Sterling B2B Integrator

Performance Management

Version 5.2

Download from Www.Somanuals.com. All Manuals Search And Download.

Download from Www.Somanuals.com. All Manuals Search And Download.

Sterling B2B Integrator

Performance Management

Version 5.2

Download from Www.Somanuals.com. All Manuals Search And Download.

Note
Before using this information and the product it supports, read the information in "Notices" on page 247.

Copyright

This edition applies to Version 5 Release 2 of Sterling B2B Integrator and to all subsequent releases and modifications until otherwise indicated in new editions.

© Copyright IBM Corporation 2000, 2015.
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Download from Www.Somanuals.com. All Manuals Search And Download.

Contents

Performance Management 1
Overview 1
Intended Audience 2
System Components 2
Performance Tuning Methodology 4
Performance Recommendations Checklists 5
Sterling B2B Integrator: General Recommendations Checklist 5
IBM Sterling File Gateway: Specific Recommendations Checklist 9
EBICS Banking Server: Specific Recommendations 12
Database Management. 13
Planning 13
Server Sizing 14
Storage and File Systems 14
Database management for Sterling B2B Integrator 16
Oracle Database Configuration and Monitoring 22
IBM DB2 for LUW Configuration and Monitoring 35
Microsoft SQL Server Configuration and Monitoring 51
Java Virtual Machine 61
Garbage Collection Statistics 62
JVM Verbose Garbage Collection 62
IBM JVM Garbage Collection Example 62
Introduction to HotSpot JVM Performance and Tuning 63
Introduction to the IBM JVM Performance and Tuning Guidelines 78
Monitoring Operations 82
Managing System Logs 82
Auditing 93
Monitoring a Business Process Thread 94
Monitoring Messages 96
Reviewing System Information 97
Monitoring Node Status 106
Monitoring Deprecated Resources 109
Soft Stop of Sterling B2B Integrator 110
Hard Stop of Sterling B2B Integrator 120
Performance Tuning 120
Performance Tuning Utility 120
Manual Performance Tuning 166
Performance Statistics 203
Turning On and Turning Off Performance Statistics 213
Reporting Performance Statistics 213
Troubleshooting 214
Database Maintenance Check 214
Full Database Issues and Resolution 218
Database Connection Issues 222
Database Down Check 224
Tracking JDBC Connections 226
Types of Cache Memory 227
Symptoms and Causes of Inefficient Cache Usage 228
Resolving Inefficient Cache Usage 228
Correcting Dropped Data Problems 229
Correcting Out-Of-Memory Errors 230
Understanding Business Process 230
Symptoms and Causes of Poor Business Process Execution Time 230
Resolving Halted, Halting, Waiting, or Interrupted Business Processes 232
Slow System: Symptoms, Causes, and Resolution 234
Performing a Thread Dump 236
Scenario-Based Troubleshooting Tips and Techniques 239
Performance and Tuning Worksheet 241
workflowLauncher: Running a Business Process from a Command Line 244

Notices 247

Index 251

© Copyright IBM Corp. 2000, 2015

Download from Www.Somanuals.com. All Manuals Search And Download.

Download from Www.Somanuals.com. All Manuals Search And Download.

Performance Management

You can manage the performance of Sterling B2B Integrator according to your needs.

Overview

Typically, performance in general, and optimal performance in particular, are associated with the following criteria: latency, throughput, scalability, and load.

- 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 B2B Integrator 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 Sterling B2B 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 "Performance Tuning Utility" on page 120.
- Add additional hardware.
- Tune your business processes to run more efficiently on Sterling B2B Integrator.
- Monitor and archive the database to free up resources.
- Create Sterling B2B 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" on page 241 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.

© Copyright IBM Corp. 2000, 2015

Download from Www.Somanuals.com. All Manuals Search And Download.

- 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 B2B 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 B2B 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.

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
- IBM® Sterling B2B Integrator Administrator

Performance management documentation includes the following information:

- Background information about the different performance and tuning issues when running Sterling B2B IntegratorSterling B2B Integrator.
- Step-by-step information that helps you:
  - Optimize the performance.
  - Diagnose and resolve performance issues, if any, 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

- The Sterling B2B 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**)
  - 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 B2B 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 B2B 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 B2B 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 B2B 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 B2B Integrator administrators read all the topics.

