (ACTION_VALUE)
Example: Message Purge
- Principal (PRINCIPAL)
Example: UserID
- Resource Name (OBJECT_NAME)
Example: Message Maintenance
- Resource Type (OBJECT_TYPE)
Example: User News
- Time (TIME)
Example: 07/15/2008 12:48:54 PM

The AUDIT_ADMIN table also includes the ARCHIVE_DATE field, which is the earliest date on which the data can be purged.

The audit process tracks the following resources:

- Accounts
- Application Configurations
- Business Processes
- Digital Certificates
- Communities
- ebXML Specifications
- Extended Rule Libraries
- Maps
- Mail Boxes
- PGP Profiles
- Proxy policies
- Proxy Setmaps
- Perimeter Servers
- Report Configurations
- Schedules
- XML Schemas
- Security Tokens
- Service Configurations
- SSH Resources

- SWIFTNet Routing Rules
- Trading Partner Data
- Web Resources
- Web Services
- WSDL
- Web Templates
- XSLTs

Monitoring a Business Process Thread

A thread is a basic unit of program execution. Threads perform the actual work in a process. A process can have several threads working concurrently, for example, transferring a file to one node using FTP, and to another node using HTTP. The Activity Engine is that part of a business process workflow engine (WFE) that calls the business process service, takes the results from the service, and immediately starts the next business process service cycle.

You can monitor the threads related to business processes in Sterling Integrator using the Thread Monitor. The Thread Monitor lists all the threads that are currently running in the Activity Engine. This information is useful for troubleshooting business processes that have either stopped or are hanging. You can also stop a thread from the Thread Monitor.

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

Monitoring a Thread

To monitor threads, from the **Administration** menu, select **Operations > Thread Monitor**. Sterling Integrator 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 at which the thread was registered

Note: You can also view thread details from the System Troubleshooting screen (**Operations > System > Troubleshooter**).

Interrupting a Thread

Important: Before using the Thread Monitor to interrupt a thread, contact Sterling Commerce Customer Support for assistance. Else, unexpected processing results

may occur, which may in turn affect other processes. It is recommended that you interrupt the business processes from the Business Process Monitor.

To interrupt a thread, click the **Interrupt** icon adjacent that thread. This instructs the thread to immediately interrupt processing without waiting for any steps to complete or exit with an error. In most cases, the thread will be interrupted immediately, but the display in the Thread Monitor will not be changed immediately. The thread will be marked as interrupted later.

Note: In some situations, it may not be possible to interrupt a thread, and your attempt may not succeed. Interrupting a thread does not stop the service the thread is using. The service is still available for Sterling Integrator to use in other business processes.

Stopping a Thread

Important: Before using the Thread Monitor to interrupt a thread, contact Sterling Commerce Customer Support for assistance. Else, unexpected processing results may occur, which may in turn affect other processes. It is recommended that you interrupt the business processes from the Business Process Monitor.

To stop a thread, click **Stop** icon adjacent to the thread that is to be stopped. 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. However, 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 Sterling Integrator 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 **Automatically refresh every 15 seconds** check box.

Monitoring Messages

In Sterling Integrator, business processes use messages to communicate with each other. Certain business processes create Produce and Consume messages. While the Produce business processes produce messages that are to be consumed, the Consume business processes wait for messages from the Produce business processes.

Sometimes the handoff from the Produce business processes to the Consume business processes does not occur. Either the produced message has no waiting Consume process, or the waiting Consume process waits without success for the Produce message. Use the Message Monitor to track unsuccessful message handoffs within Sterling Integrator.

It is recommended that you keep the following considerations in mind when monitoring messages within Sterling Integrator:

- The Message Monitor pages of the Sterling Integrator interface display only those 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 on which the associated business process step invokes the Produce or Consume activity. However, messages also have identities that are 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 having the same name. Let us for example consider that two Produce processes run at different times; both produce a message named OutMsg1. Thus, there are two messages named OutMsg1 in Sterling Integrator, 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, it produces InMsg1 for the Consume process with the earlier start date/time.
- Unsuccessful messages remain in Sterling Integrator until an appropriate Produce or Consume business process step completes the message handoff and uses the message.

Note: Unsuccessful messages never expire or time out.

Searching for a Message

To search for a message in Sterling Integrator:

1. From the **Administration** menu, select **Operations > Message Monitor**.
2. In the Message Monitor screen that is displayed, select one of the following search criteria and click **Go!**
 - **Message Name** – Displays messages whose names contain the specified character or string.
 - **Start Date** and **End Date** – Displays messages for which the related business process step invokes either the Produce or Consume activity at the specified start date and time. The start date and time that are displayed by default depends on whether Sterling Integrator has outstanding unsuccessful messages.
 - If there are no unsuccessful messages in Sterling Integrator, 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 Sterling Integrator, 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) – Displays those 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 pertaining to 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 of the business process.
- Process Name – Name of the associated Produce business process 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 Troubleshooting page (**Operations > System > Troubleshooter**) to review system information and troubleshoot system issues in Sterling Integrator.

From the System Troubleshooting page, you can perform the following tasks:

- View system information on different nodes from any node in a clustered environment.
- Refresh system status.
- Stop Sterling Integrator.
- View database usage statistics.

Note: If IBM DB2 is your database, the database usage statistics are always displayed as unavailable in the System Troubleshooting page.

- 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 adapter information.
- View perimeter server information.
- View when the archive, index, and purge tasks were completed.

The System Troubleshooting page provides quick access to system information. It provides links to key areas of system information instead of loading all of the system information at once. The system information for each area will be displayed in a pop-up window, which reduces the loading time of the System Troubleshooting page.

The following table provides a general description of each area. The subsequent sections provide further details about each area.

Area	Information
Stop the System	Stops Sterling Integrator using the softstop script.

Area	Information
Select Node	<p>The Select Node list is displayed only if you are working in a clustered environment. Your selection determines which node's information is displayed in the rest of the System Troubleshooting page.</p> <p>The select node list enables you to select a node in a clustered environment, which in turn decides the node whose information is to be displayed 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>
Host Information	<p>The Host Information displays the following information:</p> <ul style="list-style-type: none"> • Start time • Uptime • Host • Location • State • Memory available • Active threads
Gentran:Server	<p>This area is displayed only if you have configured your system for Gentran:Server Data Manager support.</p> <p>The Gentran:Server area displays the following information:</p> <ul style="list-style-type: none"> • Host • Location • State • Controllers • Adapters
Classpath	<p>Displays the Sterling Integrator system classpath and DCL (Dynamic Class Loader) configuration.</p>
JNDI Tree	<p>Displays the JNDI tree in Sterling Integrator.</p>
Database Usage	<p>Displays the database space usage, database services (business process eligibility for archive, index, and purge), and environment pool usage. Note: For Oracle, the database usage indicator calculates free space used based on Oracle extents. The name of this area changes along with the database that you are running. For example, oraclePool is for an Oracle database.</p>
Business Process Queue Usage	<p>Displays business process queue usage statistics such as cache disk usage, cache memory usage, queue statistics, and cache statistics.</p>
Business Process Usage	<p>Displays count of business processes by its state.</p>
Cache Usage	<p>Displays size and hit rate for object caches.</p>
Threads	<p>Displays active processes at a thread level.</p>
Clean-Up Processes Monitor	<p>Displays the time since the archive, index, purge, and recovery tasks were completed.</p>
Controllers	<p>The state and name of each controller or server in the Sterling Integrator installation.</p>

Area	Information
Adapters	Displays list of all the adapters in the system and their status. Note: Click an adapter name to view details about its status in the Service Settings page.
Perimeter Server Status	The information in this area is displayed only after you have added a perimeter server to Sterling Integrator. The Perimeter Servers area displays the following information: <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

Stopping Sterling Integrator from the System Troubleshooting Page

You can stop Sterling Integrator using the System Troubleshooting page. However, stopping Sterling Integrator in this manner stops Sterling Integrator using the softstop script, allowing all the business processes to complete before stopping the system.

Important: Using the Stop the System option stops only the Sterling Integrator interface immediately, while all the business processes that are in progress run until complete. After all the business processes' current services are completed, Sterling Integrator stops. To stop the system and all the processing immediately, in the *install_dir/bin* directory, run the hardstop script. All the processes that have not been completed will stop and have to be restarted.

To stop Sterling Integrator using the System Troubleshooting page:

1. From the **Administration** menu, select **Operations > System > Troubleshooter**.
2. In the System Troubleshooting page, in the System Status area, click **Stop the System**.
3. A message asking you for confirmation is displayed. Click **OK**.
The interface stops immediately, but all the business processes that are in progress are completed before the system stops.

Viewing the Host Information

Host Information displays the current operational status of the processing environment for your installation of Sterling Integrator. The System Troubleshooting page displays separate information for each installation.

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

To view Host Information:

1. From the **Administration** menu, select **Operations > System > Troubleshooter**.
2. In the System Troubleshooting page, click **Host Information**.

The following information is provided for each installation:

- The cluster node name, if you are working in a clustered installation of Sterling Integrator.

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

- Host - The name of the host on which a specific installation resides.
- Location - The location or path of the installation.
- State - The running state of the installation, either Active (available for processing) or Inactive.
- Memory in use - The amount of memory used by Sterling Integrator.
- Active threads - The number of concurrent threads that are active.

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. In the System Troubleshooting page, in the System Status area, click **Classpath**.
Information about the System Class Path and the Dynamic Class Loader is displayed.

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. In the System Troubleshooting page, in the System Status area, click **JNDI Tree**.
The system JNDI tree displays the JNDI name and class name pairs.

Viewing a Node JNDI Tree in a Clustered Environment

You can view a specific node's 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's JNDI tree in a clustered environment:

1. From the **Administration** menu, select **Operations > System > Troubleshooter**.
2. In the System Troubleshooting page, in the System Status area, click **node#**, where # is the number of the node you want to view information about.
The node's JNDI tree displays JNDI name, class name pairs, and the node name.

Viewing the Database Usage Statistics

Database usage statistics show how your database is performing, including database insert information, database capacity, and environment pool usage. Sterling Integrator uses pools to store database connections. To change pool settings, you must manually edit the configuration files and restart the system.

Note: Database usage statistics are not available for DB2.

To view database usage statistics:

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

Note: If you have the `displayGraphics` property, which is located in the `install_dir/properties/ui.properties` file, set to true, the Database Usage page is displayed in graphic format. Otherwise, the Database Usage page is displayed in text format. The default is true for Linux[®], Sun[™], HP, and Microsoft Windows[®] operating systems. The default is false for IBM AIX[®] and IBM iSeries[®] operating systems.

The Database Usage report is displayed with the following information for each unit test of the database:

- Average time taken to perform a given number of database inserts in a unit test
- Number of inserts performed to the database in a unit test

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

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

Note: You can change the value of the `dbAccessDataSize` property in the `install_dir/properties/ui.properties.in` file. After making the 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 part of the pool names changes depending on the database you are using, for example, if the database is Oracle, you will see `oraclePool`.

Viewing the Business Process Queue Usage

The Business Process Queue Usage page enables you to diagnose problems with your business process queues.

To view Business Process Queue Usage:

1. From the **Administration** menu, select **Operations > System > Troubleshooter**.
2. In the System Troubleshooting page, click **Business Process Queue Usage**.

The page provides the following queue information:

- The amount of memory available for cache and the amount consumed.
- The amount of disk space available for cache and the amount consumed.
- The average wait time based on priority.
- The average BP execution cycle time based on priority. The average BP execution cycle time may include the execution times of several steps. It captures the average time that BPs are active on threads before being rescheduled.
- The number of business processes in priority queues.
- The number of business processes that ran without being cached and the number that are currently in cache. Cache location is also specified so that you can determine the number of business processes that were found in the soft reference cache, in the disk cache, and in the memory cache.
- Number of business processes within the data size ranges that have been processed.

Viewing the Business Process Usage

The Business Process Usage page enables you to review the state of a business process and the process count. If a link is enabled in the Process Count area, you can select it to view affected business processes and manage them.

To view Business Process Usage:

1. From the **Administration** menu, select **Operations > System > Troubleshooter**.
2. In the System Troubleshooting page, click **Business Process Usage**.

The page provides the information about the following business process states and the process count:

- ASYNC_QUEUED
- Active
- Halted
- Halting
- Waiting
- Waiting_On_IO
- Interrupted_Man
- Interrupted_Auto

Viewing Cache Usage Information

Sterling Integrator uses caches to hold information that is frequently requested by the system. For each cache, the Cache Usage report displays statistics pertaining to the count, number of requests, and number of successful hits. To change cache settings, refer to the information provided in the topic Performance Tuning Utility. You can view the cache usage information to monitor the use of various cache types.

To view cache usage:

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

2. In the Application Status area of the System Troubleshooting page, 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, irrespective of whether it was found or not.
- Hits – Number of times an object was requested from the cache, and was found successfully.

Viewing Threads

The Threads pop-up window displays active processes at a thread level.

To view threads in Sterling Integrator:

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

The Threads report is displayed.

Viewing the Clean-Up Processes Monitor Details

You can view details about the time that has lapsed since the completion of different cleanup processes, including archiving, purging, and indexing. This helps you view, at a single location, information about whether these processes are running and completing.

To view the Clean-Up Process Monitor details:

1. From the **Administration** menu, select **Operations > System > Troubleshooter**.
2. In the System Troubleshooting page, in the System Status area, click **Clean-Up Processes Monitor**.

The Clean-Up Processes Monitor Details window displays the status, workflow (or cleanup process) name, the date and time at which the workflow was last run, and the workflow ID.

The Status column has the following values:

- Red – More than four times the scheduled interval has elapsed without a successful launch by the scheduler.
- Yellow – More than three times the scheduled interval has elapsed without a successful launch by the scheduler.
- Green – Less than three times the scheduled interval has elapsed since a successful launch by the scheduler.
- Gray – Either the process has never completed any scheduled instance or has never been scheduled.

Note: Red or Yellow status may mean that the process is not able to complete the data cleanup tasks. If this condition continues even after you have taken steps to resolve the errors, contact Sterling Commerce Customer Support.

Refreshing a Controller

You can refresh the controllers that are 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 System Troubleshooting page, click **Controllers**. The Controllers pop-up window appears and displays the state of the controller and name of the controller.
3. Click Refresh icons (arrows) next to the Controller you want to refresh.
The controller is refreshed.

Viewing Adapter Information

You can view the adapters by its state that are currently active or stopped, to verify accuracy or to plan changes as needed.

To view adapter settings:

1. From the **Administration** menu, select **Operations > System > Troubleshooter**.
2. In the Troubleshooting page, click **Adapters**.
3. In the Adapters page, click the name of the adapter that you want to view.

Enabling an Adapter

You can enable a disabled adapter using the System Troubleshooting page.

To enable an adapter:

1. From the **Administration** menu, select **Operations > System > Troubleshooter**.
2. In the Troubleshooting page, click **Adapters**.
3. In the Adapters page, next to the adapter that you want to enable, in the On/Off column, select the check box.

Disabling an Adapter

You can disable an enabled adapter using the System Troubleshooting page.

To disable an adapter:

1. From the **Administration** menu, select **Operations > System > Troubleshooter**.
2. In the Troubleshooting page, click **Adapters**.
3. In the Adapters page, next to the adapter that you want to disable, in the On/Off column, select the check box.

Viewing the Perimeter Server Status

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

- Name of the cluster node with which 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

Notes:

- If you are working in a clustered environment, the information that is displayed is determined by the node you select from the Select Node list.

- If you are not using a perimeter server, the perimeter server name is displayed as local, and the server state is Enabled.

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 of the System Troubleshooting page, under the On/Off column, select the check box adjacent to the perimeter server you want to enable.

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 of the System Troubleshooting page, under the On/Off column, select the check box adjacent to the perimeter server you want to disable.

Monitoring Node Status

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 information described in the following table:

Heading	Description
Name	Displays the names of the nodes.
URL	Displays the uniform resource locator of node.
Troubleshooter	Provides a link to the System Troubleshooter.
Token Node	Indicates whether this is a token node. <ul style="list-style-type: none"> • True – Indicates that this is the token node. • False – Indicates that this is not the token node.
Creation Time	Displays the date and time on which a node was created.

Heading	Description
Status	<p>Displays the status of a node.</p> <ul style="list-style-type: none"> • Active – Indicates that a node is working and is available for processing. • Node went down – Indicates that a node is not working and is not available for processing. • Node completely down – Indicates that all the services of a node have been stopped completely. • Starting Ops – Indicates that a node is starting up, but is not available for processing. <p>Notes:</p> <ul style="list-style-type: none"> • If it is indicated that a node has gone down, in the <i>install_dir/bin</i> directory of the node, run the <i>hardstop.sh</i> (for UNIX) or the <i>hardstop.cmd</i> (for Windows) script, and then run the <i>run.sh</i> (for UNIX) or the <i>startWindowsService.cmd</i> (for Windows) script. This stops all processing, and restarts the node. • Rerun the script of a node when its status is Node completely down. You can monitor the status of the restarted node in the other nodes of the cluster too.

2. To view more details about a node, click the name of the node in the Node Status page. This displays a details page containing the information described in the following table:

Heading	Description
Node Name	Displays the name of the node.
URL	Displays the uniform resource locator of node.
Token	<p>Indicates whether this is a token node.</p> <ul style="list-style-type: none"> • True – Indicates that this is the token node. • False – Indicates that this is not the token node.
Creation Time	Displays the date and time on which a node was created.
Status	<p>Displays the status of a node.</p> <ul style="list-style-type: none"> • Active – Indicates that a node is working and is available for processing. • Node went down – Indicates that a node is not working and is not available for processing. • Node completely down – Indicates that all the services of a node have been stopped completely. • Starting Ops – Indicates that a node is starting up, but is not available for processing. <p>Notes:</p> <ul style="list-style-type: none"> • If it is indicated that a node has gone down, in the <i>install_dir/bin</i> directory of the node, run the <i>hardstop.sh</i> (for UNIX) or the <i>hardstop.cmd</i> (for Windows) script, and then run the <i>run.sh</i> (for UNIX) or the <i>startWindowsService.cmd</i> (for Windows) script. This stops all processing, and restarts the node. • Rerun the script of a node when its status is Node completely down. You can monitor the status of the restarted node in the other nodes of the cluster too.
Location	Displays the directory path on which the node is installed.
Role	States the role of the node.

Heading	Description
Operation Controller Host	Displays the name of the server acting as the Operation Controller Host.
Operation Controller Port	Displays the port number of the Operation Controller.
Sterling Integrator Version	Displays the version of Sterling Integrator that is installed in the node.
JVM Version	Displays the version number for the Java Virtual Machine (JVM™).
JVM Vendor	Shows the vendor who provided the JVM™.
The information described in the following rows is displayed only if <code>queueWatcher=true</code> in the <code>install_dir/properties/noapp.properties_platform_ifcresources_ext</code> file.	
Multicast Info	<p>Provides information about the communication mechanism across nodes. This is used to provide load factor and status information about a node to all the other nodes in a cluster.</p> <p>Clustering uses the JGroups communication toolkit, which supports the following communication protocols, which can be configured during deployment:</p> <ul style="list-style-type: none"> • UDP (User Datagram Protocol) (IP multicast) • TCP
Engine Queue Content	<p>Provides a link to the Engine Queue Content page for the node and queue. This page contains details about active threads, and the items in the queue that are waiting for a thread.</p> <p>Use the drop-down menu to select the queue you want to review</p>
Distribution Size	<p>Provides a link to the Distribution Size page. This page contains the following information:</p> <ul style="list-style-type: none"> • Scheduling policy details • Soft reference cache details • Queue details (for each queue)

Monitoring Deprecated Resources

Sterling Integrator includes a tool that helps you monitor resources that will be deprecated in future releases. The deprecated resources include services, adaptors, business processes, and maps. Following are the reasons for the deprecation:

- An existing resource has been replaced by an improved version of the same resource.
- An existing resource has become obsolete because of changes in the behavior of Sterling Integrator.

Note: Resources that have been identified for deprecation will continue to be supported by Sterling Integrator.

Manage the deprecated resources as follows:

1. Identify the deprecated resources using the Sterling Integrator Resource Deprecation Report.
2. Remove and replace the deprecated resources.