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B Integrator:

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

   Note: If you purchased IBM 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 IBM 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 B2B Integrator, and on the DMZ computer if you are running perimeter servers. Both the JVM versions must match each other and your Sterling B2B 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 B2B 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 “Performance Tuning Utility” on page 120.
9. Navigate to **Operations** > **System** > **Performance** > **Statistics**. In the Statistics 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

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B 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 Benchmarks*dd/mm/yy*.
16. Conduct the same test that you conducted in step 11.
17. Review the Benchmarks*dd/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.

Performance Recommendations Checklists

Performance Recommendations Checklists provide guidelines to plan for the required elements, and enhance the performance of Sterling B2B Integrator and related components.

Sterling B2B 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 B2B 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 B2B 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 B2B Integrator on certified OS versions and levels. Refer to the *System Requirements* 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

Recommendation	Test	Production	Comments
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: *PSALLOC=EARLY* *NODISCLAIM=TRUE*
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.
Monitor Paging	C	C	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. You can monitor paging levels using standard operating system or third-party measurement tools, such as: • UNIX/Linux – SAR • Windows – System Monitor
Monitor Heap Garbage Collection Performance	C	C	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.

Download from Www.Somanuals.com. All Manuals Search And Download.

Recommendation	Test	Production	Comments
EDI Encoder Service EDI Envelope Service	R	R	As of Sterling B2B Integrator 5.2, the EDI Encoder Service and EDI Envelope Service notify the user of an incorrect value or mismatch in the Mode parameter by writing information on how to correct the problem in the status report of the business process step. This can cause a performance degradation in some very heavy usage scenarios by adding additional load on the database. It is recommended that all business processes that use the EDI Encoder Service and/or EDI Envelope Service be reviewed to make certain they are using the correct Mode parameter, and that the values match in the business process. • EDI Encoder Service should be set as `Mode`. For example: `<assign to="Mode">IMMEDIATE</assign>` • EDI Envelope Service should be set as `MODE`. For example: `<assign to="MODE">IMMEDIATE</assign>`
Planning			
Server node sizing	NA	C	You should ensure that you have sufficient computing capacity to process peak transaction volumes. Refer to *System Requirements* documentation to get an estimate of the processor, memory, and other requirements. You can also engage IBM 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.
Database disk sizing	NA	C	You should ensure that you have sufficient disk space for the database server. 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.
JVM			

Download from Www.Somanuals.com. All Manuals Search And Download.

Recommendation	Test	Production	Comments
JVM version	C	C	You must ensure that you install Sterling B2B Integrator on certified JVM versions and levels. Run the **–version** command in the command prompt to ensure that you have installed the correct version.
Verbose GC statistics	NA	C	You can enable verbose GC statistics collection. Understanding the health of GCs for each JVM is critical for performance.
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 B2B 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	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. Sterling B2B Integrator supports both 32-bit and 64-bit JVMs. You should allocate more heap space when you are running a 64-bit JVM.
Database			
Monitor and regulate indexes	C	C	Sterling B2B Integrator comes with a default set of indexes. In some cases, the indexes may not apply to your operational environment. 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

Recommendation	Test	Production	Comments
Cursor sharing	C	C	If you are using Oracle database, cursor sharing enables dynamic SQL to be reusable, thereby reducing the contention on the shared pool. You should ensure that you set cursor_sharing=EXACT
Parameters governing DB2 locking strategy	C	C	Set DB2_EVALUNCOMMITTED, DB2_SKIPDELETED, and DB2_SKIPINSERTED to reduce lock contention.
Parameters governing DB2 memory	C	C	Set parameters to manage various memory structures such as LOCKLIST, SORTHEAP, and so on, to AUTOMATIC.
Volatile table	NA	C	Mark tables that change significantly, as volatile.

IBM Sterling File Gateway: Specific Recommendations Checklist

IBM Sterling File Gateway is installed on an instance of Sterling B2B 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 B2B 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 B2B 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 property files documentation in the online documentation library.

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

Download from Www.Somanuals.com. All Manuals Search And Download.

Recommendation	Test	Production	Comments
Increase the value of Sterling File Gateway.	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)
If you are processing very large files, increase the probe values to avoid timeout conditions.	R	R	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. First, probe 120 times, with 100 milliseconds between each probe, for a total of 12 seconds. Default value: • bpCompletionProbes.1=120 • bpCompletionSleepMsec.1=100 Then, probe 600 times with 2000 milliseconds between each probe, for a total of 1200 seconds (20 minutes). Default value: • bpCompletionProbes.2=600 • bpCompletionSleepMsec.2=2000
If you have a high volume of PGP traffic, you can improve your performance by specifying a group for the file gateway.	R	R	The name of the Command Line 2 adapter to be used for PGP packaging and unpackaging. You can override this property in the `customer_overrides.properties` 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. Default value: pgpCmdline2svcname= PGPCmdlineService
If you have very large files that will be processed by PGP, increase the value of the file gateway.	R	R	Timeout value, in milliseconds, for PGP package and unpackage operations invoked by Sterling File Gateway. Default value: 240000 milliseconds (4 minutes)

Download from Www.Somanuals.com. All Manuals Search And Download.

Recommendation	Test	Production	Comments
If you have high volumes of FTP traffic, you can improve your performance by specifying a group.	R	R	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. Default value: ftpClientAdapterName= FTPClientAdapter
Decrease the value of evaluation frequency.	R	R	You can enable either MailboxEvaluateAllAutomaticRules or MailboxEvaluateAll AutomaticRulesSubMin. MailboxEvaluateAll AutomaticRulesSubMin 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 MailboxEvaluateAll AutomaticRulesSubMin business process.
Suppress Duplicate Messages	R	R	Prevents duplicate messages from using system resources.
Increase the number of steps a business process must complete prior to returning to the queue.	R	R	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. 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.
Increase the time period that a business process can use a thread, before releasing it to be used for another business process.	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.

Download from Www.Somanuals.com. All Manuals Search And Download.

Recommendation	Test	Production	Comments
Increase the number of concurrent threads.	R	R	Total number of concurrent threads that Sterling File Gateway is allowed to use. This number may be verified against the licensed number of threads. 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.	R	R	File System is more efficient. Default value: database

The following table shows the properties that control the above parameters:

Recommendation	Property
Increase the value of Sterling File Gateway.	fgRouteConcurrentSessionLimit
If you are processing very large files, increase the probe values to avoid timeout conditions.	• filegateway.bpCompletionProbes.2 • filegateway.bpCompletionSleepMsec.2
If you have a high volume of PGP traffic, you can improve your performance by specifying a group for the file gateway.	pgpCmdline2svcname
If you have very large files that will be processed by PGP, increase the value of the file gateway.	fgRoutePGPCmdLineSocketTimeout
If you have high volumes of FTP traffic, you can improve your performance by specifying a group.	filegateway.ftpClientAdapterName
Decrease the value of evaluation frequency.	• MailboxEvaluateAllAutomaticRules or • MailboxEvaluateAllAutomaticRulesSubMin
Suppress Duplicate Messages	mailbox.disallowDuplicateMessages=true
Increase the number of steps a business process must complete prior to returning to the queue.	noapp.AE_ExecuteCycle.#
Increase the time period that a business process can use a thread, before releasing it to be used for another business process.	noapp.AE_ExecuteCycleTime.#
Increase the number of concurrent threads.	noapp.MaxThreads
Set storage type.	N/A

EBICS Banking Server: Specific Recommendations

The EBICS Banking Server is installed on an instance of Sterling B2B Integrator, and shares many of the resources with the latter.

You should, therefore, tune your Sterling B2B Integrator installation first, and then perform the EBICS Banking Server-specific tuning. Be aware that the changes you make to the EBICS Banking Server can also affect the performance of Sterling B2B Integrator.

Download from Www.Somanuals.com. All Manuals Search And Download.

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

Recommendation	Comments
Allocation of additional threads to the HTTP Server Adapter	You can allocate additional threads to the HTTP Server adapter when the rate of concurrent requests to the EBICS Banking Server is moderately high. Complete the following steps to allocate additional threads to the HTTP Server adapter: 1. Modify the numOfmaxThread parameter in the `http.properties.in` file by setting numOfmaxThread 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 "View and Restore Performance Configuration Settings" on page 121. 2. Run the **`setupfiles.sh`** (**`setupfiles.cmd`** for Windows) script to apply the changes.
Reduction of the disk I/O amount	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: • handleEBICSRequest • EBICSOrderAuthorisationProcessing • EBICSOrderProcessing
Storage size	The database space required to store the files processed by the EBICS Banking Server is approximately 2.5 times the size of the transacted files. For more information about managing databases in Sterling B2B Integrator, refer to "Database management for Sterling B2B Integrator" on page 16.
Purge documents	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. 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. 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 in the online documentation library.