Identifying Deprecated Resources

To identify the deprecated resources, perform the following tasks:

1. From the **Administration** menu, select **Operations > Reports**.
2. In the Reports screen, under the Lists section, select the letter D.
3. Click **Go!**. The Reports page displays all the reports starting with the alphabet D.
4. For the Deprecation Report, click the **source manager** icon.
5. Select the format in which you want the report output, for example, PDF, HTML, or XLS.
6. Click the execute icon.

The Sterling Integrator Resource Deprecation Report is displayed, containing the columns described in the following table:

Column Heading	Description
Resource Name	The name of the resource that will be deprecated. Examples: <ul style="list-style-type: none"> • B2B_FTP_CLIENT_ADAPTER • DeenvelopeGSType • Reconcile997Type
Resource Type	The type of resource that will be deprecated. Example: service_instance
Date Deprecated	The date on which the resource will be deprecated. Example: 2/20/13 8:00 AM
Notes	Information about how the deprecated resource will be replaced. Examples: <ul style="list-style-type: none"> • Replaced by FTP Client Adapter • Use the DeenvelopeEDIFACT service • Use the DeenvelopeX12 service

Removing and Replacing the Deprecated Resources

The following table lists resources that have been identified for deprecation.

Note: Simply viewing a business process is not sufficient to identify the resources that are being used in deprecated business resources. Some predefined business processes determine the resources to be used, only at runtime. For example, some EDI enveloping and deenveloping business processes determine which map to use at runtime based on header information.

Deprecated Resource Name	Replaced By
997 Generator	X12 Deenvelope Service
997 Reconcile	X12 Deenvelope Service
B2B FTP Client	FTP Client Adapter

Deprecated Resource Name	Replaced By
B2B HTTP Client Adapter	<p>Migration of this adapter involves adding three PS HTTP Services into the business process:</p> <ul style="list-style-type: none"> • HTTP Begin Session Service • HTTP Post/Get Service • HTTP End Session Service <p>The B2B HTTP Client Adapter is invoked through the B2B Send Service, which has the option to specify the contract ID. If a business process has included the B2B HTTP Client Adapter, you must make the necessary changes to pass in the profile ID of the Consumption Profile to the PS HTTP adapter because the PS HTTP adapter does not take the contract ID.</p>
B2B HTTP Communications Adapter	HTTP Client Adapter
B2B HTTP Server Adapter	<p>The migration of the B2B HTTP Server Adapter involves:</p> <ul style="list-style-type: none"> • Migrating the URIs defined in the legacy server instance to either the default PS HTTP Server instance or a newly created instance. • Migrating the bootstrapped business processes to use the HTTP Respond Service when attempting to send a response back to the trading partner.
Connect:Direct Adapter	Connect:Direct Requester Adapter
Connect:Enterprise Adapter	Connect:Enterprise Unix Server Adapter
ebXML Manifest Creation Service	ebXML Message Handler Service
FTP Get Adapter	FTP Client Adapter
FTP Send	FTP Client Adapter
Generate CONTRL	EDIFACT Deenvelope Service
GS De-envelope	X12 Deenvelope Service
GS Envelope	X12 Envelope Service
HTTP Communications Adapter	PS HTTP Client Adapter
HTTP Send Adapter	<p>Migration of the HTTP Send Adapter involves adding three PS HTTP Services into your business process. The three services are:</p> <ul style="list-style-type: none"> • HTTP Begin Session Service • HTTP Post/Get Service • HTTP End Session Service
HTTP Client Adapter	HTTP Client Adapter
ISA De-envelope	X12 Deenvelope Service
ISA Envelope	X12 Envelope Service
LightweightJDBCEngineService	Lightweight Java Database Connectivity (JDBC) Adapter

Deprecated Resource Name	Replaced By
LightweightJDBCEngineServiceType	Lightweight Java Database Connectivity (JDBC) Adapter
Reconcile CONTRL	EDIFACT Deenvelope Service
ST De-envelope	X12 Deenvelope Service
ST Envelope	X12 Envelope Service
TA1 Generator	X12 Deenvelope Service
TA1 Reconcile	X12 Deenvelope Service
UNB De-envelope	EDIFACT Deenvelope Service
UNB Envelope	EDIFACT Envelope Service
UNG De-envelope	EDIFACT Deenvelope Service
UNG Envelope	EDIFACT Envelope Service
UNH De-envelope	EDIFACT Deenvelope Service
UNH Envelope	EDIFACT Envelope Service
XML Transformer	DocToDOM, DOMToDoc, and XSLT Service for translation.

Chapter 6. Performance Tuning

Performance Tuning Utility

Sterling Integrator provides a utility that calculates the tuning parameter values for you. Using the Performance Tuning Utility, you can tune and configure Sterling Integrator according to your business requirements. You can change the database settings, memory allocations, and pool settings to enable Sterling Integrator to perform according to your requirements. However, note that you cannot use the Utility to tune your business processes. For more information about tuning business processes, refer to *Sterling Integrator Business Process Modeling Guide*.

This topic provides information about calculating the settings required and viewing and editing performance configuration settings.

Calculating the Settings Required

Sterling Integrator calculates the settings that are required based on the number of cores 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 formulae in the `tuningFormulas.properties` file in the `install_dir/properties` directory. The calculations for determining the appropriate number of threads is only a guideline, using which you can achieve a well-tuned system.

Repeat the following procedure until your system is well tuned:

1. Run a one-hour load test on the system on which you have installed Sterling Integrator. Record the results for CPU utilization and application performance. Monitor the average utilization and peaks. If the peaks are sustained for more than approximately 20 minutes, improper sizing may result in substantial queuing of work while waiting for resources. The nature of the application and the business problem it is solving determines whether you should use average utilization or peak utilization for determining the desired performance levels.
2. Increase or decrease the number of threads by approximately 10% for each iteration. The number of threads that are increased or decreased should be a multiple of the number of cores, or one thread for every two cores, where core is an independent processing engine.

Note: Monitor all the nodes in a clustered deployment.

3. Run the one-hour load test again. Record the results for CPU utilization and application performance.
4. Compare the results of CPU utilization and application performance to determine if your new thread allocation has improved performance.

Note: An increase in the number of global threads may require you to increase the amount of memory allocated to the Sterling Integrator JVM. If out-of-memory conditions occur during testing, or if CPU utilization increases because of JVM garbage collection activities, increase the amount of memory allocated to the JVM. In general, garbage collection activity should not account for more than 3% of CPU utilization.

On a 32-bit JVM, you cannot increase the memory of the JVM beyond 2 GB.

If you are still unable to achieve a well-tuned system, you may want to engage Sterling Commerce Professional Services. Contact your Sales Representative for more information about this.

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

Performance Tuning Property (Formula)	Default Value	User-Identified Value
Number of cores in system (shipped value = 4)	2	6
Physical memory (MB) allocated to Sterling Integrator (shipped value = 1024)	768	2048

The following properties are calculated according to the number of cores and physical memory	Default 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	4096
Transactional pool connections (initial) ((4 * Processor Number) + 5)	13	29
Transactional pool connections (max) ((4 * Processor Number) + 30)	38	54
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

The following properties are calculated according to the number of cores and physical memory	Default Value	Calculated Value
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

This topic provides information about the following tasks:

- Viewing Performance Configuration Settings
- Editing Performance Configuration Settings

Viewing Performance Configuration Settings

You can view the performance configuration settings to verify if a particular performance configuration setting is correct.

To view your performance configuration settings:

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

Editing Performance Configuration Settings

To reach the level of performance you require, you may have to allocate the following items in higher or lower quantity to specific components of Sterling Integrator:

- System memory
- Database pools
- Cache
- Memory setting for MySQL (if applicable)

You can use the performance tuning utility to accomplish your goals and meet your business requirements.

You should configure the performance settings 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 topic Performance and Tuning Worksheet and tune Sterling Integrator according to the requirements specified in the worksheet.

You can use the Performance Tuning Utility to edit the Sterling Integrator performance configuration using the settings calculated by the system. Alternatively, if you have specific requirements, you can edit the existing performance configuration without using the settings calculated by the system.

To edit your performance configuration:

1. From the **Administration** menu, select **Operations > System > Performance > Tuning**.
2. In the Performance Tuning page, under **Edit**, next to **Edit Performance Configuration**, click **Go!**.
3. On clicking Go, the Performance Tuning UI lock-enabled message is displayed:
 Performance Tuning UI lock enabled.
 Lock will be released only by clicking on Cancel or Finish in the wizard.
 Click **OK**. 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. In the Edit Mode page that is displayed, select one of the following options:
 - **Generate recommended settings** – On selecting this option, Sterling Integrator calculates the recommended settings based on the number of CPUs and amount of memory you specify. Refer to the topic Calculating the Settings Required for more information about the settings that have been calculated. 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 fields under the Recommended column, based on your performance requirements.
 - **Edit settings** – On selecting this option, Sterling Integrator displays 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. In the System page that is displayed, decide whether you want to use the current settings or the recommended settings.
 - If yes, click **Next**.
 - If no, enter information in the following fields and click **Next**.

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

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

7. In the BP Queue page that is displayed, decide whether you want to use the current settings or the recommended settings.
 - If yes, click **Next**.
 - If no, enter information in the following fields and click **Next**.

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

8. In the Database Pool page that is displayed, decide whether you want to use the current settings or the recommended settings.
 - If yes, click **Next**.
 - If no, enter information in 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. 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 have to increase the value of this property to improve performance.

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 Sterling Integrator starts. If Sterling Integrator uses all the connections concurrently, it creates a new connection until it reaches the maximum number of connections.</p> <p>This field changes the MIN_TRANS_POOL property in the tuning.properties file in the <i>install_dir</i>/properties directory. Shipped value is 13.</p>

Field	Description
Transactional pool connections (max)	<p>Maximum number of connections that are reserved and available for critical data requests that Sterling Integrator can use concurrently. After a connection is no longer in use, the connection is returned to the available connections pool and can be reused. When Sterling Integrator reaches the value specified in this field, it must wait for a connection to be returned to the pool before processing the data request.</p> <p>This field changes the MAX_TRANS_POOL property 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 Sterling Integrator starts. If Sterling Integrator uses all the connections concurrently, it creates a new connection until it reaches the maximum number of connections.</p> <p>This field changes the MIN_NONTRANS_POOL property 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 Sterling Integrator can use concurrently. When a connection is no longer in use, the connection is returned to the available connection pool and can be reused. When Sterling Integrator reaches the value specified in this field, it must wait for a connection to be returned to the pool before processing the data request.</p> <p>This field changes the MAX_NONTRANS_POOL property 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 have to increase the database pools to reduce the chances of a business process or adapter having to wait for a database connection.

9. In the Memory page that is displayed, decide whether you want to use the current settings or the recommended settings.
 - If yes, click **Next**.
 - If no, enter information in the following fields and click **Next**.

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

Field	Description
JVM short-lived memory (max)	<p>Maximum amount of JVM memory that Sterling Integrator uses for short-lived objects. This field changes the MAX_AGE property in the tuning.properties file in the <i>install_dir</i>/properties directory. Shipped value is 256.</p> <p>Note: JVM short-lived memory has a fast garbage collection rate. To avoid reduced performance, set the JVM short-lived memory to one-third of the JVM long-lived memory.</p>

10. In the BP Execution page that is displayed, decide whether you want to use the current settings or the recommended settings.
- If yes, click **Next**.
 - If no, enter information in the following fields and click **Next**.

Field	Description
Number of BP steps executed before returning to queue	<p>Number of business process steps that are run before returning to the queue. This field changes the BP_STEPS property in the tuning.properties file in the <i>install_dir</i>/properties directory. Shipped value is 100.</p> <p>Typically, this property must match the number of steps in your largest business process.</p> <p>If your business processes contain loops, you may have to increase this property to improve performance. If you find that you have to increase this number, you must streamline your business processes to remove unnecessary loops.</p> <p>If you set this value too low, your 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 time and causes bottlenecks in Sterling Integrator.</p>
Initial steps in the first execute cycle	<p>Number of steps executed before returning a business process to the queue on its first execution cycle. This field changes the ASYNC_BP property in the tuning.properties file in the <i>install_dir</i>/properties directory. Shipped value is 2.</p>

11. In the Cache page that is displayed, decide whether you want to use the current settings or the recommended settings.
- If yes, click **Next**.
 - If no, enter information in the following fields and click **Next**.

Field	Description
Translation maps	<p>Amount of cache used to store translation maps. This field changes the MAP_CACHE property in the tuning.properties file in the <i>install_dir</i>/properties directory. Shipped value is 100.</p> <p>Note: Increasing the cache size for items that are not used frequently may degrade performance. Increase the cache size for items that are used frequently to improve performance.</p>

Field	Description
Envelopes	Amount of cache used to store envelopes. This field changes the ENVELOPE_CACHE property in the tuning.properties file in the <i>install_dir</i> /properties 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 in order to improve performance.
Other EDI	Amount of cache used to store EDI data. This field changes the EDI_CACHE property in the tuning.properties file in the <i>install_dir</i> /properties 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. In the Confirm page that is displayed, review the changes you made to the performance configuration, and perform one of the following tasks:
 - To change your performance configurations, click **Back**.
 - To abandon the performance configurations you carried out, click **Cancel**.
 - To save the performance configurations you carried out, click **Finish**.
13. From the **Administration** menu, select **Operations > System > Troubleshooter**.
14. From the System Troubleshooting page, click **Stop the System**. This runs the softstop script, allowing a fixed amount of time for the business processes to complete before stopping the system.
15. From the *install_dir*/bin directory, run one of the following commands:
 - (UNIX or Linux) **setupfiles.sh**
 - (Windows) **setupfiles.cmd**

The performance configuration changes populate the initial setup files that are used during startup.
16. From the *install_dir*/bin directory, start Sterling Integrator by running one of the following commands:
 - (UNIX or Linux) **run.sh**
 - (Windows) **startWindowsService.cmd**

Sterling Integrator starts and the updated performance configuration settings are applied.

Manual Performance Tuning

You may find that you have to tune some application server-independent (ASI) properties that are not tuned into the Performance Tuning Utility. You can manually tune these properties in the *install_dir*/properties/noapp.properties file or extension files (ext) of that file.

This topic provides information about tuning your ASI environment, scheduling policies, tuning queues, and changing persistence levels. It also provides information about concepts such as customizing property files, changing system logging, enabling trusted domains for schemas, managing system recovery, and optimizing system performance for Sterling e-Invoicing.

Note: For large volume operations pertaining to FTP, Mail Box, Advanced File Transfer, and Sterling File Gateway, set the number of open file descriptors to at least 4096, using the command **ulimit -n 4096**.

Tuning Your ASI Environment

About this task

You can tune your application server-independent (ASI) system using the properties found in the *install_dir/properties/noapp.properties_platform_ifresources_ext* file, or in the *install_dir/properties/noapp.properties_gis_ext.in* file (for the *persistence_level* property).

Properties that you may have to tune include:

Property	Description
SchedulingPolicyName	Set of rules that Sterling Integrator uses to make decisions about how to manage the workload. Valid values are: <ul style="list-style-type: none"> BasicSchedulingPolicy FairShareSchedulingPolicy Note: For more information about scheduling policy, refer to the topic Scheduling Policy.
MemCacheSize	Size, in megabytes, of the in-memory cache used for small objects. This cache can speed up execution by eliminating the necessity for 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 large, the value of MemCacheSize should also be large. A moderate multiple of the disk block size seems to work well in many cases.
MaxThreads	Total number of concurrent threads that Sterling Integrator is allowed to use. Verify this number 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 enable smoother multitasking. 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. .# indicates the queue number.
AE_ExecuteCycleTime.#	Maximum time period, in milliseconds, for which a business process can use a thread for, before releasing it for use by another business process. This value will override the value set for the AE_ExecuteCycle. This is intended to ensure that a series of unusually slow steps do not tie up a thread completely. This value can be different for each queue. .# indicates the queue number.
QueueDepth.#	Maximum number of business processes that can be in a queue at a given time. Generally, this value should be left at the default value of 10,000, 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. .# indicates the queue number.
MaxPoolSize.#	Maximum number of threads executed for a 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. .# indicates the queue number. Note: Setting all the queues' MaxPoolSize parameters to the maximum threads available does not mean that 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 fewer threads.

Property	Description
CacheThreshold.#	<p>Number of business processes that must be in the queue before any business process is 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.</p> <p>.# 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 Sterling Integrator System Troubleshooter. For more information about reviewing system, refer to the topic Reviewing System Information.</p> <p>There is a tradeoff between performance and memory consumption. Setting this value too high can result in one of the following:</p> <ul style="list-style-type: none"> • Leave Sterling Integrator with insufficient memory in some circumstances • Cause some business processes to remain in the queue for 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.</p>
MinPoolSize.#	<p>Minimum number of threads reserved for a 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.</p> <p>.# 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 user interface. This value can be different for each queue. Not all JVMs handle this the same way, with some not respecting it at all (particularly, the IBM AIX® JVM).</p> <p>.# indicates the queue number.</p>
EnableDeadlines.#	<p>Enables or disables deadline support for a particular 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 a queue. • true – Enables deadline support for a queue. <p>This value can be different for each queue.</p> <p>.# indicates the queue number.</p>
Rescheduling.#	<p>Enables or disables rescheduling support for a queue. Valid values are:</p> <ul style="list-style-type: none"> • false – Business processes are not rescheduled in a queue. • true – Business processes are rescheduled in a queue. <p>This value can be different for each queue.</p> <p>.# 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.</p> <p>.# indicates the queue number.</p>
MaxWaitTime.#	<p>Maximum time, in milliseconds, for which 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.</p> <p>.# indicates the queue number.</p> <p>Example: MaxWaitTime.8=60000 sets the maximum wait time for queue number 8 as 60 seconds.</p>

Property	Description
ResourceAllocation.#	<p>Amount of resources to allocate to a 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 for each queue by a percentage. You can set queue 1 with a resource allocation of 10 and queue 2 with a resource allocation of 90. This ensures that queue 2 has more resources for processing, which in turn improves processing time and ensures efficient use of system resources. This value can be different for each queue.</p> <p>Note: Set the resource allocations according to the percentage 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 in order to increase processing capability. Apply a lower percentage of resources to lower priority queues or less used queues in order to keep resources free for higher priority processing.</p> <p>.# indicates the queue number.</p>
InitialCycles.#	<p>Number of cycles to be executed the first time a business process gets to execute. Normally, this value should be 5 or less. This 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.</p> <p>.# indicates the queue number.</p>
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) steal only from queues that have the AllowStealing property set to true. This value can be different for each queue.</p> <p>.# 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 – Allows other queues to steal threads. • false – Does 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.</p> <p>.# indicates the queue number.</p>
persistence_level	<p>Persistence (storage to the database) level for Sterling Integrator. Set the persistence_level value to PERSISTENCE_NONE during day-to-day processing.</p> <p>persistence_level 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 the documents, status report (if any), and instance data. • PERSISTENCE_MINIMAL/PERSISTENCE_STEP_STATUS – From user interface tracking, every step will be persisted, but the document, status report (if any), and instance data will not be persisted in some of the steps (for example, assign and sleepservice). Most of the services that create documents will be persisted in full.

Property	Description
	<p>persistence_level valid values continued:</p> <ul style="list-style-type: none"> • 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 and sleepservice). Services that are defined as full will be persisted in full. Documents created through a stream and stored in the file system will be persisted. Note: If you have configured a business process with INVOKE_MODE set to SYNC, which in turn invokes business processes with INVOKE_MODE set to INLINE and the persistence_level set to BP Start Stop Engine May Override, the business processes may fail. To counter the business process failure, follow either of these workarounds: <ul style="list-style-type: none"> – Change the INVOKE_MODE of the second business process from INLINE to SYNC. – Add a dummy <assign> statement after the <operation> statement as shown: <pre data-bbox="537 541 1073 961"> <process name = "Call_BP_withError.kj"> <sequence name="Sequence_Start"> <operation name="Invoke Sub-Process"> <participant name="InvokeSubProcessService"/> <output message="InvokeSubProcessServiceTypeInputMessage"> <assign to="INVOKE_MODE">INLINE</assign> <assign to="WFD_NAME">Called_BP_withError.kj</assign> <assign to="." from="*"></assign> </output> <input message="inmsg"> <assign to="." from="*"></assign> </input> </operation> <assign to="." from="*"></assign> </sequence> </process> </pre> • PERSISTENCE_OVERRIDE_MINIMAL/PERSISTENCE_STEP_STATUS_ONLY – The first step will be persisted in full. The last step will be persisted minimally. Documents created through 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 minimally. Some of the services will not be persisted (for example, assign and sleepservice). Services that are defined as full will not be persisted. Documents created through a stream and stored in the file system will be persisted.
	<p>persistence_level valid values continued:</p> <ul style="list-style-type: none"> • 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 minimally. Documents created through 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 Sterling Integrator engine, and if the service is creating documents in the document stream, or if FS is document storage, the document will be persisted. If it is running in an 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 performing document stream FS as document storage, only the first step will be persisted, or, when the business process generates errors, the first step (initial workflow context) and the error step will be persisted. The error step will be persisted as minimal in order to avoid the problem of resuming. If there is an onfault in the business process and the business process generates errors, the first step (initial workflow context), error step, and last step from the onfault will be persisted. If the onfault block runs fine, or if the onfault block generates errors too, the error step in onfault will be persisted minimally instead of the last step in onfault. If a step is persisted, the first step's initial workflow context is always persisted in full. <p>Note: This property is set in the noapp.properties_gis_ext.in file. If the persistence level property is a comment in the noapp.properties_gis_ext file (#persistence_level=value), the persistence level is the default value of full (PERSISTENCE_FULL).</p>

Scheduling Policy

About this task

Sterling Integrator uses scheduling policies to manage workload. A scheduling policy is a set of rules that Sterling Integrator uses to determine how workload is managed. Workload management decisions include, but are not limited to:

- The number of threads to run at a given time.
- The order in which to run the business processes.
- The number of steps a business process should run before releasing the thread for use by other business processes.
- The caching requirements when a business process is returned to the queue.
- The workload distribution, if you are working in a clustered environment.

Scheduling policy choices include:

- Basic Scheduling Policy – Recommended for consistent workloads that do not have data processing peaks or change in data or processing types.
- Fair Share Scheduling Policy – Recommended for mixed workloads that include both batch and online processing, and in environments that have data processing peaks (Default). For more information about this policy, refer to the topic Queues in an Application Server Independent Environment.

Changing a Scheduling Policy

To change a scheduling policy in Sterling Integrator:

Procedure

1. Open the *install_dir*/properties/noapp.properties_platform_ifcresources_ext file using a text editor.
2. In the noapp.properties_platform_ifcresources_ext file, locate the following code: # SCHEDULING POLICY CONFIGURATION
3. Under # SCHEDULING POLICY CONFIGURATION, locate the following property: SchedulingPolicyName=
4. Adjacent to the SchedulingPolicyName parameter, enter one of the following code samples, depending on your choice of either basic or fair-share scheduling policy. Next to the SchedulingPolicyName parameter, enter:
 - 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_platform_ifcresources_ext file. Changing either the name or the location of the noapp.properties_platform_ifcresources_ext file may prevent your changes from taking effect.

Queues in an Application Server Independent Environment

About this task

Sterling Integrator uses fair-share scheduling, which is in-memory queuing and caching, to enable you to set a global thread limit and allocate resources on a

queue-by-queue basis. This enables flexible management of your workload. The global thread limit is the maximum number of business processes that can run simultaneously in Sterling Integrator. You can change the thread limit and queue the resource allocations in the *install_dir/properties/noapp.properties_platform_ifcresources_ext* file. The thread limit is set in the MaxThreads property. For more information about the MaxThreads property, refer to the section on tuning queues.

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

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

The following example procedure explains the fair-share scheduling process:

1. Open the *install_dir/noapp.properties_platform_ifcresources_ext* file using a text editor.
2. Set the global number of threads as 16 by entering 16 as the new MaxThreads value.
3. Determine that queue Q1 is for small items and set the MinPoolSize.1 value to 10 and the MaxPoolSize.1 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.
4. Determine that according to the number of small items requiring processing, queue Q1 requires a resource allocation of 70% of the total resources, and you enter 70 as the new ResourceAllocation.1 value.
5. Determine that queue Q8 is for batch processing and set the MinPoolSize.8 value to 1 and the MaxPoolSize.8 value to 4.
6. Determine that according to the number of batch items requiring processing, queue Q8 requires a resource allocation of 30% of the total resources, and you enter 30 as the new ResourceAllocation.8 value.
7. During processing, different conditions arise, as shown in the following examples:
 - If there is no batch processing for queue Q8, but a lot of processing for queue Q1, queue Q1 uses all the 16 threads to complete processing faster.
 - If there is some batch processing for queue Q8, but also some processing for queue Q1, queue Q1 uses up to 12 threads to complete processing faster, while queue Q8 uses up to 4 threads to complete processing faster. These numbers are according to the percentages applied to the MaxThreads value of 16. In the case of queue Q1 with 70% 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.8 is set to 4, only 4 threads can be used by queue Q8, enabling the 12th thread to be used by queue Q1.
 - If there is no processing for queue Q1, but a lot of processing for queue Q8, queue Q8 uses all 4 threads to complete processing faster, because the MaxPoolSize.8 for queue Q8 is 4.

Tuning Queues

To tune the queues in Sterling Integrator:

Procedure

1. Open the `install_dir/noapp.properties_platform_ifcresources_ext` file using a text editor.
2. In the `noapp.properties_platform_ifcresources_ext` file, under the `# 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 guarantee that 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_platform_ifcresources_ext` file, under the `# QUEUE CONFIGURATION, Queue # - ALL Policies`, where `#` indicates the queue number, locate the properties described in the following table and change the values for each as appropriate. Complete this step for each queue that you are using. For each property, `#` indicates the queue number, for example, `AE_ExecuteCycle.3` is the `AE_ExecuteCycle` property for the third queue. Also, property values can be different for each queue.

Property	Description
<code>AE_ExecuteCycle.#</code>	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 enables smoother multitasking. Interactive use favors a lower number of steps, while batch processing favors a higher number of steps.
<code>AE_ExecuteCycleTime.#</code>	Maximum time period, in milliseconds, that a business process can use a thread, for releasing it for use by another business process. This value will override the value set for <code>AE_ExecuteCycle</code> . This is intended to ensure that a series of unusually slow steps do not tie up a thread completely.
<code>QueueDepth.#</code>	Maximum number of business processes that can be in a queue at a given time. Generally, this value should be left at the default value of 10,000, unless you anticipate having more that 10,000 business processes in a queue at the same time.
<code>MaxPoolSize.#</code>	Maximum number of threads executed for a specific queue. The <code>MaxThreads</code> value overrides the <code>MaxPoolSize</code> value if the <code>MaxPoolSize</code> value is set higher than the <code>MaxThreads</code> value. Note: Setting all the queues' <code>MaxPoolSize</code> parameters to the maximum threads available does not mean you will have faster processing. Depending on your system resources, setting the <code>MaxPoolSize</code> 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.

Property	Description
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.</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 by using the Sterling Integrator System Troubleshooter. For more information, refer to the topic Reviewing System Information.</p> <p>There is a trade-off between performance and memory consumption. Setting this value too high can lead to the following:</p> <ul style="list-style-type: none"> • Leave Sterling Integrator with insufficient memory in some circumstances • Cause some business processes to remain in the queue for 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 when the context is placed in the queue.</p>
MinPoolSize.#	<p>Minimum number of threads reserved for a specific queue. The sum total of all the queues' MinPoolSize values must be equal to or less than the MaxThreads value. The MinPoolSize value can be different for each queue.</p> <p>.# indicates the queue number. Note: Setting the MinPoolSize parameter too low may cause queues to back up and degrade performance.</p>
JavaPriority.#	<p>Java priority of the threads that are running business processes. This enables you to set some queues to run more slowly in order to reduce the issues associated with heavily used computers being unresponsive to the user interface. Not all JVMs handle this the same way, with some not respecting it at all (particularly, the IBM AIX® JVM).</p>
EnableDeadlines.#	<p>Enables or disables deadline support for a queue. Having the deadlines enabled controls both the 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 a queue. • true – Enables deadline support for a queue.
Rescheduling.#	<p>Enables or disables rescheduling support for a queue. Valid values are:</p> <ul style="list-style-type: none"> • false – Business processes are not rescheduled in a queue. • true – Business processes are rescheduled in a queue.
ReschedulingInterval.#	<p>Time interval, in milliseconds, to wait before the rescheduler passes through the queue to reschedule business processes.</p>
MaxWaitTime.#	<p>Maximum time for which 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.</p>

4. In the noapp.properties_platform_ifcresources_ext file, under the # QUEUE CONFIGURATION, Queue # - FairShareSchedulingPolicy, where # indicates the queue number, locate the properties described in the following table and change the value for each as appropriate. Complete this step for each queue you are using. For each property, # indicates the queue number, for example, ResourceAllocation.3 is the ResourceAllocation property for the third queue. Also, property values can be different for each queue.

Property	Description
ResourceAllocation.#	Amount of resources to be allocated to a specific queue for fair share scheduling. This ensures that the queue has more resources for processing, which improves processing time and ensures the efficient use of system resources. Note: Set the resource allocations according to the percentage 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 in order to increase processing capability. Apply lower percentages of resources to lower priority queues or less used queues in order to keep resources free for higher priority processing.
InitialCycles.#	Number of cycles to be executed the first time a business process gets to execute. Normally, this value should be 5 or less. This 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.
StealThreads.#	Enables or disables the ability to steal threads from other queues. Valid values: <ul style="list-style-type: none"> • true – Enables the ability to steal threads • false – Disables the ability to steal threads Queues that are configured to steal threads from other queues (StealThreads property set to true) steal from only queues that have the AllowStealing property set to true.
AllowStealing.#	Enables or disables the ability of other queues to steal threads from this queue. Valid values: <ul style="list-style-type: none"> • true – Allows other queues to steal threads • false – Does not allow other queues to steal threads Queues that are configured to steal threads from other queues (StealThreads property set to true) steal from only queues that have the AllowStealing property set to true.

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

Important: Saving the `noapp.properties_platform_ifcresources_ext` file under a different name or in a different location will prevent the changes from being applied to Sterling Integrator.

6. Stop Sterling Integrator.
7. In the `install_dir/bin` directory, apply these changes by running one of the following commands:
 - (For UNIX or Linux) `setupfiles.sh`
 - (For Windows) `setupfiles.cmd`
8. Start Sterling Integrator. The changes are applied to Sterling Integrator, and override the changes made in the `install_dir/properties/tuning.properties` file.

Monitoring Queues using Queue Watcher

Queue Watcher monitors various components in Sterling Integrator as well as manages queue configuration settings.

Accessing Queue Watcher

About this task

To access Queue Watcher, do the following:

Procedure

1. Open your web browser to `http://host:port/queueWatcher`, where *host:port* is the IP address and port number where Sterling Integrator resides on your system. A login page appears.

Note: Any user with Administrator privileges can login to the Queue Watcher application, provided he has all the necessary permissions or is a part of the Sterling Integrator Administrator group.

2. Type your username and password. The Queue Watcher displays the following information:

Heading	Description
View Active Threads for All Queues	<p>Displays a list of all active queue threads. When selected, you can review the following information:</p> <p>QueueName – Displays the queue name.</p> <p>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.</p> <p>Used – Number of business processes currently running on a thread.</p> <p>Calc – Fairshare thread calculation for the queue. Fairshare is based on concurrent activities on all queues and is dynamically updated.</p> <p>Pool – Number of threads in a queue’s pool. Threads timeout if they are not used.</p> <p>Max - Maximum number of threads used by the queue. <i>Calc</i> determines the maximum concurrent threads that is dynamically calculated.</p> <p>Queue Depth - Number of business processes waiting for a thread in the queue.</p> <p>List of Working Threads – List of business processes currently running on a thread.</p>
Pause All Queues	Use this option to stop queues. Stopping individual queues is not possible.
Restart All Queues	Use this option to restart queues. Restarting individual queues is not possible. DBResources will use this command if the database becomes unavailable.
View Default Queue Configuration Parm	Displays the parameters set for all of the queues.
View Active Queue Configuration Parm	Displays the current queue configuration.
View list of Workflow IDs that recover would see in the queue	<p>Displays the workflow ID when it is run or moved to another node in the cluster. Valid values are:</p> <p>Executed</p> <p>Moved to another (cluster) node</p>

View Context Cache Entries	<p>Displays the context cache entries. Note: If entries show up as invalid they are still correct and do not indicate an error.</p> <p>Soft Reference Cache Slots in use - Workflow Context (wfc) is saved into this queue (hashtable) and can be recovered from it. This is the fastest back queue. If required, the garbage collector can acquire more heap space from this queue. The workflow contexts are not serialized on this queue.</p> <p>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.</p> <p>Disk Cache Bytes in use - This cache holds workflow contexts larger than the defined threshold. The workflow contexts are serialized on this queue.</p>
Wait Queue	Displays the workflow IDs when the Wait Service is being processed. The Wait Service will only appear if the wait interval is less than 30.
Queue_1 – Queue_9	Displays running and waiting (for available thread) business processes.
View Heap Memory Level	Displays heap usage in the system. Business processes can run if heap space and CPU resources are available.
View Memory Generation	Displays JVM information specific to garbage collection and memory generation.
View System Information	Displays system level information from the JVM.
View VM Status	Displays Java Virtual Machine status.
View Manager Properties	Displays the list of properties from the noapp.properties file.
View Queue Threads	Displays a list of all queue threads.
View All Threads	Displays a list of all active threads.
View Stateful Adapters	Displays a list of stateful adapters running in the system. Stateful Adapters are adapters with an adapterType of STATEFUL, for example, the HTTP adapter.
View Stateless Adapters	Displays a list of stateless adapters running in the system. Stateless adapters are adapters with an adapterType of STATELESS, for example, the File System Adapter.
View Disabled Adapters	Displays a list of adapters that are currently marked as disabled (not running).
View DB Pool Information	Displays usage information for the configured DB pools.
View Cluster Multicast Data	Displays load data broadcast from the nodes when running in a cluster.
Config Queue	<p>Configure the queue parameters to tune performance. The parameters are not persisted and are reset when Sterling Integrator restarts.</p> <p>Note: The Config Queue, Reset Queue, and Step Monitor fields can only be used one at a time. To submit the data entered, you must click enter.</p>

Reset Queue	Resets the queue to default values. The parameters are not persisted and are reset when Sterling Integrator restarts. Note: The Config Queue, Reset Queue, and Step Monitor fields can only be used one at a time. To submit the data entered, you must click enter on your keyboard.
Step Monitor	Displays the list of business processes and workflow contexts in the queue. Note: The Config Queue, Reset Queue, and Step Monitor fields can only be used one at a time. To submit the data entered, you must click enter on your keyboard.
View Properties	Displays a list of all available property file names. Select a property from the list, then click Send .
View Common Properties	Displays a list of the named common property files. Select a property from the list, then click Send .

Enabling Queue Watcher

About this task

Queue Watcher allows you to enable the monitoring and management functionality from Sterling Integrator without having to restart the system for it to take affect.

To enable Queue Watcher without restarting Sterling Integrator:

Procedure

1. Access the Queue Watcher tool. See *Accessing Queue Watcher* for additional information.
2. Click **Enable Queue Watcher**. The page refreshes and displays the Queue Watcher page.

Disabling Queue Watcher

Queue Watcher allows you to disable the monitoring and management functionality from Sterling Integrator without having to restart the system for it to take affect.

About this task

To disable Queue Watcher without restarting Sterling Integrator:

Procedure

1. Access the Queue Watcher tool. For more information about accessing Queue Watcher, refer to the topic "Accessing Queue Watcher" on page 137.
2. Click **Disable Queue Watcher**. The Queue Watcher tool is disabled.

Compress Cache Contents

About this task

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

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

Procedure

1. Open the *install_dir*/properties/workflows.properties file using 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.

Important: Saving the workflows.properties file under a different name or in a different location will prevent the changes from being applied to Sterling Integrator.

4. Stop Sterling Integrator.
5. In the *install_dir*/bin directory, run one of the following commands:
 - (For UNIX or Linux) **setupfiles.sh**
 - (For Windows) **setupfiles.cmd**
6. Start Sterling Integrator. The changes are applied to Sterling Integrator and the cache contents are compressed before being written to disk and decompressed after being read from the disk.

Set the Document Body Serialization Threshold Property

About this task

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

If the business process requires the use of the document body at almost every step, the number of database reads may use more resources than serialization or deserialization. In such a scenario, 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, perform the following actions:

Procedure

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 the **tune.noapp.documentInlineSerializationThreshold** property in the *install_dir*/properties/tuning.properties file according to your document size.
3. From the **Administration** menu, select **Operations > System > Troubleshooter**.
4. In the System Troubleshooting page, click **Stop the System**. This stops Sterling Integrator.
5. From the *install_dir*/bin directory, run one of the following commands:
 - (For UNIX or Linux) **setupfiles.sh**
 - (For Windows) **setupfiles.cmd**

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

6. From the *install_dir*/bin directory, start Sterling Integrator by running one of the following commands:

- (UNIX or Linux) **run.sh**
- (Windows) **startWindowsService.cmd**

Sterling Integrator starts, and the updated performance configuration settings are applied.

Allocate JNDI Contexts

About this task

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

To change the JNDI allocation:

Procedure

1. From the **Administration** menu, select **Operations > System > Troubleshooter**.
2. In the System Troubleshooting page, click **Stop the System**. This runs the `softstop` script, allowing a fixed amount of time for the business processes to be completed before the system is stopped.
3. In the `install_dir/properties` directory, open the `tuning.properties` file using a text editor.
4. Locate the `tune.noapp.jndi.contextpoolsizes` property, and change the value to the level you require.
5. In the `install_dir/bin` directory, run one of the following commands:
 - (For UNIX or Linux) **setupfiles.sh**
 - (For Windows) **setupfiles.cmd**

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

6. In the `install_dir/bin` directory, start Sterling Integrator by running one of the following commands depending on your operating system:
 - (For UNIX or Linux) **run.sh**
 - (For Windows) **startWindowsService.cmd**

Sterling Integrator starts, and the updated performance configuration settings are applied.

Changing Persistence Levels

About this task

Persistence level is the level of detail written to the database even as the business process runs. Sterling Integrator can persist either all the data or only a small amount of data for each activity that is processed. The default value for Sterling Integrator is Full persistence. You may find that you have to change the persistence level to a lower level in order 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 both your processing and your system. You may have some business processes for which you have to persist all the steps and data, and other business processes for which you do not have to persist any of the steps or data.

There are three levels at which you can set persistence:

- Global – Change the value of the `persistence_level` property in the `install_dir/properties/noapp.properties_gis_ext.in` file. The default value is `SYSTEM DEFAULT`. You cannot set the value to `PERSISTENCE_WF_NONE` or `PERSISTENCE_ERROR_ONLY`.

Note: It is more efficient to set persistence levels at either the business process level or the activity level.

- Business Process – Change the persistence level during the check-in process. The business process persistence level overrides the global persistence level.
- Activity Level – Change the persistence level in the service or adapter. The activity persistence level overrides the business process persistence level and the global persistence level. However, changing the persistence level at the Activity level is rarely necessary.

Changing Persistence at the Global Level

To edit the global persistence levels:

1. Open the `install_dir/properties/noapp.properties_gis_ext.in` file using a text editor.
2. In the property file, locate the `persistence_level` property, and change the value to the persistence level you want.

Note: If the persistence level property is a comment in the `noapp.properties_gis_ext.in` file (`#persistence_level=value`), the persistence level is the default value of full (`PERSISTENCE_FULL`).

- `PERSISTENCE_FULL` – Saves a complete copy of process data along each step of the process.
- `PERSISTENCE_MINIMAL` – Saves all the steps in a business process and selected copies of process data.
- `PERSISTENCE_NONE` – Saves the first and last steps of a business process and any steps with an override persistence level.

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

3. Stop Sterling Integrator.
4. In the `install_dir/bin` directory, run one of the following commands:
 - (For UNIX or Linux) `setupfiles.sh`
 - (For Windows) `setupfiles.cmd`
5. Start Sterling Integrator.