Database Management

For optimal performance in Sterling B2B Integrator, you must properly plan, configure and manage your database.

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:

Download from Www.Somanuals.com. All Manuals Search And Download.

- Selecting a certified database management server software and version. Refer to the *Sterling B2B 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 IBM Professional Services Project Manager or a IBM Sales Executive. This study starts with the Sterling B2B Integrator Server Sizing Questionnaire. IBM Performance Engineering creates a sizing document that provides information about the estimated processor, memory, and network requirements for the standard/baseline Sterling B2B 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	• 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

Goal	Strategy
Scalability and Performance	• 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 stripping (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
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-6	Excellent for read (similar to RAID-0). Potentially poor for write performance as parity calculations need to happen. The performance of RAID-6 and RAID-5 is about the same and dependent on the controllers.	Better. Can survive a double-disk failure. This gives it an edge over RAID-5. A failure of more than two disks creates an outage.	Excellent. Low disk utilization skew. Possible LUN utilization skew.	12
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

Download from Www.Somanuals.com. All Manuals Search And Download.

Technology	Scalability	Survivability	Maintainability	Num Disks
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 B2B Integrator

Sterling B2B Integrator uses a database server as a repository for transactional, reference, and history data that it generates and uses.

Refer to the system requirements for a list of supported database configurations.

This topic covers the following concepts:

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

JDBC connection pool overview

Sterling B2B Integrator uses internal connection pooling 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. Internal connection pooling improves performance by removing the need to go to the driver and creating and delete a new connection each time one is needed.

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 must 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 must establish 120 database connections.

Download from Www.Somanuals.com. All Manuals Search And Download.

With this implementation, the database reports any idle connections to the system, which is the expected behavior. JDBC properties can be tuned in the `jdbc.properties` file to delete 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, see the properties file documentation.

Implementation of connection pooling

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 B2B Integrator requests additional connections up to the maximum number defined by the **maxsize** property. When Sterling B2B Integrator finishes with a connection, the connection is returned to the pool of connections for later use.

If Sterling B2B Integrator requires connections beyond the maximum size limit for each database pool, and every pool is in use, Sterling B2B 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 that are defined by the **maxsize** property) are created as needed, and deleted when the calling process is completed. To detect this condition, search for warning messages such as connections exceeded in the `noapp.log` file. Performance drops if Sterling B2B Integrator runs in this range for a long time.

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

Summary

Each JVM can create connections for each of the pools that are 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 buffer size limit. Connections beyond **maxsize** and below the **buffersize** are not reused. Each connection is created when needed, and deleted when no longer needed by the calling process. This method is expensive and harms performance if Sterling B2B 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, see "Dynamically Manage JDBC Pools" on page 20.

Download from Www.Somanuals.com. All Manuals Search And Download.

Configuring parameters for a database pool

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

Parameter	Description and recommended value
initsize	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 pool. Connections that are created from the **initsize** are part of the connections that exist in the pool for reuse, as needed. Since various JVMs can load the initial number of connections, but might not really need them, it is recommended that you do not set the **initsize** to a large value. Default value: 0 or 1 Recommended value: 1
maxsize	Maximum size of the connection pool that pertains to a database pool. After a connection is created up to this value, it is not deleted. It remains idle until needed.
Buffersize	Maximum number of connections that can be created. Connections that are created beyond the **maxsize** value are created and deleted as needed. After 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 is exhausted.
maxConn	This parameter is not used in Sterling B2B 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 the `noapp.log` file to ensure that you are not close to the maximum size of the connection pool (**maxsize**). Monitor the usage over a period of time and observe if the usage is increasing. If the usage limits exceed 25% of the settings that are recommended by the tuning wizard and the demand for connections is increasing, contact the IBM Professional Services onsite consultant or IBM 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 to be available when the system needs them, which means that connections are often shown as being idle. This method 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

Download from Www.Somanuals.com. All Manuals Search And Download.

- Placement of Indexes
- Sterling B2B Integrator Database Tables

Placement of Schema Table Objects

The Sterling B2B 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 need to review, and possibly modify, these statements before production.

Placement of Indexes

The DDL statements create a minimal set of indexes for general use. You might need to create more indexes to suit your business needs. Contact IBM Professional Services or IBM Customer Support to create more indexes. You should also maintain a list of the indexes added and ensure that they are present after you upgrade Sterling B2B Integrator and monitor its usage. These indexes may or may not be added to the Sterling B2B Integrator and is specific to your operating environment.

You might want to create more table spaces for storing indexes in addition to those table spaces for the base tables. Table spaces should have multiple data files for I/O load balancing.

Important: Indexes can also be separated out into different storage areas. This action 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 IBM Professional Services to prevent upgrade issue in the future.

Database tables that can substantially grow

The following table lists some of the tables that are defined in Sterling B2B Integrator experience substantial input, output, and rapid growth.

Table name	Table function
DATA_TABLE	These tables hold the binary payload of documents that are used for processing within Sterling B2B Integrator.
• DOCUMENT • CORRELATION_SET	Contain document metadata that is used for searchability and enablement of various document processing tasks.
WORKFLOW_CONTEXT	Contains step status and business process flow information.
ARCHIVE_INFO	Holds lifespan information that pertains to all business processes and document storage in the system.

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

Important: Move these tables only in coordination with IBM Professional services. When patches or upgrades that need to re-create a table are installed, the installation attempts to create them in the default table space.

Download from Www.Somanuals.com. All Manuals Search And Download.

Controlling entries in the CORRELATION_SET table

You can control when correlation entries for different activities of Sterling B2B Integrator are written to the CORRELATION_SET table. Reducing the number of records in the CORRELATION_SET table improves the performance of Sterling B2B Integrator.

The `doc_tracking.properties` file includes properties that enable or disable correlation entries for the following activities. The properties are in parentheses.

- Mailbox (**`makeCorrelationForMailbox`**)
- FTP (**`makeCorrelationForFTP`**)
- SFTP (**`makeCorrelationForSFTP`**)
- System tracking extension (**`makeTrackingExtensions`**)

The default value for each property is `true`, which means that correlation entries are written to the CORRELATION_SET table.

Attention: If `makeTrackingExtensions=false`, then no mailbox correlation entries are created, even if `makeCorrelationForMailbox=true`.

The `doc_tracking.properties` file is in the *`install_dir`*`\install\properties` directory.

Dynamically Manage JDBC Pools

Sterling B2B 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

This procedure shows how to add JDBC pools using the customer_overides.properties file.

Procedure

1. In the customer_overrides.properties file, specify the database server name, port number, database/catalog name, user ID, and password. To encrypt your database password, use the encrypt_string.sh or encrypt_string.cmd utility in the bin directory. Then place the encrypted password, prefixed by an encryption indicator, in the customer_overrides.properties file.
2. To update Sterling B2B Integrator with this new pool information, restart the application.

Modifying 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

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:
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, customer_overrides.properties files and system_overrides.properties.
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: You can now change database passwords on a scheduled basis in Sterling B2B 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 B2B Integrator when to start using that credential for new connections. This applies primarily to external pools, although Sterling B2B 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>	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: • 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

Download from Www.Somanuals.com. All Manuals Search And Download.

newDBPool.effective.1 = <The date for the new password starts to take affect>	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: • 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: 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.

Oracle init parameter configuration checklist

The Oracle init parameters have mandatory and recommended settings for Sterling B2B Integrator performance with an Oracle database.

Mandatory Oracle init parameters:

The Oracle init parameters have mandatory settings for Sterling B2B Integrator performance with an Oracle database.

Parameter	Mandatory value
cursor_sharing	Exact
Character set	AL32UTF8

Download from Www.Somanuals.com. All Manuals Search And Download.

Recommended Oracle init parameters:

The Oracle init parameters have recommended settings for Sterling B2B Integrator performance with an Oracle database.

Parameter	Recommended value
processes	Must be greater than the number of connections that are required by Sterling B2B Integrator (sum of transactional or local and **NoTrans** pools in the `jdbc.properties` file), and operational management tools.
• sga_max_size • sga_target • pga_aggregate_target	1 GB to *n* 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, see "SGA pool sizing" on page 25.
timed_statistics	True
optimizer_mode	All_rows
open_cursors	2000 or higher if prepared statement caching is to be used.

Oracle init parameter descriptions:

The Oracle init parameters have mandatory and recommended settings for the performance of Sterling B2B Integrator with an Oracle database.