Changing Persistence at the Business Process Level

To edit the persistence at the business process level:

Procedure

1. From the **Administration** menu, select **Business Processes > Manager**.
2. In the **Search** box, enter either the entire business process name for which you want to edit the persistence level or a part of the business process name, and click **Go!**
3. Click **source manager** adjacent to the business process you want to edit.
4. Click **edit** adjacent to the corresponding business process in order to lock the business process and access the Graphical Process Modeler.

5. In the BPML Specification page that is displayed, enter a new description that will help you identify this version of the business process.
6. Click **Next**.
7. In the Process Levels page, indicate the level of persistence you want to assign for the business process:

- Full - Retains all the data for this business process, including associated documents, activities, and all the process data.

Note: This choice uses the maximum amount of system resources compared to other persistence settings.

- Step Status (Engine May Override) - Retains the process data and documents for services configured to always persist data for a business process. Sterling Integrator may persist the additional data that is requested by the process engine or service parameters.

This setting is recommended for most processes because it helps optimize performance and provides a level of process recovery that is suitable for most business requirements.

- BP Start Stop (Engine May Override) - Retains additional data as requested by the process engine or service parameters.
- System Default - Assumes that Sterling Integrator is already configured 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. Sterling Integrator does not persist additional data.

Select this setting if you want to persist only the status information for each step, you are certain that your process will execute successfully without persisted data, and your recovery requirements do not depend on the step data being saved.

- BP Start Stop Only - Retains only error information. Select this setting when you are certain that your process will execute successfully without persisted data, and when your recovery requirements do not depend on the process step data being saved.

- Zero - Does not retain any data.

- Error Only - Retains process information only when an error occurs. If nothing is wrong and streamed document data is not used, nothing is retained.

- BP Start Stop Only (No Errors) - Retains information only on business process start and stop steps only. The first step of the business process is persisted in full and the last step is persisted in minimal. However, when the business process encounters an error and stops, the error step is not persisted and the business process stays in the ACTIVE state.

8. Click **Next**.
9. In the Deadline Settings page that is displayed, click **Next**.
10. In the Life Span page that is displayed, click **Next**.
11. In the Default Version page that is displayed, click **Next**.
12. In the Confirm page that is displayed, click **Finish**.

Changing System Logging

About this task

Sterling Integrator writes various types of messages to the log files on disk. The default is to log only the errors. For more information about system logs, refer to the topic Managing System Logs.

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

If your system logging level is set to ERROR, you may find that you require additional messages from more detailed logging.

For information about how to change system logging levels, refer to the topic Managing System Logs.

Enable Trusted Domains for Schemas

About this task

Sterling Integrator has document type definitions (DTDs) and schemas in a repository for XML components. When a DTD or schema is required, Sterling Integrator searches the database for the appropriate file. If Sterling Integrator cannot find the appropriate file, you can configure the *install_dir/properties/noapp.properties_platform_ifcesources_ext* file to have Sterling Integrator go to trusted Web sites and download the appropriate DTD or schema.

To configure Sterling Integrator to go to trusted Web sites for the DTD and schema data:

Procedure

1. In the *install_dir/properties* directory, locate the *noapp.properties_platform_ifcesources_ext.in* file.
2. Locate the *schemaResolver.trustedDomain.#* property and type the URL for the Web site. For example, the first trusted domain could be *schemaResolver.trustedDomain.1 = http://www.gdsregistry.org*. The second trusted domain could be *schemaResolver.trustedDomain.2=URL*, and so on. Sterling Integrator will search through these Web sites to find the appropriate DTD or schema.
3. Save the *noapp.properties_platform_ifcesources_ext.in* file under the same name and in the same location.
4. Stop Sterling Integrator.
5. In the *install_dir/bin* directory, run one of the following commands:
 - (For UNIX or Linux) **setupfiles.sh**
 - (For Windows) **setupfiles.cmd**
6. Start Sterling Integrator. The changes are applied and you can now download the DTD and schema data from the trusted Web sites.

Enable the Console Listener

About this task

You can enable the Console Listener to append the events identified in the Events Viewer to the *noapp.log* file, 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 will then have to view each component log file individually to see the events that are logged.

To enable the Console Listener:

Procedure

1. In the *install_dir*/properties directory, locate the listenerStartup.properties.in file.
2. Locate the following line:
`#Listener.Class.1=com.sterlingcommerce.woodstock.event.listeners.console.ConsoleListener`
3. Delete the # symbol in the first line to make the line active.
4. Save the listenerStartup.properties.in file under the same name and in the same location.
5. Stop Sterling Integrator.
6. In the *install_dir*/bin directory, run one the following commands:
 - (For UNIX or Linux) **setupfiles.sh**
 - (For Windows) **setupfiles.cmd**
7. Start Sterling Integrator. The changes are applied and you can now view the events as they are appended to the noapp.log file during processing.

Enable the JMX Event Listener

About this task

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

If you have an existing infrastructure for your enterprise applications to perform JMX-based management and event monitoring, you can establish a uniform way to handle event notifications from Sterling Integrator or Sterling Platform, along with your other applications, through your JMX framework, using the JMX console and tools.

To enable the JMX Event Listener:

Procedure

1. In the *install_dir*/properties directory, locate the jmx.properties file.
2. Set the jmx.start and jmx.eventListener.start properties to true. If necessary, restart the Sterling Commerce Platform to pick up these changes.
3. In the *install_dir*/properties directory, locate the noapp.properties file.
4. Set the startup.class4 property to `com.sterlingcommerce.woodstock.jmx.event.SPEventJMXAgentStartup` by either adding this line or editing the existing line. This starts the JMX Agent when the Sterling Commerce Platform starts. To view the MBeans configured in the Sterling Commerce Platform, go to `http://host:port`, where port is the HTMLAdapterServer port configured in the jmx.properties file.
5. Configure a business process to send an event at startup, completion, or at any other time.

6. In 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 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 the system-wide events. It is possible to listen to select events, rather than all of them. By adding a (comma-separated) list of event schema IDs to the jmx.eventListener.forward property, the JMX agent will restrict JMX notifications to only those events specified in the list.

Enhance EDI Performance

About this task

You can boost EDI performance in Sterling Integrator by adjusting the tracking level in the enveloping.properties file. The tracking level setting affects the following EDI functionalities:

- EDI Correlation Search
- EDI Document Tracking
- EDI Reporting

To modify the tracking level:

Procedure

1. In the install_dir/properties directory, locate the enveloping.properties.in file.
2. Locate the following line: TRACKING_LEVEL=*setting*
3. Change the value of the setting to achieve the required 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 Acknowledgment 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 Sterling Integrator.
6. In the install_dir/bin directory, run one of the following commands:
 - (For UNIX or Linux) **setupfiles.sh**
 - (For Windows) **setupfiles.cmd**
7. Start Sterling Integrator. The changes are applied to the enveloping.properties file, and the new settings will determine the default EDI performance and functionality of your business processes. The default EDI performance settings

that are 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. Refer to *Sterling Integrator Services and Adapters* documentation 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, `TRACKING_LEVEL` settings in both the `enveloping.properties` file and the EDI service parameter are overridden.

Customizing Property Files

About this task

The property files in Sterling Integrator help you set up, customize, and diagnose problems pertaining to the installation and operation of Sterling Integrator. You can maintain your customized property file settings with a customer override property file (`customer_overrides.properties`) that overrides the default property settings in the property files. The customer override property file is not changed during the installation of Sterling Integrator upgrades or patches.

To prevent your customized settings from being overwritten, use the customer override property file whenever possible rather than directly editing the Sterling Integrator property files or the associated `.in` files. For information about how to use the customer override property file to change property files, refer to the *Sterling Integrator Property Files* documentation.

CAUTION:

Property files directly affect the operation of Sterling Integrator. Contact Sterling Commerce Customer Support before changing any properties.

The following tasks use the `customer_overrides.properties` file:

- Changing Advanced File Transfer Settings
- Change the SoftStop Time Limit

Managing System Recovery

About this task

Sterling Integrator 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 on which Sterling Integrator is installed to a stable state, and then restart Sterling Integrator. After it is restarted, the system invokes the Recovery business process to search for and requeue any unprocessed transactions. If the database goes down, Sterling Integrator stops pulling jobs from queues. The system recovers these jobs when the database becomes available.

Sterling Integrator uses multithread operation to restart or resume business processes following a system crash. Unfinished business processes, up to a preconfigured number, are stored in a pool. Each of the preconfigured number of threads receive a batch of business processes to restart or resume. The size of the

batch is also preconfigured. 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, each of which resumes or restarts 10 business processes, and then requests and processes another 10, up to a total of 1000 business processes for all threads. You can configure the system settings to adjust recovery performance according to your business requirements.

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

Procedure

1. In the *install_dir*/properties directory, locate (or create, if necessary) the *customer_overrides.properties* file.
2. Open the *customer_overrides.properties* file using a text editor.
3. Add override statements as required, using the following format for each new line: `bprecovery.property=new_value` These statements will override the property settings in the *bprecovery.properties* file. The applicable properties are described in the following table:

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 have 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 the <i>batchSize</i> property).</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 the <i>batchSize</i> property. 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 have 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 <i>numberOfThreads</i> 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 on your business requirements.</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 on your business requirements.</p>

Let us 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 Sterling Integrator and restart it to use the new values. For assistance in determining the proper settings for your business requirements, contact Sterling Commerce Customer Support.

Change the SoftStop Time Limit

About this task

By default, the time limit for a softstop of Sterling Integrator is 5 minutes in order to give the business processes that are running the time to finish. At the end of the time limit, the system performs a hardstop that immediately shuts down Sterling Integrator. You can change the time limit for a softstop by overriding the default value.

To override the default softstop time limit:

Procedure

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

Note: This line overrides the `shutdown_timeout` property in the `noapp.properties_platform_ifcesources_ext` file.

In this line, *new_value* refers to the new time limit, in seconds, for a softstop of Sterling Integrator. 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 Sterling Integrator and restart it to use the new values.

Changing Advanced File Transfer Settings

About this task

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 (`persistent_batching_interval`)
The AFT Worker thread stores the AFT events in the database. It consumes all the current events in an input event queue, and then sleeps for the duration of 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.
- The capacity of the AFT database listener event queue (`event_input_queue_capacity`)

Note: For large volume operations of FTP, Mailbox, Advanced File Transfer, and Sterling File Gateway, the number of open file descriptors should be set to at least 4096, using the command `ulimit -n 4096`.

You can specify the size of the event input queue. By default, the AFT backend will store up to 2048 AFT events while waiting for the AFT Database worker

thread to return and process them. If the thread does not return before Sterling Integrator 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, such as copying a file, until the AFT Database worker thread empties the queue. If the AFT Database worker thread returns and the event queue is filled 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 sleep operations and service the event queue. The Worker thread will also monitor the number of events it is retrieving from the input queue. When the queue is no longer filled to capacity, 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 the AFT Visibility performance.

- The amount of AFT coverage to enable visibility_coverage.

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 the AFT Visibility components. Specifying the proper settings according to your requirements allows you to disable unnecessary components and improve performance, while preserving core file transfer and data flow capabilities.

To specify AFT Visibility settings, perform the following steps:

Procedure

1. Navigate to the *install_dir*/properties directory and locate (or create, if necessary) the customer_overrides.properties file. For more information about customizing property files, refer to the topic Customizing Property Files.
2. Open the customer_overrides.properties file using a text editor.
3. Add the property settings that you want to use in the following format, using the information provided in the following table as reference:

```
dmivisibility.PROPERTY_NAME=PROPERTY_VALUE
```

Here, PROPERTY_NAME is the name of the property.

PROPERTY_VALUE is 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:

```
dmivisibility.persistent_batching_interval=3000
```

Property	Description
persistent_batching_interval	Specifies the time interval, in milliseconds, between batching AFT events to the database. Default is 2000 milliseconds (2 seconds). Example: dmivisibility.persistent_batching_interval=2000

Property	Description
event_input_queue_capacity	<p>Specifies the capacity of the bounded queue (AFT database listener event queue) through which all the AFT events must pass en route the database. Default is 2048. Example: <code>dmvisibility.event_input_queue_capacity=2048</code></p> <p>Notes:</p> <ul style="list-style-type: none"> • If the value of this property is set too high, it may negatively impact system performance because of excessive memory usage. This property should be tuned with the <code>persistent_batching_interval</code> property for optimum results. • For large volume operations of FTP, Mailbox, Advanced File Transfer, and Sterling File Gateway, the number of open file descriptors should be set to at least 4096, using the command <code>ulimit -n 4096</code>.
visibility_coverage	<p>Specifies which AFT Visibility components are enabled. Valid values are:</p> <ul style="list-style-type: none"> • All – Enables all the AFT Visibility components. This includes CommBase, Authentication, Authorization, NonFileXfer, and AdminAudit. • CommBase – Enables 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 – Enables Authentication events and records. • Authorization – Enables Authorization events and records. • NonFileXfer – Enables Non File Transfer events and records (such as FTP CWD). • CommAll – Enables CommBase, Authorization, Authentication, and NonFileXfer. • AdminAudit – Enables Administrative Audit trail events and records. • None – Disables AFT Visibility. <p>Example: <code>dmvisibility.visibility_coverage=All</code></p> <p>Visibility coverage property settings can be also combined using commas as delimiters, for example:</p> <p><code>dmvisibility.visibility_coverage=CommBase,Authentication,AdminAudit</code></p>

4. Save and close the `customer_overrides.properties` file.
5. Stop Sterling Integrator and restart it to use the new values.

Optimize System Performance for Sterling e-Invoicing

Sterling e-Invoicing is installed on an instance of Sterling Integrator, and shares many of the resources with the latter.

You should, therefore, tune your Sterling Integrator installation first, and then perform the Sterling e-Invoicing-specific tuning task. However, be aware that the changes you make to Sterling e-Invoicing can also affect the performance of Sterling Integrator.

Note: Do not edit the property files directly. Make the changes in the `customer_overrides.properties` file.

To optimize the performance of Sterling e-Invoicing:

1. From the **Administration** menu, select **Business Processes > Manager**.
2. Enter `TrustWeaverEInvoicing` in the **Search** box and click **Go!**
3. Click **source manager**.
4. Click **edit** to lock the business process.
5. In the BPML Specification page that is displayed, enter a new description that will help you identify this version of the business process.

6. Click **Next**.
7. In the Process Levels page that is displayed, in the **Set Queue** field, select a single workflow queue to dedicate solely to the TrustWeaverEInvoicing process, for example, queue 5.
8. Click **Next**.
9. In the Deadline Settings page that is displayed, click **Next**.
10. In the Life Span page that is displayed, click **Next**.
11. In the Default Version page that is displayed, select the version being edited and click **Next**.
12. In the Confirm page that is displayed, click **Finish**.
13. Stop Sterling Integrator.
14. Modify the tuning.properties file so that the selected queue's minimum and maximum pool sizes are equal to the number of simultaneous connections allowed to TrustWeaver. For example, if you are using queue 5, and are using up to 10 simultaneous connections to TrustWeaver, you must set the following properties in the tuning.properties property file:
 - NOAPP.MAX_POOL_SIZE_5=10
 - NOAPP.MIN_POOL_SIZE_5=10
15. After you edit the property file, run **setupfiles.sh** (**setupfiles.cmd** for Windows).
16. Restart Sterling Integrator to apply the change.

The Order of Segments in Maps and Data

The benefit of correctly constructed maps is faster processing of input data. A correctly constructed map contains segments on the input (left) side of the map in the same order in which those segments occur in the input data.

In order to locate a matching segment identifier, a search is conducted in the map for each input record in the data stream, starting with the last segment entry plus one. The search is sequential, and wraps to the beginning of the map if the target segment is not found when the search reaches the end of the map. If the target segment identifier is not found even after the entire map has been searched, an error is generated, and the processing continues with the next record in the data stream. You can minimize the processing time by organizing the map segments in the same order as that of the input document segments.

The following scenario demonstrates a situation where an entire map is searched to match each input segment stream.

The following table lists the order of the identifiers in the map:

Segment Identifier
00001
00002
00003
00004

The following table lists the order of the input data segments in the document:

Segment Identifier
00004
00003
00002
00001

Let us assume that a translator searches three records in the map before finding the fourth map record matching the first segment in the input data stream. Similarly, to process the second segment in the input data stream, the translator begins the search after the fourth segment, wraps to the beginning of the file, and searches three more segments, before finding the correct match. The same search sequence is followed until all the segments are processed in the input stream. If map segments are missing in an input document, it results in lower throughput.

Purge Performance Recommendations

This topic provides information about purge performance recommendations. Purge process removes data from the live database tables.

The following table provides information about purge performance recommendations:

Recommendation	Comments
Purge business processes	<p>The Schedule_IndexBusinessProcessService runs every 10 minutes. By default, it is configured to index 5,000 business processes every time it runs, which translates to 720,000 business processes in a 24-hours time span.</p> <p>If your system is handling a high volume of business processes, you can increase the number of business processes indexed per cycle. You can modify the Schedule_IndexBusinessProcessService business process by modifying variable as shown in the following example:</p> <pre><assign to="max_business_processes">500000</assign></pre>
Batch purge	<p>The Schedule_PurgeService scans the ARCHIVE_INFO database table to purge data in the system. By default, it purges 2,000 rows of data in batches to limit the database transaction size. You can set the batch size in the jdbc_purge.properties file that can be located in <i>install_dir/install/properties</i> (<i>install_dir\install\properties</i> for Windows) directory. Modify the batch size by changing the variable as shown in the following example:</p> <pre>query.purge_archiveInfoRowLimit.<i>database_name</i></pre> <p>Note: Replace the <i>database_name</i> with the database you are using along with Sterling Integrator, for example, mssql.</p> <p>If you modify the value to increase the batch purge rate, ensure that you configure your database with sufficient log space to handle large transactions.</p>

Recommendation	Comments
Document lifespan	<p>By default, lifespan for information persisted by a business process is two days. The Schedule_PurgeService purges data in competition with business activity on a system with predictable processing periods. You can defer document lifespan by increasing the business process lifespan. For example, consider that the system is processing information from 8 a.m. to 5 p.m., you can increase the document lifespan by 12 hours, the purge activity will occur between 8 p.m. on the same day and 5 a.m. on the next day.</p> <p>You can modify the business process lifespan by logging in as a Sterling Integrator administrator. Navigate to Operations > Archive Manager > Configure Archive Settings and perform the changes.</p>
Database compression	<p>The Schedule_IndexBusinessProcessService and the Schedule_PurgeService runs complex queries which results in significant disk I/O on the database server. You can compress the database tables provided in the following list to achieve significant reduction (up to 50% in some cases) in disk usage:</p> <ul style="list-style-type: none"> • CORRELATION_SET • WORKFLOW_CONTEXT • TRANS_DATA • DOCUMENT • ARCHIVE_INFO

Optimizing and Controlling the System Threads

This topic provides information about optimizing and controlling the system threads. It also provides information about threads created, their source, and procedures to control them.

Optimizing System Threads

Out of memory situations are very difficult to diagnose. Sterling Integrator creates around 300 threads that can be grouped under system threads, adapter threads, common JVM threads, third party software threads, and several other threads that occur only once for different purposes.

The following table lists the threads created in Sterling Integrator and their source.

Thread created by	Thread name	Count
ActiveMQ	ActiveMQ transport	80
	ActiveMQ Session Task	11
Jetty	SessionScavenger	47
	ConduitStreamListener	16
	SocketListener	10
JGroup	Various Jgroup Handlers	22
JetSpeed	RunnableThread	10
Perimeter PS Dispatcher	Various Adapters	13

Thread created by	Thread name	Count
Business process queues	ReschedulingThread	10
B2B	B2B http Servlet Thread	3
	FIFOEventListener	10
	QueueThread:queue	11
System	RMI	7
	Timer	7
Others	From various components	30
Total		287

Controlling the Threads

Several threads created by Sterling Integrator for various purposes may not be required always and they can be controlled wherever required. This will enhance the Sterling Integrator's performance considerably.

Following are the concepts described in this topic:

- ActiveMQ Threads
- Jetty Threads
- JGroup Threads
- JetSpeed Threads
- Adapter Threads
- Business Process Queue Threads
- FIFOEventListener and Queue Threads
- RMI Threads
- Timer Threads

ActiveMQ Threads

ActiveMQ threads can be controlled by running ActiveMQ broker in a separate JVM. No additional setup or configuration is necessary to run ActiveMQ in a separate JVM. Sterling Integrator build installation process configures the system to use it out of the box for both cluster ActiveMQ and non-cluster ActiveMQ.