Parameter	Description
Number of processes	Maximum number of processes that the Oracle server can create. Each inbound connection from a client requires an available process on the Oracle serverand internal processes that run in the Oracle server itself. This setting needs to be set high enough to accommodate the expected peak connections from Sterling B2B Integrator as well as additional processes for internal server processes plus the possible usage of buffer connections from the Sterling B2B Integrator pools. Note that the expected peak connections from Sterling B2B Integrator are per node, and this number should be multiplied by the number of nodes in a cluster. **Important:** Exceeding the allocated processes can destabilize the entire Oracle server and cause unpredictable issues. Normally, the Sterling B2B Integrator logs a report if this has occurred, by logging an "ORA-00020 maximum number of processes exceeded" error.

Download from Www.Somanuals.com. All Manuals Search And Download.

Parameter	Description
open_cursors	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. 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. A simple rule of thumb is to set this to four times the number of processes, four times the number of nodes of Sterling B2B Integrator, plus an additional 10%. If the number of cursors is exceeded, Sterling B2B Integrator logs "ORA-01000: maximum open cursors exceeded."
Character Set	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. AL32UTF8 is the preferred setting for Oracle database.
cursor_sharing	Controls whether SQL sent from a client is reparsed each time, or if the server can reuse ("share") the complied plan. Sterling B2B Integrator requires the setting to be "exact" for both performance and data integrity.
sga_max_size	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. 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.
sga_max_target	Target value that Oracle uses for automatic allocation of SGA pool resources. Recommended setting is equal to the sga_max_size. The exception is for extremely large databases, where this may need to be adjusted manually.
pga_aggregate_target	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. Recommended setting is a non-zero value. This allows the Oracle instance to size the SQL working areas automatically as needed.
timed_statistics	Controls whether database statistics for particular times are logged by Oracle. This information may be useful to monitor system or application performance. 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

Parameter	Description
optimizer_mode	Controls the mode that the optimizer uses to select the proper execution path for fulfilling query results, among other functions. Setting this to an incorrect mode can dramatically affect the overall performance of the server. Recommended setting for Sterling B2B Integrator: all_rows

SGA pool sizing:

As an alternative to manually sizing the main SGA pools in an Oracle database, IBM recommends that you use Automatic Shared Memory Management (ASMM).

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:

Download from Www.Somanuals.com. All Manuals Search And Download.

```
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.

Rollback Transaction Errors

If Sterling B2B Integrator is running with Oracle and is under load, the **snap shot too old** SQL error may be found in the wf.log file. The SQL error can be resolved with the following:

- Tune the Oracle database by increasing undo_retention and redo table space.
- If the SQL error is from WFReportService (from the stack trace in wf.log), change the `useNewStateAndStatusLogic` property in jdbc.properties and jdbc.properties.in to **false**. This will prevent the **snap shot too old** SQL error, but it might cause the Recover process to run longer.

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.

Download from Www.Somanuals.com. All Manuals Search And Download.

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 discontiguous, 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:

Download from Www.Somanuals.com. All Manuals Search And Download.

- 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.

I/O Sub-System Response Time

Sterling B2B Integrator is an On-Line Transaction Processing (OLTP) application. As an OLTP application, database response time to the I/O sub-system needs to be in the single digit range even during the peak periods. The database response time to the I/O sub-system should be less than:

- 5 ms for logs. 1ms or better is recommended.
- 10 ms or better for data. 5ms is recommended.

Monitoring with the AWR Report

Important: The information in an Automatic Workload Repository (AWR) report needs to be considered in relation to Sterling B2B 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 B2B 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:

Download from Www.Somanuals.com. All Manuals Search And Download.

```
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 B2B 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 B2B 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 Elapsd %:	233.33	% Non-Parse CPU:	90.82

The following figure shows the buffer pool statistics.

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B 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 B2B 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
PX Deq Credit: send blkd	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.

Download from Www.Somanuals.com. All Manuals Search And Download.

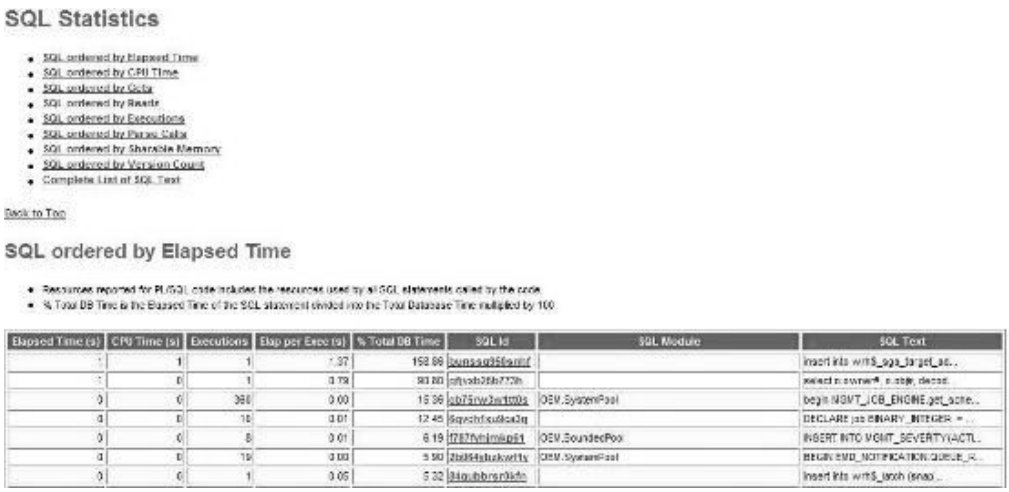

SQL Statistics

- 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 Parse Calls
- SQL ordered by Sharable 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	158.86	bvunsag856amf		insert into wmS_aga_target_ac...
1	0	1	0.79	93.80	qfvsh28b772h		select o.owner#, o.obj#, decod...
0	0	360	0.00	19.96	qb75rw3w1tm0s	OEM.SystemPool	begin MGMT_JOB_ENGINE.get_scre...
0	0	10	0.01	12.45	6qvchfxu6ca2g		DECLARE job BINARY_INTEGER := ...
0	0	8	0.01	8.19	f787fvtrjm4g1	OEM.BoundedPool	INSERT INTO MGMT_SEVERITY(ACTL...
0	0	10	0.00	5.50	2bd84duukwt1y	OEM.SystemPool	BEGIN EMD_NOTIFICATION.QUEUE_R...
0	0	1	0.05	5.32	94qubbrsr0kfn		insert into wmS_lanch (snap...

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 one time every day). You can refresh statistics by using the default job that is 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 with sufficient changes. Oracle bypasses statistics generation for tables that did not change 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 B2B Integrator schema owner.

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

Speak with the Sterling B2B Integrator Administrator to find out more about document lifespans. Generally, updating statistics one time every 24 hours is sufficient.

Index Rebuilds

Index rebuilds are also recommended for the Sterling B2B Integrator schema when the data in the tables changes by approximately 20%. However, the rebuild must be validated by looking at the system performance as a whole. Online index rebuilds, which cause minimal impact to the system, is possible.

Download from Www.Somanuals.com. All Manuals Search And Download.

The Sterling B2B Integrator system administrator must work with the DBA to identify the indexes that must be rebuilt. You must always rebuild indexes at schema level.

Sterling B2B Integrator is a high-volume OLTP application, and rarely, there might be times when you might need to shrink the tables or rebuild indexes. Oracle documentation provides details about the following tasks:

- Shrinking database segments online
  - http://docs.oracle.com/cd/B28359_01/server.111/b28310/schema003.htm#ADMIN10161 (Oracle Database 11g Release 1)
  - http://docs.oracle.com/cd/E18283_01/server.112/e17120/schema003.htm#ADMIN10161 (Oracle Database 11g Release 2)
  - http://docs.oracle.com/database/121/ADMIN/schema.htm#CBBBIADA (Oracle Database 12c Release 1)
- Rebuilding an existing index
  - http://docs.oracle.com/cd/B28359_01/server.111/b28310/indexes004.htm#ADMIN11734 (Oracle Database 11g Release 1)
  - http://docs.oracle.com/cd/E18283_01/server.112/e17120/indexes004.htm#insertedID2 (Oracle Database 11g Release 2)
  - http://docs.oracle.com/database/121/ADMIN/indexes.htm#i1006864 (Oracle Database 12c Release 1)

Enable BLOB Caching

By default, Oracle database configures itself with the BLOB caching disabled when you choose Oracle as your database during Sterling B2B 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.