However, if you plan to use clustering, you may choose a different configuration by editing the `activemqconfig.xml` file. Before editing this file, read the `readme_cluster.txt` file. It contains information about how to use the options in the `activemqconfig.xml` file. Both files are located in the `install_dir/install/activemq/conf` folder.

Mandatory Startup for ActiveMQ

The `startActivemqMandatory` parameter in the `install_dir/install/properties/activeMQ.properties` file controls the remaining processes and starts them if ActiveMQ fails to start. The default value for this parameter is `false`. To change this, you can create an extension file (for example, `activeMQ.properties_clumpName_ext.in` or `customer_overrides.properties` file) and specify the following entry:

```
startActivemqMandatory=true|false
```

Where:

true = If activemq fails to start, the rest of processes will not be started.

false = If activemq fails to start, continue to start the rest of the processes.

Standalone ActiveMQ Commands

You can start and stop standalone ActiveMQ server by running the following commands.

To start the standalone ActiveMQ server, ensure that ActiveMQ dynamic configuration file (`activemqconfig.xml.in`) and ActiveMQ configuration XML file (`activemqconfig.xml`) are present in the `install_dir/install/activemq/conf` directory. Run the following command from `install_dir/install/bin` directory:

- For UNIX, run **startActiveMQ.sh**
- For Windows, run **startActiveMQWindowsService.cmd**

To stop the standalone ActiveMQ server, run the following command from `<install_dir>/install/bin` directory:

- For UNIX, run **stopActiveMQ.sh**
- For Windows, run **stopActiveMQWindowsService.cmd**

Note: You can also start or stop ActiveMQ service from Windows Service Manager.

Using an External ActiveMQ Environment

ActiveMQ is bundled along with Sterling Integrator. However, you can use a different ActiveMQ environment by modifying certain files.

Note: It is recommended that users who are familiar with ActiveMQ environment perform this task.

To use an external ActiveMQ environment in UNIX:

1. Shut down Sterling Integrator.
2. Change `ACTIVEMQ_PORT` in `sandbox.cfg` and point to your own ActiveMQ environment.
3. Remove **startActiveMQ.sh** from `install/bin/run.sh.in`.
4. Remove **stopActiveMQ.sh** from `install/bin/hardstop.sh.in`.
5. Change `remote.protocol_config=client connection` in the `install/event.properties.in` file to your ActiveMQ environment.
6. Run **install/bin/setupfiles.sh**.
7. Restart Sterling Integrator.

To use an external ActiveMQ environment in Windows:

1. Shut down Sterling Integrator.
2. Change `ACTIVEMQ_PORT` in `sandbox.cfg` and point to your own ActiveMQ environment.
3. Remove **"net start "%ACTIVEMQ_SERVICE_NAME%" >NUL"** from `install/bin/startWindowsService.cmd`.
4. Remove **"net stop /y "%ACTIVEMQ_SERVICE_NAME%"** from `install/bin/stopWindowsService.cmd`.

5. Change **remote.protocol_config=client connection** in the `install/event.properties.in` file to your ActiveMQ environment.
6. Run **install/bin/setupfiles.cmd**.
7. Restart Sterling Integrator.

Changing the Cluster Setting for Bundled ActiveMQ

The configuration file for the bundled ActiveMQ is `install/activemq/conf/activemqconfig.xml`. You can manually change the broker setting to fit your business requirements. You can also extend this file with `activemqconfig_clumpname_ext.xml` to configure your own beans.

Note: Read `install/activemq/conf/readme_cluster.txt` file before making any changes.

Jetty Threads

Sterling Integrator uses Jetty version 4.2.24. Jetty version 4.2.24 when compared to latest versions like Jetty version 6.1.8 offers limited control on the number of threads created. However, you can control the numbers of threads created by Jetty listeners. Further, the large numbers of SessionScavenger and ConduitStreamListener threads are not controlled by listener thread parameters. They are created for web applications and HTTP Servlet adapters.

You can control the number of threads created by Jetty Listeners by modifying the following configuration parameters in `noapp.properties` file:

```
# specify the minimum number of threads for Socket Listeners
for Jetty jetty_min_threads = 5
# specify the maximum number of threads for Socket Listeners
for Jetty jetty_max_threads = 100
```

Note: You cannot modify the `jetty_min_threads` value. However, you can modify the `jetty_max_threads` value in the available range from 5 - 100.

JGroup Threads

JGroup is a reliable multicast communication toolkit and is used in Sterling Integrator cluster environment. You cannot control the number of threads created by JGroup.

JetSpeed Threads

Jetspeed is the portal engine used in Sterling Integrator dashboard interface. The `jetspeedresources.properties` file controls the number of threads created by JetSpeed.

You can control the number of threads created by JetSpeed by modifying the following configuration parameters in `install/noapp/deploy/dashboard/webapp/WEB-INF/conf/JetspeedResources.properties` file.

```
#Specify the initial number of threads to create
services.ThreadPool.init.count=5
#Specify the maximum number of threads to create
services.ThreadPool.max.count=20
#Specify the minimum number of threads to keep as
spare until you hit the maximum
services.ThreadPool.minspare.count=5
```

Note: You cannot modify the `services.ThreadPool.init.count` value. However, you can modify the `services.ThreadPool.max.count` value in the available range from 5 - 20.

After modifying, you should remove the `install/noapp/deploy/dashboard/webapp/WEB-INF/conf/JetspeedResources.properties` from `install/noapp/deploy/dashboard.war` file to make your change take effect.

Adapter Threads

Several Jetty and Timer threads are created by adapters. You can disable the adapters that are not required to run your business processes thereby controlling the number of threads created by the adapters.

The following adapters can be disabled to reduce the number of threads created:

Note: Disabling an adapter in the following list can reduce at least one or two threads in most cases.

- FIFO Routing
- FIFO Error Queue Listener
- HTTP Communications Adapter
- B2B HTTP Communications Adapter
- SFTP Client Adapter
- FTP Client Adapter
- Map Test Http Server
- ebXML Http Server Adapter
- MBI Http Server Adapter
- SOA Http Server Adapter
- SOA SSL Http Server Adapter
- RN Http Server Adapter
- Http Server Adapter
- SWIFTNet Http Server Adapter

Business Process Queue Threads

Sterling Integrator creates nine regular business process queues and one internal queue called wait queue for wait service. You cannot control the number of threads created for business processes.

FIFOTaskListener and Queue Threads

The FIFORouting adapter creates and controls ten queues for FIFO processing. Each FIFO queue creates a FIFO task listener and every task listener creates a consumer at startup. You can configure the number of queues to reduce the number of threads. Additionally, you can disable the FIFORouting adapter if you are not using it thereby turning off all the queues created by the adapter.

You can control the number of threads by modifying the following configuration. The number of queues configured depends on the system load.

```
#In customer_overrides.properties, additional queues can be added
by adding, for example:
#fifo.workflow.taskqueue.11=FIFO.GIS.QUEUE.11
#fifo.workflow.taskqueue.12=FIFO.GIS.QUEUE.11
```

```
#Note, queues cannot be reduced in customer_overrides.properties
but the names can be changed and must be unique
workflow.taskqueue.1=FIFO.GIS.QUEUE.1
workflow.taskqueue.2=FIFO.GIS.QUEUE.2
workflow.taskqueue.3=FIFO.GIS.QUEUE.3
workflow.taskqueue.4=FIFO.GIS.QUEUE.4
workflow.taskqueue.5=FIFO.GIS.QUEUE.5
workflow.taskqueue.6=FIFO.GIS.QUEUE.6
workflow.taskqueue.7=FIFO.GIS.QUEUE.7
workflow.taskqueue.8=FIFO.GIS.QUEUE.8
workflow.taskqueue.9=FIFO.GIS.QUEUE.9
workflow.taskqueue.10=FIFO.GIS.QUEUE.10
```

RMI Threads

The RMI threads are system generated threads for JNDI. You cannot control the number of RMI threads.

Timer Threads

The timer threads are created when Sterling Integrator starts. It is not recommended to control these threads as they are necessary for Sterling Integrator to run smoothly.

The following timer threads are created when Sterling Integrator starts:

- Check Sterling Integrator component licenses and generate messages for users when one or more licenses is about to expire.
- Roll the log service files.
- Gather YCP statistics used by the entity framework.
- Monitor resources and detect database connections or database connection leaks.
- Schedule business processes.
- JNDI service timer.
- ActiveMQ timer.

Chapter 7. Performance Statistics

Manage Performance Statistics

The Performance Statistics Report is the key to managing your Sterling Integrator's performance.

Before you can generate or view a performance statistics report, you must turn on performance statistics. For more information about how to turn it on, refer to the topic *Turning On and Turning Off Performance Statistics*. You can turn on and turn off performance statistics as many times as you want to. When performance statistics is turned off, the statistics are simply discarded, and when performance statistics is turned on, the current report resumes the process of collecting statistics. If you do not generate a report before you turn on performance statistics, the default report is generated automatically, and collects the performance statistics.

Run the Performance Statistics Report to obtain benchmarks directly after you perform each of these tasks:

- Install Sterling Integrator.
- Tune Sterling Integrator using the information provided in the “Performance Tuning Utility” on page 121 topic.
- Complete manual tuning, if any, to the *install_dir/properties/noapp.properties* file and any extension files (ext).
- Make improvements to your business processes, if necessary, including changing the persistence levels.

After you receive your initial benchmark Performance Statistics Report, generate a new report every few weeks or at least once a month to have a snapshot of the way your Sterling Integrator 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 troubleshoot and resolve issues before they cause 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 generate a new report, the new report becomes the current report.

Viewing a Performance Statistics Report

To view a Performance Statistics Report:

1. From the **Administration** menu, select **Operations > System > Performance > Statistics**.
2. In the Performance Statistics Report Manager page, under **View Report**, select a name from the list, and click **Go!**.

Reading a Performance Statistics Report

Performance Statistics Report provides you with information that can help you maintain your Sterling Integrator system, and reduce performance issues before they become problems. Performance Statistics Report helps you identify and analyze bottlenecks in your business processes and the time taken for the completion of each activity.

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

Performance Statistics Report consists of the following sections:

- Business Process Statistics
- Internal System Statistics

Business Process Statistics

The following figure displays the Business Process Statistics section of a 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)
CIIDeveloper	19183	19275	19229.0	2	38458	Invocations per hour: 371.583	
--> 1 DeenvloperCII	18474	18872	18573.0	2	37146	152.0	176.0
--> 2 DecisionEngineService	0	0	0.0	2	0	5.0	100.0
ED Developer	16142	16142	16142.0	1	16142	Invocations per hour: 237.743	
--> 1 ECIIDeveloper	0	14859	14859.0	1	14859	0.0	50.0
Sched.Job_nce:BusinessProcessService	733	1732	1144.0	7	8008	Invocations per hour: 2.573	
--> 1 AssignService	0	1	0.0	7	3	1.0	7.0
--> 2 StoreLockService	16	82	33.0	7	235	0.0	30.0
--> 3 AssignService	4	24	10.0	7	70	1.0	5.0
--> 4 nce:BusinessProcessService	249	1034	838.0	7	5867	0.0	14.0
--> 5 AssignLockService	13	26	20.0	7	151	0.0	12.0
Sched.Job_PurgeService	447	1226	809.0	7	5667	Invocations per hour: 5.380	
--> 1 AssignService	1	17	11.0	7	77	1.0	0.0
--> 2 StoreLockService	17	126	63.0	7	373	1.0	12.0
--> 3 AssignService	1	24	11.0	7	111	1.0	6.0

The following table describes the columns in the Business Process Statistics section:

Note: Business process statistics are more varied than what is described in the Description column of the following table.

Statistic Name	Description
Business Process Statistics	
Name	<p>Name of a business process, followed by numbered rows pertaining in the individual activities in the business process. The number of each row indicates the corresponding 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 DeenvloperCII • --> 2 DecisionEngineService

Statistic Name	Description
Min (ms)	<p>Minimum time (in milliseconds) it took for a business process or activity to complete. If there are multiple invocations, this is the lowest minimum time taken for all the invocations.</p> <p>For example, in the figure preceding this table, it can be seen that the CIIDevelope business process took 19183 ms to complete, while the Deenveloping step (1 DeenveloperCII) took a minimum time of 18474 ms. Both 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 down. If you notice this statistic increasing over your benchmark value, it may be indicative of a performance issue. If you see an occasional variance for a business process or activity, it does not necessarily indicate a performance issue. If, however, you notice a continuous variance between production statistics and your benchmarks, you probably have a real issue that should be addressed.</p> <p>For more information about slow systems, refer to the topic A Slow System: Symptoms, Causes, and Resolution.</p> <p>For more information about improving business process execution time, refer to the topic Symptoms and Causes of Poor Business Process Execution Time.</p>
Max (ms)	<p>Maximum time (in milliseconds) it took for a business process or activity to complete. If there are multiple invocations, this is the highest maximum time taken for all the invocations.</p> <p>For example, in the figure preceding this table, it can be seen that the CIIDevelope business process took 19275 ms to complete, while the Deenveloping step (1 DeenveloperCII) took a maximum time of 18672 ms. Both 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 down during processing. If you notice this statistic increasing over your benchmark value, it may be indicative of a performance issue. If you see an occasional variance for a business or activity, it does not indicate a performance issue. If, however, you notice a continuous variance between production statistics and your benchmarks, you must resolve the issue.</p> <p>For more information about slow systems, refer to the topic A Slow System: Symptoms, Causes, and Resolution.</p> <p>For more information about improving business process execution time, refer to the topic Symptoms and Causes of Poor Business Process Execution Time.</p>

Statistic Name	Description
Average (ms)	<p>Average processing time taken for a business process or service to be completed.</p> <p>For example, in the figure preceding this table, it can be seen that 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 these are the average statistics; there were two invocations of the business process.</p> <p>Use this statistic to track the overall averages of business processes or activities that begin to slow down during processing. If you notice this statistic increasing over your benchmark value, it may be indicative of a performance issue. If you see an occasional variance for a business or activity, it does not necessarily indicate a performance issue. If, however, you notice a continuous variance between production statistics and your benchmarks, you probably have a real issue that should be addressed.</p> <p>For more information about slow systems, refer to the topic A Slow System: Symptoms, Causes, and Resolution.</p> <p>For more information about improving business process execution time, refer to the topic Symptoms and Causes of Poor Business Process Execution Time.</p>
Invocations	<p>Number of times the business process or service was invoked since the report was generated first. The number of invocations should be the same for both the business process and each activity in the business process, unless you have a business process that loops several times. In such a situation, you can set the number of invocations lower than the number of activities in the business process.</p> <p>For example, in the figure preceding this table, it can be seen that the CIIDevelope business process had two invocations, and the Deenveloping step (1 DeenveloperCII) had two invocations.</p> <p>Use this statistic to determine if the number of business processes expected are running, and if all the activities in the business processes are running. If you see a variance between the business process invocation number and the activity number, this may indicate that a business process has an error, and is in a waiting, interrupted, or halted state. You can also use this statistic to determine the processing load on your system during different processing periods, such as peak and nonpeak processing hours.</p>

Statistic Name	Description
Total (ms)	<p>Total time for (in milliseconds) a business process or service has taken to process since the time you last generated the report. This is the sum of total time taken for the completion of all the business processes.</p> <p>For example, in the figure preceding this table, it can be seen that the CIIDevelope business process took 38458 ms to process since the report was last generated and the Deenveloping step (1 DeenveloperCII) took a total of 37146 ms.</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, it indicates that the processing time is longer. If the processing time is longer than your benchmark, you must resolve the issue.</p> <p>Note: Because statistics are accumulated from the beginning of monitoring until its end, you might have to turn monitoring on and off and collect the statistics at regular intervals. This allows you to monitor performance at certain times during processing. A change in the total work a business process performs during the intervals might indicate a problem, but only if the demand for that business process is constant throughout the measurement period.</p> <p>For more information about improving business process execution time, refer to the topic Understanding Business Process.</p>
Pre-Service (ms)	<p>Total time taken by the business process engine to complete activities before a business process or service is run.</p> <p>For example, in the figure preceding this table, it can be seen that the DeenvelopeCII activity took 37146 ms for preservice activities. Every time an activity is run, Sterling Integrator must complete system service invocations before running the activity.</p> <p>Use this statistic to see if the system services are taking a long time to run before a business process is run. If you see a variance between the report time and the benchmark time, this may indicate that there is a database, operating system, or some other system issue that must be resolved.</p> <p>For more information about operating system issues, refer to your vendor's documentation.</p> <p>For more information about slow systems, refer to the topic A Slow System: Symptoms, Causes, and Resolution.</p>

Statistic Name	Description
Post-Service (ms)	<p>Total time taken by the business process engine to complete activities after the business process or service is run.</p> <p>For example, in the figure preceding this table, it can be seen that the DecisionEngineService activity took 100 ms for postservice activities. Every time an activity runs, Sterling Integrator completes some persistence of information to the database. This persistence takes time to complete, and the postservice time includes this persistence time.</p> <p>Use this statistic to see if the system services are taking a long time to run after a business process is run. If you see a variance between the report time and the benchmark time, this may indicate that there is a database, operating system, or some other system issue that must be resolved.</p> <p>For more information about operating system issues, refer to your vendor's documentation.</p> <p>For more information about slow systems, refer to the topic A Slow System: Symptoms, Causes, and Resolution.</p>

Internal System Statistics

The Internal System Statistics section of the Performance Statistics Report includes a lot of information that you can use to monitor your Sterling Integrator system and ensure that it meets the benchmarks that you have established.

The following figure displays the Internal System Statistics section of a Performance Statistics Report:

Internal System Statistics					
Name	Min (ms)	Max (ms)	Average (ms)	Frequency	Total (ms)
Persist a Document to the Database	9	212	111.0	2	223
Persist a Business Process Step to the Database	5	391	59.0	162	9556
Assign from an XPath Statement	2	43	10.0	29	291
Execute an XPath Statement	1	43	5.0	111	556
Acquire a Database Connection	1	1080	3.0	18702	56106
Put a Business Process on the Queue	1	49	3.0	38	45

The following table describes the columns in the Internal System Statistics section:

Statistic Name	Description
Internal System Statistics	
Name	<p>Name of the system activity, including the following:</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

Statistic Name	Description
Min (ms)	<p>Minimum time (in milliseconds) the system activity took to complete. If there are multiple invocations, this is the lowest time for all the invocations.</p> <p>For example, in the figure preceding this table, the Persist a Document to the Database activity took 9 ms to complete. This is the shortest time of the nine invocations of the activity.</p>
Max (ms)	<p>Maximum time (in milliseconds) the system activity took to complete. If there are multiple invocations, this is the highest time for all the invocations.</p> <p>For example, in the figure preceding this table, the Persist a Document to the Database activity took 212 ms to complete. This is the longest time of the nine invocations of the activity.</p>
Average (ms)	<p>Average processing time taken for the system activity to complete. This is the average time for all invocations.</p> <p>For example, in the figure preceding this table, the Persist a Document to the Database activity took an average of 111 ms to complete. This is the average processing time taken by the nine invocations of the activity.</p>
Invocations	<p>Number of times the system activity was invoked since you last generated the report.</p> <p>For example, in the figure preceding this table, the Persist a Document to the Database activity was invoked nine times.</p>
Total (ms)	<p>Total time (in milliseconds) the system activity has taken to process since the report generation began. This is the sum total time of all the system activity completion times.</p> <p>For example, in the figure preceding this table, the Persist a Document to the Database activity took a total of 1003 ms to complete.</p>

The following table describes each internal system activity, and how to use the statistics to reduce the chances of a performance issue occurring:

Statistic Name	Description
Persist a Document to the Database	<p>The length of time and the number of times a document was persisted to the database.</p> <p>For more information about persistence, refer to the topic Changing Persistence Levels.</p> <p>For example, in the figure preceding this table, 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>The Persist a Document to the Database statistic is helpful in determining database issues and slow processing issues. When you compare this information with your benchmarks, verify if you notice the following:</p> <ul style="list-style-type: none"> • The Min (ms), Max (ms), and Average (ms) times are increasing. This indicates that the database is becoming full, or that you have a connection leak. • The number of invocations are increasing. This indicates that your persistence level is set too high. • The number of invocations is small and the Min (ms) and Max (ms) times are increasing. This indicates that you are persisting large documents to the database, which can be moved to the disk to save database space.

Statistic Name	Description
Persist a Business Process Step to the Database	<p>The length of time and the number of times a business process step was persisted to the database.</p> <p>For more information about persistence, refer to the topic Changing Persistence Levels.</p> <p>For example, in the figure preceding this table, 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>The Persist a Business Process Step to the Database statistic is helpful in determining database issues and slow processing issues. When you compare this information with your benchmarks, verify if you notice the following:</p> <ul style="list-style-type: none"> • The Min (ms), Max (ms), and Average (ms) times are increasing. This indicates that either the database is becoming full, or you have a connection leak. • The number of invocations are increasing. This indicates that your persistence level is set too high. <p>For more information about slow systems, refer to the topic Changing Persistence Levels.</p> <ul style="list-style-type: none"> • The number of invocations is small and the Min (ms) and Max (ms) times are increasing. This indicates that you are persisting large documents to the database, which can be moved to the disk to save database space.
Assign from an XPath Statement	<p>The length of time and the number of times an assign activity was completed from an XPath statement in a business process.</p> <p>For example, in the figure preceding this table, 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>The Assign from an XPath Statement statistic is helpful in determining if you have well-structured XPath statements in your business processes. When you compare this information with your benchmarks, verify if you notice the following:</p> <ul style="list-style-type: none"> • The Min (ms), Max (ms), and Average (ms) times are increasing. This indicates that the XPath statement is not written efficiently, and may slow down your process time. 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, because this causes the entire process data to be traversed.

Statistic Name	Description
Execute an XPath Statement	<p>The length of time and the number of times an XPath statement ran in a business process.</p> <p>For example, in the figure preceding this table, Assign from an XPath Statement had the following statistics:</p> <ul style="list-style-type: none"> • Min (ms) – 1 • Max (ms) – 43 • Average (ms) – 5.0 • Invocations – 111 • Total (ms) – 598 <p>This statistic is helpful in determining if you have well-structured XPath statements in your business processes. When you compare this information with your benchmarks, verify if you notice the following:</p> <ul style="list-style-type: none"> • The Min (ms), Max (ms), and Average (ms) times are increasing. This indicates that the XPath statement is not written efficiently, and may slow down your process time. 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, because this causes the entire process data to be traversed.
Acquire a Database Connection	<p>The length of time and the number of times a database connection was made.</p> <p>For example, in the figure preceding this table, 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>The Acquire a Database Connection statistic is helpful in determining whether you have database issues, resource leaks, or whether you have to increase the number of database pools. When you compare this information with your benchmarks, verify if you notice the following:</p> <ul style="list-style-type: none"> • The Min (ms), Max (ms), and Average (ms) times are increasing. This indicates that the database is quite active if it is used by other applications outside of Sterling Integrator. • The number of invocations are increasing. This indicates that your persistence level is set too high, causing too much data to be stored in the database, or your cache levels are not tuned correctly, causing the data to be stored and retrieved from the database instead of from the cache. • The number of invocations are low and not increasing. This indicates that either 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.

Statistic Name	Description
Put a Business Process on the Queue	<p>The length of time and the number of times a business process was placed in a queue.</p> <p>For example, in the figure preceding this table, 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>The Put a Business Process on the Queue statistic is helpful in determining if you are using a queue instead of your cache during processing. When you compare this information with your benchmarks, verify if you notice increasing invocation times. If yes, it indicates that you are not using your cache efficiently.</p>

Turning On and Turning Off Performance Statistics

About this task

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, you may want to turn on the performance statistics when you are testing or running your weekly or monthly report in order to compare your benchmarks, and turn off performance statistics during the day-to-day operations in order to save your system resources for processing.

To turn on performance statistics:

1. From the **Administration** menu, select **Operations > System > Performance > Statistics**.
2. In the Performance Statistic Report Manager page, under **On/Off**, select the check box adjacent **Enable Performance Statistics**. Click **OK** at the confirmation message.

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

To turn off performance statistics:

Procedure

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

Reporting Performance Statistics

About this task

Reporting performance statistics includes the following tasks:

- Creating a New Performance Statistics Report
- Deleting a Performance Statistics Report

Creating a New Performance Statistics Report

After enabling performance statistics, create a Performance Statistics report to view the statistics:

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

The new report is created and the name is displayed in the list under **View Report**. This displays the latest report at the top and the oldest report at the bottom.

Deleting a Performance Statistics Report

You may have to delete a Performance Statistics Reports to free up storage or to simply remove old reports from Sterling Integrator:

Procedure

1. From the **Administration** menu, select **Operations > System > Performance > Statistics**.
2. In the Performance Statistic Report Manager page, under **Delete Reports**, next to **Delete All Inactive Reports**, click **Go!**. All the reports, except the latest report, are deleted from Sterling Integrator.

Chapter 8. Troubleshooting

Database Maintenance Check

Sterling Integrator performs a scheduled maintenance check on the database every Monday at 4:00 a.m. This is the default setting and is configurable. This maintenance check is performed by the DB Monitor service and its associated Schedule_DBMonitorService business process. The system verifies if any of the following conditions have occurred:

- Database is more than 80 percent full.
- Database has grown more than 30 percent since the last maintenance check.
- Database has to be indexed.

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

If you are notified that the database has to be indexed, contact your database administrator or use the `db_optimization_tool` script to rebuild the indexes and optimize the database.

Important: The reports and tools described in this section are not a replacement for standard database tools and maintenance scripts, and should be used only if your DBA is not involved with regular maintenance on the Sterling Integrator database.

View the Database Statistics Report

To view information about business processes in the database:

1. From the **Administration** menu, select **Operations > Reports**.
2. In the Reports page, under **Search**, enter DBStats in the **By Configuration 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** next to the DBStats report.
4. Select the appropriate report format from the drop-down list and click **execute**.

Database Statistics Report

The Database Statistics Report provides statistical information about business processes in the database. You should review the Database Statistics Report if you experience full database or failed database connection problems, or if you receive automated notification that the database's health needs to be checked.

The Database Statistics Report is divided into the following sections:

Report Section	Description and Fields
Business Process Runs in Active System	Lists all the business process definitions that have had at least one instance executed. Fields are: <ul style="list-style-type: none"> • Business Process Name – Name of the business process instance. • Number of Runs – Number of times the instance has run. • Number of Persisted Steps – Number of steps persisted per business process definition. • Average Number of Persisted Steps – Average number of steps persisted.
Business Process Runs in Active System - Last 24 Hours	Lists all the business process definitions that have had at least one instance executed in the last 24 hours.
Business Process Runs in Active System - Interval of Last 48 to 24 Hours	Lists all the business process definitions that have had at least one instance executed during the previous day (24-48 hours before the current time).
Number of BPs by Removal Method	Shows the number of business processes flagged to be archived, purged, or indexed. It contains the following fields: <ul style="list-style-type: none"> • Archive Flag – Shows the number of business processes in each category: <ul style="list-style-type: none"> – 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	Number of eligible business processes that are flagged to be archived, purged, or indexed. 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. These business processes will not be purged until they become eligible for indexing.
Archive Dates by Removal Method	Shows the date range of business processes that are flagged to be archived, purged, or indexed. Fields are: <ul style="list-style-type: none"> • Archive Flag – Shows one of the following conditions: <ul style="list-style-type: none"> – 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	Shows the number of rows for each table in the active system. Fields are: <ul style="list-style-type: none"> • Table Name – Name of the table • Count – Number of rows in the table
Table Row Counts - Last 24 Hours	Shows the number of rows for each table in the active system during the last 24 hours.

Report Section	Description and Fields
Table Row Counts - Interval of Last 48 to 24 Hours	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).
Table Row Counts - Index	Shows the number of rows that have not been indexed for each table in the active system.
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.
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.
Number of Rows with No Matching Records in Archive Info	Shows the number of rows in a table that are orphaned (have no matching records in the archive data). It contains the following fields: <ul style="list-style-type: none"> • Table Name – Name of the table • Orphaned Record Count – Number of orphaned rows
No Matching Workflow ID (Null)	The section shows the number of records in TRANS_DATA with NULL workflow IDs or -1 workflow IDs. It contains the following fields: <ul style="list-style-type: none"> • 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.

Optimize the Database with the db_optimization_tool Script

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

Note: The db_optimization_tool script does not work for DB2.

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

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

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.

Option	Description
-t	Date and time cutoff to look for expired data. Format: yyyyMMdd-HH:mm:ss.SSS. 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, 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 be displayed only 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 is displayed on the screen. Must be used with -l or -r. Optional.
-h or -?	View the help screen.

Command Examples

To view a list of table indexes that have to be rebuilt, and to 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 the tables, run the following command:

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

Full Database Issues and Resolution

For a variety of reasons, you may find that your Sterling Integrator database has become full. This may cause performance issues that appear to be related to memory or other resource usage.

If your database is set to auto extend, the database may attain 99 percent capacity before extending automatically. In this situation, the 99 percent capacity does not indicate a performance issue.

After you run Sterling Integrator 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 Sterling Integrator uses, verify that the other applications are also experiencing slow database returns before making any database adjustments.

If the other applications are not experiencing database issues, the performance issue may be related to a cause other than the database.

Symptoms and Causes of a Full Database

The symptoms of a full database may include:

- Database does not accept new data.
- Database Usage Report shows that your database usage is increasing to capacity.
- Sterling Integrator interface is slow.
- Business process execution times are slow.

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

- The database is not sized properly. Determine if the current database size exceeds the expected size determined with your Sterling Commerce Sales Representative at the time of sales. (Customers must have this information in the form of a sizing sheet.)
- 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.
- 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.
- Backup 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 on which 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 on which 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 that can be archived elsewhere are being stored in the database, 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.
- Schedule_BPRcovery business process is either not running or failing, causing business processes to remain in the database in an interrupted state.
- IndexBusinessProcess business process is either not running or failing, causing no business processes to be flagged for archival or purging.
- AssociateDocsToBP business process is either not running or failing, causing expired documents to not be associated for being purged from the system.

To determine the cause of a full database:

- Monitor the Database Usage Report in Sterling Integrator 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, refer to the documentation pertaining to your database.

- Review the archive.log and database-specific log by navigating to **Operations > System > Logs** for information about archiving and purging activities and errors.

Resolve Full Database Issues

Try the following approaches to resolve full database issues.

Approach	Details
Change persistence levels	For information about how to change persistence levels, refer to the topic Changing Persistence Levels.
Complete indexing activities	<p>If your index business process is scheduled, but not running, no data is flagged for archival or purging, and your database may 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.</p> <p>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.</p> <p>If the index business process is running and failing, contact Sterling Commerce Professional Services for assistance.</p>
Complete archiving activities	<p>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.</p> <p>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.</p> <p>If the archival business process is running and failing, contact Sterling Commerce Professional Services for assistance.</p>
Complete purging activities	<p>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.</p> <p>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.</p> <p>If the PurgeService business process is running and failing, contact Sterling Commerce Professional Services for assistance.</p>

Approach	Details
Complete Associate Docs to BP activities	<p>If an associate document to BP business process is scheduled but not running, your database can fill up because expired document data will not be associated to be removed.</p> <p>Review the business processes and verify that the schedule is defined and turned on in the business processes, and that business processes are purged after a reasonable amount of time.</p>
Reduce the amount of large or old data stored in the database	<p>If old data and large files are using database resources and causing your database to fill up, archive those files to another file system instead of to database tables.</p>
Optimizing the Size of Your Database	<p>As your business and processing needs change, you may find that your initial database size requirements are different from your requirements. If your database is filling and your volume is increasing:</p> <ul style="list-style-type: none"> • Increase the size of your database to reflect your processing needs. • Compress your database, if your database has this feature, in order to save more data to the same size database. <p>For more information about compressing your database, refer to your database documentation.</p>
Optimizing the Data Retention Period	<p>You may be retaining information in the database for longer periods of time than necessary.</p> <p>Set the data retention period in the database to lower levels as appropriate.</p> <p>For more information about setting data retention periods in your database, refer to your database documentation.</p>
BPREcovery Business Process is Not Running or is Failing	<p>If your BPREcovery business process is either not running or is failing:</p> <ol style="list-style-type: none"> 1. Contact Sterling Commerce Professional Services for assistance with the BPREcovery business process 2. Run the corrected BPREcovery. 3. Restart, resume, or terminate interrupted business processes. <p>You can perform a mass termination of business processes if many are in an interrupted state.</p>

Document Persistence

Sterling Integrator offers two options for persisting documents: a database option and a file system option. When dealing with database sizing, it is important to understand these differences and ensure that you have chosen the correct option for your environment.

- Database option – The payload of each document in the system is stored in the DATA_TABLE or TRANS_DATA tables. When documents (payloads) become large, it leads to inefficient use of database disk space and results in greatly increased network traffic.
- File system persistence option – This option was created to store the payload of documents out on disk. For documents with file system persistence, the DATA_TABLE/TRANS_DATA tables are still utilized. The difference is that the

content of the BLOB data is not the payload, but a serialized Java HashMap that contains the file name of the payload. This is used to tie the file with the document.

Use of the database or the system for persisting the document payload is specified in the jdbc.properties file with the defaultDocumentStorageType property. This setting is the system default and can be inherited or overridden at the WFD (workflow definition) level. The location of document payloads is also configurable in the jdbc.properties file with the document_dir property.

Database Connection Issues

You may find that your business process throughput slows down in Sterling Integrator during data bursts (sudden increase in data volumes) or connection or cursor leaks (unreleased database connections). These can also result in an increased number of failed pool requests.

Symptoms and Causes of Failed Database Connection Requests

The 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 as seen in the Database Usage report by navigating to **Operations > System > Troubleshooter**. This report shows the number of pool requests and failed requests.

Determining the cause of an increasing number of failed pool requests necessitates you to investigate many areas of the system and the way in which you are implementing Sterling Integrator. The 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 being released back to Sterling Integrator for use by other business processes. This leads to fewer available connections to the database.
- Improperly tuned performance properties – Fewer database pools configured than is required by your processing volumes. For more information, refer to the “Calculating the Settings Required” on page 121 topic.

To determine the cause of failed database connection requests:

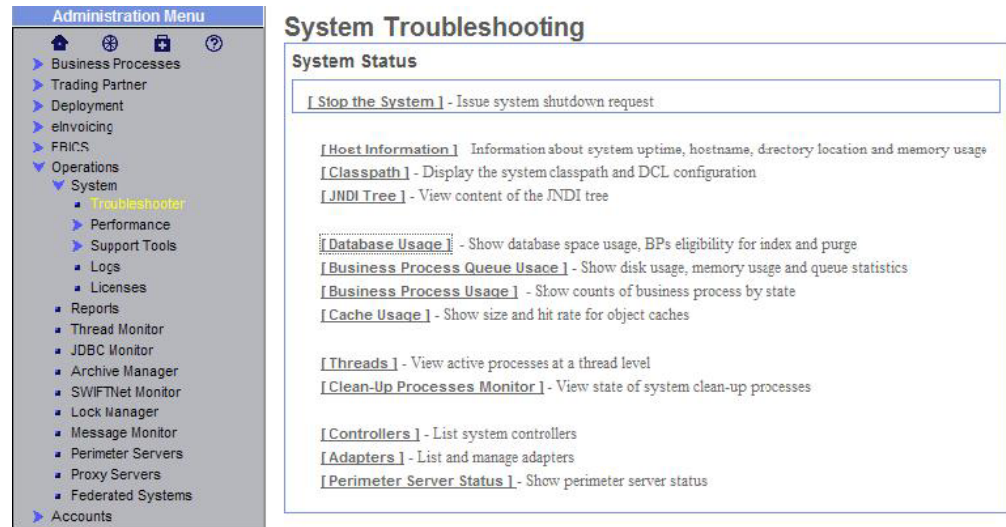
- Monitor the Database Usage report.
- Monitor database tools specific to your database. For more information about this, refer to your database documentation.

Monitor the Database Usage Report

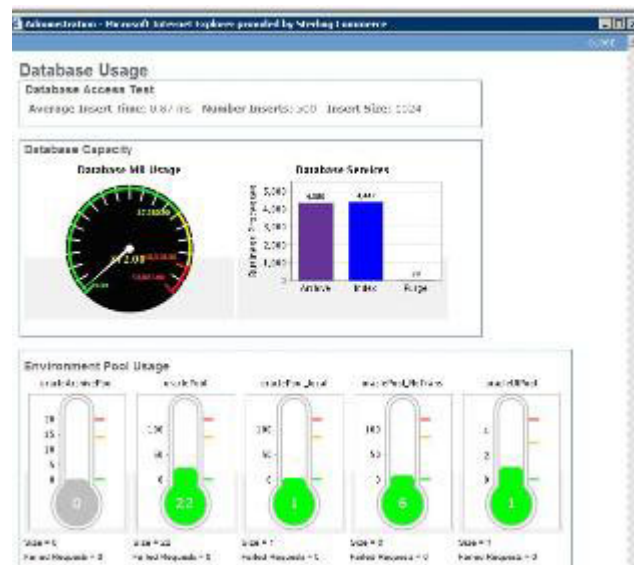
This report shows the number of failed requests for each pool type. You can use this information to determine which pool settings to change in the Performance Tuning Utility.

1. From the Administration menu, navigate to **Operations > System > Troubleshooter** as shown in the following figure.

- The System Troubleshooting page is displayed as shown in the following figure. Click **Database Usage**.



The Database Usage Report (details) is displayed, as shown in the following figure.



Resolve Database Connection Issues by Increasing the Database Pool Size

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

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

To increase the size of specific database pools in Sterling Integrator, navigate to **Operations > System > Performance > Tuning**. For specific information about increasing the database pool size, refer to the topic Performance Tuning Utility.

Database Down Check

Sterling Integrator performs a scheduled check on the database connection every 120 seconds to verify that it 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 stating 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 Sterling Integrator 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 increasing intervals of 1, 5, 15, 30, 45, and 60 minutes unless Sterling Integrator 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)
```

Configure DB Resource Monitor

The DB Resource Monitor can be configured to perform the database down check at intervals other than the default interval of every two minutes. To change the interval, perform the following tasks:

1. In the `install_dir/properties` directory, locate (or create, if necessary) the `customer_overrides.properties` file.
2. Open the `customer_overrides.properties` file using a text editor.
3. Add the following override statement. This statement will set the value of the `DBResourceMonitor.delay` property of the `resource_monitor.properties` file.

```
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 120000 runs the check every 120 seconds, or once every two minutes. If you want to configure the check to run every 10 minutes (as opposed to the default), you

should set the `new_interval` 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 Sterling Integrator to use the new values.

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

Tracking JDBC Connections

Sterling Integrator 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 the 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 to perform troubleshooting by yourself, 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 the JDBC report is viewed or saved.

Enable or Disable JDBC Tracking

JDBC tracking that is enabled through the JDBC Monitor will be reset to Disabled if Sterling Integrator is restarted. This feature helps prevent degraded performance if JDBC Tracking is accidentally left Enabled after you have finished the troubleshooting activity. You should always disable JDBC Tracking when you have finished the troubleshooting activity.

To enable or disable JDBC tracking:

1. From the **Administration** menu, select **Operations > JDBC Monitor**.
2. To change the enable/disable state, select one of the following in the **On/Off** box:
 - If JDBC tracking is disabled, click (**Enable**) to enable JDBC tracking.
 - If JDBC tracking is enabled, click (**Disable**) to disable JDBC tracking.

If you want to enable JDBC tracking for extended troubleshooting so that it remains enabled even if Sterling Integrator is restarted, you can use the internal **use tracking** parameter.

To control JDBC tracking with the **use tracking** parameter:

1. Stop Sterling Integrator.
2. In the `install_dir/properties` directory, locate the `jdbc.properties.in` file.
3. Open the `jdbc.properties.in` file using a text editor.
4. In the `jdbc.properties.in` file, locate the **use tracking** parameter.
 - To enable JDBC tracking, set the `databasePool.useTracking` property to true.
 - To disable JDBC tracking, set the `databasePool.useTracking` property to false.
5. Save the `jdbc.properties.in` file.

6. Restart Sterling Integrator to apply the change.

Note: Always disable JDBC tracking when you have completed the troubleshooting activity in order to improve performance.

View or Save a JDBC Report

To view a report, you must first enable JDBC tracking.

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

Types of Cache Memory

Sterling Integrator utilizes the memory cache types described in the following table to execute processes:

Memory Types	Description
Soft reference cache	When objects are removed from the memory cache in order to keep the memory cache size constant, they are moved to a soft reference cache, which can grow or shrink based on the available memory. If the JVM has 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. You can configure this in the noapp.properties file.
Disk cache	Objects can be read more quickly from the disk than from the database. When objects are no longer in the soft reference cache because they have been garbage collected, the disk cache provides a faster access mechanism than object retrieval from the database. You can configure this in the noapp.properties file.

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 becomes constrained. 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 reread from there.

This behavior is similar to an operating system disk cache. The disk cache retains recently read files or sections of files in memory in order to speed up subsequent access to that data. However, if the system requires 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 fall when you open a large file.

Sterling Integrator follows this practice with soft caches. Most soft caches in Sterling Integrator are implemented by Java Soft References.

JVM short-lived min/max memory values are used for generational garbage collection. For more information about generational garbage collection, refer to the following Web sites:

- 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 Sterling Integrator and the number of CPUs, the Sterling Integrator tuning wizard will allocate disk cache and in-memory cache. In-memory cache is a combination of the soft reference cache and memory cache displayed on the cache usage screen.

Symptoms and Causes of Inefficient Cache Usage

The symptoms of inefficient use of cache include:

- Slow-running business processes
- Longer completion time for business processes
- Reduced number of cache hits, as reported in the Cache Usage report, with a corresponding increase in the number of requests

Determining the reasons behind the inefficient use of cache may require you to investigate many areas of your system and the way you are implementing Sterling Integrator.

The reasons for inefficient use of cache may include:

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

To determine the reason behind the inefficient use of cache, review the Cache Usage report that is displayed in the System Troubleshooting (**Operations > System > Troubleshooter**) page for the number of counts, requests, and hits for each cache in Sterling Integrator.

An increase in the number of requests and a decrease in the number of hits indicates that you have to increase the cache size for the affected cache. Each request that is not met with a hit requires a call to the database for data, which slows down the processing time and uses more resources, which could be used by other components of Sterling Integrator.

For example, if the SchemaCache on the Cache Usage report has 5 counts, 5 requests, and 0 hits, it indicates 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. Zero hits means that the business processes requesting the schema from the cache did not find the requested data and called the database to receive the schema data, which in turn slowed the processing time.

Resolving Inefficient Cache Usage

Following are the different ways in which inefficient use of cache can be resolved, depending on the identified cause:

- Improperly Tuned Cache Performance Properties
- High Cache Allocation for Less Frequently Used Large Objects
- Low Cache Allocation for More Frequently Used Small Objects

For more information about changing the cache values in the Performance Tuning Utility, refer to the topic Viewing or Editing Performance Configuration Settings.

For more information about changing the cache values manually using the `noapp.properties` file, refer to the topic Manual Performance Tuning.

Improperly Tuned Cache Performance Properties

If you review the Cache Usage Report and notice that the number of requests are increasing and the number of hits are 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.

Following are the possible consequences of increasing the cache size:

- Increasing the size too much does not make the system any more effective than reading from the disk. Depending on your system and your configuration, you may have 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 that the number of requests to the cache are low, you can 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.

Decreasing the cache size too much may cause a reduced number of hits to the caches, which results in a call to the database for data, and increases the processing time. Depending on your system and your configuration, you may have 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 that the number of requests to the cache are high and the number of hits to the cache are low, you can 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 being used frequently with your business processes. Review the values for the cache properties 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.

Increasing the cache sizes too much may cause a reduced number of hits to the caches for larger objects. This results in a call to the database for data, and increases the processing time. Depending on your system and your configuration, you may have to adjust the cache settings to attain peak performance without losing the benefit of cache over disk.

Correcting Dropped Data Problems

Proper requirements planning and performance tuning usually prevent most memory-related issues from occurring. 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 you adjust memory allocations after initial performance tuning, you must increase the memory amount by the smallest increment to resolve the issue.

Because Sterling Integrator interacts extensively with other systems, sometimes, problems with dropped data occur because of the configuration settings of the other systems.

For example, when Sterling Integrator is under sufficient load and requires business processes to be placed in queue, the HTTPSyncSend adapter attempts to respond with a Message Disposition Notification (MDN). If this is not completed in a timely manner, the client that initiated the session may release the connection because of a timeout. If this occurs, the HTTPSyncSend operation appears to complete successfully by handing off to perimeter servers, but there is a possibility of the data being dropped.

In this example, the timeout settings in the client can be adjusted to resolve the problem. Alternatively, the amount of memory available to Sterling Integrator can 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 Sterling Integrator, use the Performance Tuning wizard (**Operations > System > Performance > Tuning**) and change the value in the **Physical memory (MB) allocated to SI** box to a higher value. For specific information, refer to the topic Viewing or Editing Performance Configuration Settings.

Correcting Out-Of-Memory Errors

Occasionally, a specific operation may require more memory than the current configuration allocates to Sterling Integrator. For example, with the default settings, attempting to view a large code list (over 100,000 entries) may result 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 Sterling Integrator, use the Performance Tuning wizard (**Operations > System > Performance > Tuning**) and change the value in the **Physical memory (MB) allocated to SI** box to a higher value. For specific information, refer to the topic Viewing or Editing Performance Configuration Settings.

Understanding Business Processes

You may find that business process execution times are longer than expected, which might be the result of an increase in volume or improperly written business processes. This indicates that either the performance does not match your requirements and corrective action is required, or that you have business processes containing errors and require manual attention.

The following concepts are explained under this topic:

- Symptoms and Causes of Poor Business Process Execution Time

- Resolving Halted, Halting, Waiting, or Interrupted Business Processes

Symptoms and Causes of Poor Business Process Execution Time

The symptoms of increasing business process execution time include:

- Performance Statistics Report showing the Min (ms), Max (ms), and Average (ms) time increasing for business processes, activities, and internal system activities. To access this report, navigate to **Operations > System > Performance > Statistics**.
- Increasing business process execution time in the wf.log in the install_dir/logs directory.

Determining the cause of increasing business process execution time may require you to investigate many areas of the system and the way in which you are implementing Sterling Integrator.

Typically, the cause of increasing business process execution time is because of:

- The database is full or is receiving too many requests
- Improperly designed business processes

Improperly designed business processes may involve the following issues:

- Business process persistence levels are set too high, which causes too much data to be stored to the database.

For more information about persistence levels, refer to the topic [Changing Persistence Levels](#).

- Unnecessary service invocations, which use resources that can otherwise be used by Sterling Integrator's other business processes or components.
- Improperly tuned cache properties, which cause the business process to query the database or disk for information that can be provided in a queue.

For more information about caches, refer to the topic [Resolving Inefficient Cache Usage](#).

- A full database or database maintenance issues, which cause longer query times and increased execution times.

If you find that all your business processes are experiencing an increase in execution time, you may also find that your database is full or is receiving many requests, causing the entire processing to slow down.

- Unnecessary loops, which cause the use of resources that can otherwise be used by Sterling Integrator's other business processes or components.
- Improperly written XPath statements, which use resources that can otherwise be used by Sterling Integrator's other business processes or components.

Use relative paths when 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 because this causes the entire process data to be searched, and slows down the search time significantly.

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

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

Determining the Causes of Poor Business Process Execution Time

To determine the cause of increasing business process execution time:

- Review the wf.log file in the *install_dir*/logs file for time and date stamp information showing increases in execution time, and errors in each business process.
- Review the Performance Statistics Report by navigating to **Operations > System > Performance > Statistics** and then accessing the report. This report shows the number of invocations and processing time for each business process and system activity.
- Review business processes that show increasing execution time for improper design, mode, or XPath statements, for example, using a Command Line adapter to retrieve data instead of a 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, which will replace the need for both the Command Line adapter and the 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 because this causes the entire process data to be traversed.

Resolving Halted, Halting, Waiting, or Interrupted Business Processes

A business process in a Halting, Halted, Interrupted_Man, or Interrupted_Auto state requires immediate attention because the business process has stopped processing. Business processes remain in a halted or interrupted state until some action is taken on the business process. Business processes that fail receive a state of Halted, which enables you to take manual corrective action without the business process being archived or purged.

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

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

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 one of the following:
 - Resources
 - A service or activity that is disabled, but will be enabled
- Terminate the business process.
- Restart the business process.

The following sections show how to diagnose and resolve issues with halting, halted, waiting, or interrupted business processes:

- Symptoms of Halting, Halted, Waiting, or Interrupted Business Processes
- Causes of Halting, Halted, Waiting, or Interrupted Business Processes

Symptoms of Halting, 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, refer to the topic “Slow System: Symptoms, Causes, and Resolution” on page 191.
- The database is getting full or having performance issues.
For more information on resolving a database issues, refer to the topic “Full Database Issues and Resolution” on page 176.
- 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.

Causes of Halting, Halted, Waiting, or Interrupted Business Processes

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.
- Sterling Integrator stops running. For example, your site experiences a power outage and you must restart Sterling Integrator 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.

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 System Troubleshooting page (**Operations > System > Troubleshooter**). 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 applicable 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, refer to the topic “Turning On and Turning Off Performance Statistics” on page 171.

- 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 ASI (application server-independent) activity. Review business processes to verify that they are effectively designed to meet your business needs.

Slow System: Symptoms, Causes, and Resolution

You may find that your Sterling Integrator is working slowly because of a variety of reasons. An increase in volume or improperly designed business process may cause slower than expected performance.

The symptoms of a slow system include:

- Slow transaction processing
- Unusually long delay in starting a business process
- Unusually large number of business processes waiting in the queues
- Slow user interface log in
- Slow business process searches
- Slow system services, such as indexing and purging
- Slow user interface performance

Potential Causes of a Slow System

A slow system can be caused by many areas of your system and the in which way you are implementing Sterling Integrator.

Potential causes of a slow system include:

- Improperly designed business processes.
- Various system services being turned off. For more information about turning off system services, refer to the topic "Disabled Service".
- Improperly tuned performance properties. For more information about tuning performance properties, refer to the topic "Improperly Tuned Performance Properties".

Determining the Cause of a Slow System

To determine the cause of a slow system, complete some, if not all, 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, it indicates that the performance issue is not specific to Sterling Integrator. 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 about logs and system information, refer to the topic [Managing System Logs and Reviewing System Information](#).

- Review the Sterling Integrator Internal System Statistics Report by navigating to **Operations > System > Performance > Statistics**.

This report shows the number of invocations and the processing time taken 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, refer to the topic [Manage Performance Statistics](#).

- Perform a thread dump to identify business process bottlenecks in Sterling Integrator.

For more information about thread dumps, refer to the topic [Conducting a Thread Dump](#).

Resolving a Slow System

A system may have slowed down because of the following reasons:

- Improperly Designed Business Processes
- Disabled Service
- Improperly Tuned Performance Properties

Improperly Designed Business Processes

Improperly designed business processes can cause performance issues with Sterling Integrator. These issues can be averted or resolved by implementing well-designed business processes.

Disabled Service

If a service that Sterling Integrator 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 these issues cause your system to run slowly.

For example, if the index service is disabled, the data in the database cannot be flagged for archiving or purging. If the index service is disabled, archiving 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

Sterling Integrator may run slowly if you have not tuned the system performance using the Performance Tuning Utility or have manually tuned the system performance in the *install_dir/properties/noapp.properties* file and any extension files (ext).

To resolve this, you may have to make changes to the appropriate performance tuning properties either manually or using the Performance Tuning Utility. For more information about performance tuning, refer to the topic Performance Tuning Utility or Manual Performance Tuning.

Conducting a Thread Dump

A thread dump is a snapshot of every thread that is running in Sterling Integrator at the time the thread dump is generated. The thread dump shows exactly what each thread is doing at the time the thread dump is generated. You can conduct a thread dump to determine bottlenecks in Sterling Integrator. Typically, problem threads are in a state of waiting. Threads with a state of runnable or R indicate a healthy thread.

For information about the different states of a thread dump, refer to the topic "Reading a Thread Dump".

The subsequent sections describe command-line methods for working with thread dumps. You can also work with thread dumps from the user interface by accessing the JVM Monitor page that is available by navigating to **Operations > System > Performance > JVM monitor**.

Note: The JVM Monitor page is available only for Linux, IBM AIX, Sun Solaris™, and HP-UX platforms. The heap dump option is available only on Linux and AIX platforms.

The JVM Monitor page enables you to:

- Take a thread dump.
- Download and delete the dumps conducted from the system.
- Enable GC (garbage collector) output. Enabling causes the GC output to be written into the noapp.log.
- (Linux and AIX only) Enable or disable heap dumps. Enabling causes a heap dump to be generated along with thread dumps.

Tip: For the Command Line Adapter 2, when you use **Operations > System > JVM Monitor > Take Thread Dump**, the default Command Line Adapter 2 must match the *CLA2_PORT* in the *sandbox.cfg* to take thread dumps from the User Interface.

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

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

1. Access the computer on which Sterling Integrator is installed.
2. Change your working directory to *install_dir*.
3. In the command line, enter **cat noapp.pid**.
This lists the parent thread process ID.
4. Enter **ps -ef | grep noapp.pid**.

This returns all the child process IDs associated with the parent process ID.

5. Enter **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.

Performing a Thread Dump in UNIX (IBM AIX)

To perform a thread dump in IBM AIX:

1. Access the computer on which Sterling Integrator is installed.
2. Change your working directory to *install_dir* directory.
3. In the command line, enter **cat noapp.pid**.
This lists the parent thread process ID.
4. Enter **ps -ef | grep noapp.pid**.
This returns all the child process IDs associated with the parent process ID.
5. Enter **kill -QUIT noapp.pid child.pid**.
(You must include both the parent PID and the child PID.)
The thread dump is placed in the javacore*.txt file in the *install_dir*/noapp/bin directory.

Performing a Thread Dump in Linux

To perform a thread dump in Linux:

1. Access the computer on which Sterling Integrator is installed.
2. Change your working directory to *install_dir* directory.
3. In the command line, enter **kill -QUIT noapp.pid child.pid**.
(You must include both the parent PID and the child PID.)
The 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

Performing a Thread Dump in Windows

To perform a thread dump in Windows:

1. Navigate to *install_dir*\bin directory.
2. Run **hardstop.cmd** command to stop Sterling Integrator.
3. Run **start_GIS_console.cmd** command from the *install_dir*\bin directory. It will run Sterling Integrator in console mode.
Running start_GIS_console.cmd command does not start Sterling Integrator services. It will run Sterling Integrator services in a series of command prompt windows, which cannot be closed while it is running. You can track the progress by watching si_exe.log as running this command takes longer time for Sterling Integrator to start.
4. Perform a thread dump using the following steps:
 - Go to the debug_noapp.cmd command prompt window.
 - Press **Ctrl + Break**.
 - Open the noapp_exe.log file in the *install_dir*\logs directory and scroll to the bottom.
 - Verify that the thread information is present in the log file.

Performing a Thread Dump in IBM iSeries®

To perform a thread dump in iSeries:

1. Access the computer on which Sterling Integrator is installed.
2. In the iSeries command line, enter **WRKACTJOB**.
3. Locate the two QP0ZSPWP jobs that are running in Sterling Integrator and note the name of the user associated with these jobs.
4. Press **F11** twice to locate the job numbers.
5. In the iSeries command line, enter **DMPJVM** and press **F4**.
6. In the Dump Java Virtual Machine (DMPJVM) panel, enter QP0ZSPWP as the job name and user as the user prompt and jobnumber as the number prompt.
7. Press **Enter**. The thread dump is completed.
8. Repeat step 6 for the second job number.
9. In the command line, enter **WRKSPLF** and locate the two QDMPJVM spooled files. These files are the output from your thread dump.

Reading a Thread Dump

A thread dump contains the following components:

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

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

- 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 virtual machine 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..bbb ffc28]
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$Lock)
at sun.rmi.transport.DGClient$EndpointEntry$RenewCleanThread.
run(DGClient.java:500)
at java.lang.Thread.run(Thread.java:534)

```

Scenario-Based Troubleshooting Tips and Techniques

This topic contains a list of scenario-based troubleshooting tips and techniques. This list is not exhaustive and does not cover all scenarios.

Situation	Explanation/Resolution
<p>During Oracle RAC failover, DMI visibility takes time to recover under high volume test when default queue sizes and thread counts are used.</p>	<p>You can increase the number of visibility queues and event input queue capacity in <code>customer_overrides.properties</code> to ensure that adapters are responsive.</p> <p>To increase the default queue size and the thread counts, perform the following tasks:</p> <ol style="list-style-type: none"> In the <code>customer_overrides.properties</code> file, set the following values: <ul style="list-style-type: none"> <code>dmivisibility.number_visibility_queues=6</code> <code>dmivisibility.event_input_queue_capacity=3072</code> From the <code>install_dir/bin</code> directory, run one of the following commands: <ul style="list-style-type: none"> (UNIX or Linux) setupfiles.sh (Windows) setupfiles.cmd Restart Sterling Integrator.
<p>Load balancing between two nodes after Oracle RAC failover</p>	<p>Connections are equally distributed between two Oracle RAC nodes. When one Oracle RAC node goes down, the active connections from the node that went down will failover to the active node. When the node that went down is up and running again, the connections are not redistributed.</p> <p>To redistribute the connections between the two RAC nodes again, you should set the lifespan to a value greater than 0 for all Oracle connection pools by performing the following tasks:</p> <ol style="list-style-type: none"> In the <code>customer_overrides.properties</code> files, override all Oracle pools lifespan to a value greater than 0. For example, <code>jdbcService.oraclePool.lifespan=180000</code>. From the <code>install_dir/bin</code> directory, run one of the following commands: <ul style="list-style-type: none"> (UNIX or Linux) setupfiles.sh (Windows) setupfiles.cmd Restart Sterling Integrator.

Situation	Explanation/Resolution
<p>While running Oracle 11g in non-RAC or RAC environment, you might encounter index hot spots, which may cause performance degradation.</p>	<p>It is recommended that you create the following indexes in a global partition to improve system performance: Note: If you have not deployed Sterling File Gateway in your environment, you can ignore the indexes starting with the name FG.</p> <ul style="list-style-type: none"> • PERF_IDX_1 • PERF_IDX_2 • PERF_IDX_3 • SCI_IDX_125 • SCI_IDX_195 • PERF_IDX_5 • SCI_IDX_56_2 • PERF_IDX_4 • SCI_IDX_120 • SCI_IDX_109 • SCI_IDX_DMI17 • SCI_IDX_3 • FG_RTE_MODTS_IDX • FG_EA_DF_IDX • FG_RE_DF_IDX <p>Note: If there are any changes in a patch for the partitioned indexes, the patch installation will drop the index and create a non-partitioned index, treating it as a normal index. You should analyze the indexdrop SQL's and recreate the partitioned indexes again if they have changed.</p> <p>The following example explains the procedure to drop an index and create the same in a global partition:</p> <ol style="list-style-type: none"> 1. To drop index PERF_IDX_1, run the following sql command: drop index PERF_IDX_1 ; 2. Create the PERF_IDX_1 in global partitions by running the following command: create index PERF_IDX_1 on WORKFLOW_CONTEXT (BASIC_STATUS,START_TIME,WORKFLOW_ID,WFD_ID) global partition by hash (BASIC_STATUS,START_TIME,WORKFLOW_ID,WFD_ID) partitions 16 store in (SI);
<p>In some cases, perimeter services may fail with the following error. The error will be recorded in the perimeter services log (perimeter.log).</p>	<p>You can improve the situation by reducing high timeout values in communications adapters and by increasing the delayed event queue size (delayed.event.queue.size) in perimeter.properties file.</p> <p>The default value of the delayed.event.queue.size property is 30000. You can increase the value based on your requirement.</p>
<p>In Sterling Integrator release 5.1.2, when you invoke a BP in sync mode using a queue where the MaxPoolSize and MinPoolSize properties in the noapp.properties_platform_ifcresources_ext file are not equal, you may experience performance degradation. This issue will be addressed in a later patch.</p>	<p>To improve performance, set MinPoolSize and MaxPoolSize with the same value for the relevant queue for the affected BP. Care should be taken, as there could be an impact on other BPs .</p>

Situation	Explanation/Resolution
<p>An error or exception in the Sterling Integrator logs like:</p> <pre>com.sterlingcommerce.perimeter.api.gmm. BufferAllocationException: Perimeter Services is full</pre> <p>or</p> <p>System has detected a deadlock, rejecting 1 allocation request(s)</p>	<p>The Sterling Integrator perimeter server uses the <code>gmm.maxAllocation</code> property in the <code>perimeter.properties</code> file to control how much memory can be used for all sessions transferring data. When concurrent sessions reach a certain level, the memory controlled by this value may be exhausted and needs to be increased. The indication of such memory exhaustion is the error message <code>com.sterlingcommerce.perimeter.api.gmm.BufferAllocationException: Perimeter Services is full</code> in one of the log files. To increase <code>gmm.maxAllocation</code> to be more than 384 MB, the JVM heap should be increased accordingly. For more information about the <code>gmm.maxAllocation</code> property, refer to the inline comments in the <code>perimeter.properties</code> file.</p>

Chapter 9. Appendix A

The Performance and Tuning Worksheet

The Performance and Tuning Worksheet can assist you in tuning Sterling Integrator for optimal performance. It helps you take an inventory of the different parts of your installation that affect performance.

Note: This worksheet is optional. It does not replace the *Sterling Integrator System Requirements* documentation that you used to install Sterling Integrator.

Use the Performance and Tuning Worksheet to help you determine your hardware requirements 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 there 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 there in a day?	Enter the number of hours you allow for processing all the data.	
How many translations are required for each transaction?	Enter the largest number of translations you will have if multiple translations are required for a business process. For example, is data mapped to a common format, and is the common format used to interface with all the external systems?	
Is content-based routing required?	If content-based routing is required in your business processes, you require maps, XML encoder, BPML choice, assigns, adapters that are used for routing, and possibly other components.	
Does the content of a transaction have to be verified?	This requires translating data into XML. The data will then have to be put in process data and verified using XPath.	
What percentage of the content of the transaction should be verified?	Typically, only a small part of a transaction has to be verified. However, if most of a transaction requires verification, you may have significant resources depending on the verification path.	

Request	Description	Your Answer
How many rules have to be applied to verify the content?	This is used to calculate the Service Latency.	
How many content elements have to be verified?	Enter the number of elements that have to be verified.	
What is the future system growth rate?	Current daily volume addition 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 that is required?	<p>Enter the level of persistence (storage to the database) you require for your business process.</p> <p>Choices include:</p> <ul style="list-style-type: none"> • Full – Saves a complete copy of process data along each step of the process. • Minimal – Saves all the steps in a business process and selected copies of process data. • None – Saves the first and the last step of a business process, 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.	
Who is the hardware vendor hosting Sterling Integrator?	Enter the hardware vendor's name, for example, IBM, Sun, HP, or Dell.	
What is the speed of CPUs in the host?	Enter the speed in MHz, for example, 450, 1900 (1.9 GHz).	
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 the amount of memory in MB, for example, 500 MB, 2048 MB.	
What is the operating system that you are using to host Sterling Integrator?	Enter the host operating system, for example, Solaris, Windows Server® 2008, AIX, or Linux.	
What is the operating system release and version number?	For example, Red Hat Enterprise Linux, Release 5, or Windows Server® 2008 Server Standard Edition R2.	

Request	Description	Your Answer
What are the JVM version numbers?	For example, IBM JDK 1.6.0 SR8 (64 bit) build pxa6460sr8-20100409_01(SR8). Your JVM version must match the requirements for your version of Sterling Integrator. The build date and lower release numbers must also match.	
What are the Sterling Integrator version and patch numbers?	Enter the version and patch number for your copy of Sterling Integrator.	
Who is your relational database vendor?	Enter the name of the relational database vendor Sterling Integrator uses, for example, MySQL, Oracle, or DB2.	
What is the relational database version and patch number?	Enter the version number and patch number of the relational database.	
What are the other business applications that have 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 require to complete processing.	
What are the other business applications' version number and patch number?	Enter the version number and patch number of the other business applications.	
What data format do you use to 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 to increase system resource requirements.	
What type of inbound transactions do you have?	Enter the types of inbound transactions that Sterling Integrator will process.	
What is the average size of each inbound transaction type?	Enter the average size of your inbound transactions that Sterling Integrator will process. Base your planning on the larger transactions at peak times for better performance.	

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

Chapter 10. Appendix B

workflowLauncher: Running a Business Process from a Command Line

About this task

Sterling Integrator provides a utility called workflowLauncher that allows you to launch a business process from a 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 a command line:

Procedure

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

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 user name from which to include the userToken. 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.

Option	Description
-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>	<p>Specifies a specific version of the business process definition to launch. Optional.</p> <p>Example: -v <i>business_process_version_number</i></p> <p>The version number is not 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 (as described using the Graphical Process Modeler and/or XML). If there is only one version, the version number will be 1. An error will be generated if the version number is not found.</p> <p>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 Sterling Integrator Business Process Source Manager.</p>
-f <i>inputFile</i>	<p>Specifies a file to use as the primary document when launching the business process. Optional.</p> <p>Example: -f <i>path/filename</i></p> <p>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 -e parameter.</p> <p>An error will be generated if the file is not found.</p>
-e <i>encoding</i>	<p>Specifies the encoding of the primary document. Optional.</p> <p>Example: -e <i>encodingName</i></p> <p>Used in conjunction with the -f parameter.</p> <p>To see what encoding names are supported, see the encodings.properties file in the <i>install_dir</i>/properties directory. The full listing is shown in the encodings_large.properties file.</p>
-w <i>count</i>	<p>Specifies the number of times to sleep. Default is 10.</p> <p>Example: -w <i>count</i></p> <p>Used in conjunction with the -i parameter to define the total timeout value. For an example showing the total timeout value, see the -i parameter.</p>
-i <i>time</i>	<p>Specifies the length of time, in milliseconds, to sleep. Default is 3000 (3 seconds).</p> <p>Example: -i <i>time</i></p> <p>Used in conjunction with the -w parameter to define the total timeout value.</p> <p>For example, the default value for -i time is 3000 milliseconds and the default value for -w count is 10. Therefore, the default timeout value is 30,000 milliseconds (3000 milliseconds x 10), or 30 seconds.</p>
-x	Specifies not to return an exit code. Optional. By default, the utility returns an exit code.
-h -?	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.

Index

A

- active threads 110
- activities
 - monitoring 104
 - stopping 105
- AE_ExecuteCycle 129, 135
- AE_ExecuteCycleTime 129, 135
- aftrouting.log 98
- alerterlogger.log 99
- AllowStealing 131, 137
- analyzing log files 93
- archive.log 99
- ASI 134
 - performance tuning 129
 - queue tuning 134
 - queues 133
- ASI queue 134
- ASM 47
- audit.log 103
- auditing 102
- Authentication.log 101

B

- Backupplogs.log 102
- Basic Scheduling policy 133
- bizIntel.log 100
- BLOB 34
- bpdeadline.log 102
- business process
 - resolving 189
 - waiting business process 189
 - resolving halted business process 189
- business process execution time
 - causes 188
 - improving 187
 - symptoms
 - increasing 188
- business process threads 2

C

- cache 2
 - determining cause of inefficiency 185
 - disk 184
 - inefficient 185
 - memory 184
 - soft 184
 - symptoms of inefficiency 185
- cache contents
 - compressing 140
- cache usage 112
- CacheThreshold 130, 136
- capacity requirements 17
 - example 18
- causes
 - slow system 191
- cdinterop_cdjava.log 99
- cdinterop.log 99
- cdsp.log 99

- ceuinterop.log 99
- checklist
 - Sterling File Gateway 10
 - Sterling Integrator 7
- Clean-Up Processes Monitor 113
- cluster environment and JNDI 110
- cluster node status 2
- common3splogger.log 101
- compressing cache contents 140
- configuring
 - database pool parameters 21
 - DB resource monitor 182
 - DB2 37
 - Microsoft SQL Server 52
 - Oracle 26
- console listener
 - enabling 145, 146
- Consume business process 105
- controllers 108
- creating
 - performance statistics report 172
- crypto.log 99
- cspftp.log 99
- csphhttp.log 99

D

- database
 - buffer pool statistics 32
 - connection issues 180
 - failed requests 180
 - resolve 181
 - db_optimization_tool 175
 - down 182
 - full database
 - issues 176
 - resolution 178
 - symptoms 177
 - instance efficiency 32
 - JDBC connections 183
 - maintenance 173
 - management 19
 - monitoring disk space 17
 - optimizing 175
 - planning 17
 - statistics report 173
 - storage 29
 - ASM 29
 - raw disks 29
 - tablespaces 30
 - tables 22
 - ARCHIVE_INFO 23
 - CORRELATION_SET 22
 - DATA_TABLE 22
 - DOCUMENT 22
 - TRANS_TABLE 22
 - WORKFLOW_CONTEXT 22
- DB2
 - configuring 37
 - connections 40
 - event monitors 50

- DB2 (*continued*)
 - indexes 51
 - logs 48
 - LOGFILSIZ 48
 - NUM_LOG_SPAN 48
 - monitoring 37, 49
 - snapshots 49
 - parameters 40, 41
 - performance manager 51
 - registry variables 37
 - settings 37
 - storage 45
 - tablespaces 45
 - volatile tables 52
- DB2 logs
 - LOGPRIMARY 48
- deadlock 50, 53, 61
- defrag 60
- delete.log 99
- deleting
 - performance statistics report 172
- deprecated resources
 - identifying 118
 - monitoring 117
 - removing 118
 - replacing 118
- descriptors
 - file 150, 152
 - file descriptors 128
- determining cause
 - inefficient cache 185
 - slow system 192
- directory naming convention 91
- disabling
 - perimeter server 115
- disk cache size 129
- DiskCachePath 129
- document persistence 179
- DTDs
 - repositories 145
 - trusted domains 145

E

- ebXML.log 99
- EDI 147
- edi.log 99
- EDIINT.log 100
- embeddedEngine.log 100
- EnableDeadlines 130, 136
- enabling
 - console listener 145, 146
 - container JVM 74
 - garbage collection
 - HotSpot 73
 - garbage collection container JVM 74
 - perimeter server 115
- enveloping.properties.in 147
- event.log 100

F

- Fair Share Scheduling policy 133
- file descriptors 150, 152
- file system 17
 - native 19
- ftp.log 100
- ftpclient.log 100
- full database
 - causes 177

G

- garbage collection 65, 73, 74
 - example 66
 - HotSpot 67
 - tools 68

H

- HotSpot
 - tools
 - HPjmeter 69
 - visualgc 69
- HotSpot tools
 - Java Garbage Collector 69
- HotSpot troubleshooting
 - looping process
 - diagnosing 81
- http.log 100
- httpclient.log 100

I

- improving business process execution
 - time 187
- indexes
 - DB2 51
 - Microsoft SQL Server 59
 - Oracle 34
 - placement 22
 - rebuild 34, 59
- inefficient cache
 - resolving inefficiency 185
- InitialCycles 131, 137
- installation path 110
- interrupting threads 104

J

- JavaPriority 130, 136
- JBOD 18
- JDBC
 - connection pools 20
 - implementing 20
 - monitor 1
- JDBC tracking
 - disable 183
 - enable 183
- jdbc.properties.in 141, 183
- jetty.log 100
- JMX event listener 146
- jmx.log 100
- JNDI 110
 - contexts 142

- JNDI and cluster environment 110
- JVM 63
 - garbage collection 69
 - example 66
 - HotSpot 67, 73
 - HotSpot tools 68
 - statistics 65
 - garbage collection tools
 - GCViewer 68
 - heap dumps 83
 - HotSpot
 - deadlock 79
 - monitoring 76
 - optimize performance 69
 - HotSpot monitoring 76
 - HotSpot thread dump
 - tools 78
 - HotSpot troubleshooting 79
 - crash 81
 - exceeds VM limit 80
 - heapdumps 89
 - hung process 81
 - hung process diagnosing 82
 - looping process 81
 - native method 81
 - out of swap space 80
 - out-of-memory 80
 - out-of-memory exceptions 87
 - PermGen space 80
 - stack overflow 83
 - trigger thread dump 88
 - HotSpot tuning parameters 69
- Hprof 83
- IBM 84
 - parameters 85
 - troubleshooting 87
 - tuning 85
- profiling 83
- version 63
- view settings 63, 64

- JVM garbage collection 73
 - enabling container 74
- JVM heap memory 67
- JVM HotSpot
 - blocked thread 79
 - optimize performance
 - runtime 69
 - startup 69
 - statistics
 - HP 74
 - Sun 75
 - thread dump tools
 - Thread Dump Analyzer 78
 - VisualVM 79
 - tuning 67
- JVM HotSpot troubleshooting 87
 - application hangs 87
 - noapp 88
- JVM monitor 1

L

- latency 2
- lifecycle.log 100
- listenerStartup.properties.in 146
- load 3

- log file
 - changing location 97
 - naming convention 92
 - types 98
- logs 2
 - analyzing 93
 - changing system 145
 - DB2 48
 - LOGSECOND 48
 - naming conventions 91
 - overview 91
 - redo 28
 - types 98
 - viewing contents 92

M

- mailbox.log 100
- manageability 17
- managing
 - database 19
 - performance statistics 161
 - system logs 91
- MaxPoolSize 129, 134, 135
- MaxThreads 129, 135
- MaxWaitTime 130, 136
- MemCacheSize 129
- MemCacheThreshold 129
- memory in use 110
- memory leaks 87
- Message monitor 1
- messages 2
 - monitoring 105
- mgmtdash.log 99
- Microsoft SQL Server
 - autogrow 56
 - configuring 52
 - dynamic management views 58
 - indexes 59
 - rebuild 59
 - maintenance plan 58
 - Management Studio 57
 - monitoring 52, 57
 - performance monitor 58
 - settings
 - DB-specific 53
 - instance-specific 52
 - snapshot 53
 - SQL Server Profiler 57
 - storage 56
 - system monitor 58
 - table statistics 59
 - tips 60
- MinPoolSize 130, 134, 136
- modifying log levels 94
- monitoring
 - activities 104
 - AWR report 31
 - database usage report 180
 - DB2 37, 49
 - deprecated resources 117
 - disk space 17
 - HotSpot JVM heap 76
 - HotSpot JVM thread 76
 - messages 105
 - Microsoft SQL Server 52, 57
 - node status 115

monitoring (*continued*)
Oracle 26
threads 104

N

naming conventions 91
noapp.log 101, 146
noapp.properties 135
node monitoring status 115

O

OdetteFTP.log 100
oftp.log 100
open file descriptors 128, 150, 152
ops_exe.log 102
ops.log 98
opsSecurity.log 98
opsServer.log 98
Oracle
 BLOB caching 34
 character set 26
 configuring 26
 cursor_sharing 26
 indexes 34
 rebuild 34
 init parameter 26
 monitoring 26
 table statistics 34
 undo changes 28

P

performance concepts 2
performance configuration settings
 editing 123
 ASI node 64
 container JVM 65
 viewing 123
performance statistics
 creating report 172
 deleting report 172
 managing 161
 turning on/off 171
Performance statistics report 1
performance tuning 142
 ASI 129
 calculating the settings 121
 document body serialization
 threshold 141
 editing 123
 JNDI contexts 142
 manual 128
 methodology 4
 prerequisites 3
 queues 134
 Sterling e-Invoicing 152
 system logging 145
 utility 121
 worksheet 199
Performance tuning wizard 2
performance. 17
perimeter server
 area on troubleshooting page 114
 disabling 115

perimeter server (*continued*)
 enabling 115
Perimeter.log 100
persistence level 143
 business process 143
 global 143
persistence_level 131
pipeline.log 100
planning worksheet
 performance and tuning
 worksheet 199
platform.log 100
Produce business process 105
PSLogger 102

Q

QueueDepth 129, 135
queues
 ASI 133

R

RAID 18
recommended settings 123
redo
 file size 28
 logs 28
report.log 100
requirements planning worksheet 4
rescheduling 130, 136
ReschedulingInterval 130, 136
resolutions
 cache 185
 slow system 192
resolving
 interrupted business process 189
ResourceAllocation 131, 137
resourcemonitor.log 100
rnif.log 101
run.sh 142

S

sap.log 101
sapxi.log 101
scalability 2, 17
schedule.log 101
ScheduleBackup.log 102
schedulemonitor.log 101
ScheduleStopOps.log 102
scheduling policies 2
scheduling policy
 changing 133
 overview 133
SchedulingPolicyName 129, 133
schema objects 22
schemas
 repositories 145
 trusted domains 145
searching messages 106
secureproxy.log 101
security.log 101
segment
 order
 data 153

segment order
 maps 153
server sizing 17
servers 108
servicesctl.log 98
setupfiles.cmd 141, 142, 145, 146, 147
setupfiles.sh 141, 142, 145, 146, 147
sftpclient.log 101
sftpserver.log 101
SGA pool sizing 28
si_exe.log 102
slow system
 causes 191
 determining cause 192
 resolutions 192
 resolving 187, 191
 symptoms 191
sql.log 101
start time 108
starting system logs 94
StartupPS.log 102
startWindowsService.cmd 142
state 110
status viewing
 system logs 91
StealThreads 131, 137
stopping
 activities 105
 system logs 94
stopping threads 105
storage
 DB2 45
 Microsoft SQL Server 56
survivability 17
SWIFTNet monitor 2
swiftnet.log 101
symptoms
 cache 185
 slow system 191
system classpath 110
system information 2, 107
system logs 2
 managing 91
 modifying log level 94
 starting 94
 stopping 94
 turning on and off 94
system recovery 2
system troubleshooting 107
System troubleshooting 2
system troubleshooting page 107
system.log 100, 101

T

T-SQL 60
tablespaces
 DB2 45
 DMS 46
 SMS 45
thread
 blocked 79
thread dump
 conducting 193
 example 76
 IBM iSeries 195
 Linux 194

- thread dump (*continued*)
 - performing 78
 - reading 195
 - tools 78
 - UNIX (HP-UX or Sun Solaris) 193
 - UNIX (IBM AIX) 194
 - Windows 194
- Thread monitor 2
- Thread Monitor 104
- threads 2
 - interrupting 104
 - monitoring 104
 - stopping 105
 - viewing 104
- throughput 2
- tracking.log 99
- trigger thread dump 88
- troubleshooting
 - system 107
- trusted domains
 - DTDs 145
 - schemas 145
- tuning 134
 - HotSpot JVM
 - parameters 69
 - IBM JVM 85
- tuning HotSpot JVM 67
- tuning.properties 64, 137, 142
- tx.log 101
- txtrace.log 101

U

- ui.log 100
- uptime 108

V

- viewing
 - JDBC report 184
 - log files 92
 - server status 91
 - service status 91
 - threads 104
- visibility.log 101

W

- wait events 32
- webdav.log 101
- WebSphereMQSuite.log 102
- webx.log 101
- wf.log 99
- wfexception.log 99
- wfstatistics.log 99
- workflowLauncher 203
- workflows.properties file 141
- worksheet 199
- wstrm.log 102
- wssec.log 101

Notices

This information was developed for products and services offered in the U.S.A.

IBM® may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing

IBM Corporation

North Castle Drive

Armonk, NY 10504-1785

U.S.A.

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing

Legal and Intellectual Property Law

IBM Japan Ltd.

19-21, Nihonbashi-Hakozakicho, Chuo-ku

Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be

incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation

J46A/G4

555 Bailey Avenue

San Jose, CA 95141-1003

U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© IBM 2014. Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. 2014.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com)[®] are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Linear Tape-Open, LTO, the LTO Logo, Ultrium and the Ultrium Logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Connect Control Center®, Connect:Direct®, Connect:Enterprise®, Gentran®, Gentran®:Basic®, Gentran:Control®, Gentran:Director®, Gentran:Plus®, Gentran:Realtime®, Gentran:Server®, Gentran:Viewpoint®, Sterling Commerce™, Sterling Information Broker®, and Sterling Integrator® are trademarks or registered trademarks of Sterling Commerce®, Inc., an IBM Company.

Other company, product, and service names may be trademarks or service marks of others.



Product Number:

Printed in USA