Using Oracle Real Application Clusters (RAC) for High Availability

Oracle Real Application Clusters (RAC) ensures high availability by supporting the 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

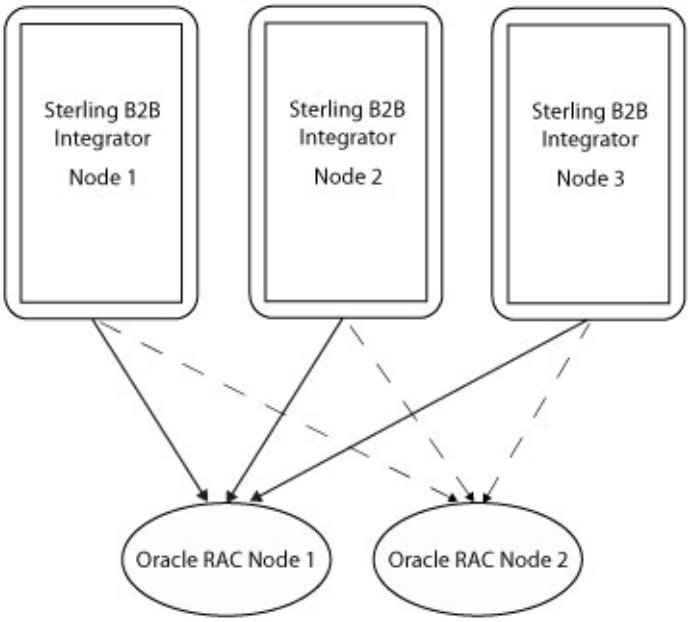

The following table lists the features and benefits of Oracle RAC when used with Sterling B2B Integrator:

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.
SCAN (Single Client Access Name)	Available in Oracle RAC 11g R2, SCAN allows you to set up RAC failover without specifying nodes. Using traditional RAC in Sterling B2B Integrator, you configure the JDBC URL to name an initial node and a failover node. With SCAN, you configure the JDBC URL so that SCAN selects the initial node and fails over to the other available node. A benefit to using SCAN over traditional RAC is that you do not need to update the URL when a node name changes. There is, however, a potential performance degradation as SCAN searches for an available node.

Download from Www.Somanuals.com. All Manuals Search And Download.

Features	Benefits

High availability (Failover)	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.
	• **Initial connection time failover**
	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.
	Initial connection time failover can be configured on the JDBC multi-instance URL on the client side by setting this parameter:
	FAILOVER=ON
	Sterling B2B Integrator supports two-node Oracle RAC setup.
	To establish initial time connection failover in Sterling B2B Integrator, perform the following tasks:
	1. 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.
	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.
	`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>`
	2. From the *install_dir*/install/bin (*install_dir*\install\bin for Windows) directory, run (UNIX or Linux) **setupfiles.sh** or (Windows) **setupfiles.cmd**.
	3. Restart the ASI and container JVMs to apply the changes.

Download from Www.Somanuals.com. All Manuals Search And Download.

Features	Benefits

High availability (Failover), continued	• **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	You can increase the number of visibility queues and event input queue capacity in customer_overrides.properties to ensure that adapters are responsive.
	To increase the default queue size and the thread counts, perform the following tasks:
	1. Set the following values in the ASI node and in customer_overrides.properties.in file in containers:
	`dmivisibility.number_visibility_threads=6`
	`dmivisibility.event_input_queue_capacity=3072`
	2. From the *install_dir*/install/bin (*install_dir*\install\bin for Windows) directory, run (UNIX or Linux) **`setupfiles.sh`** or (Windows) **`setupfiles.cmd`**.
	3. Restart the ASI and container JVMs to apply the changes.

Fast Connection Failover (FCF)	Sterling B2B Integratordoes not support Fast Connection Failover (FCF).
Transparent Application Failover (TAF)	Sterling B2B Integratordoes 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.

Mandatory settings for IBM DB2 registry variables

Mandatory IBM DB2 registry values are critical for IBM DB2 performance with Sterling B2B Integrator.

Download from Www.Somanuals.com. All Manuals Search And Download.

Variable	Mandatory value
DB2_SKIPDELETED	ON 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. 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. 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. Recommended value: ON
DB2_SKIPINSERTED	ON 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.

Recommended settings for IBM DB2 registry variables

IBM DB2 registry values include recommended settings for IBM DB2 performance with Sterling B2B Integrator.

Variable	Recommended value
DB2_USE_ALTERNATE_ PAGE_CLEANING	ON
DB2_ EVALUNCOMMITTED	ON 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. 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

Variable	Recommended value
DB2_PARALLEL_IO	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. Set the DB2_PARALLEL_IO variable if you have implemented containers on stripped 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. In this example, prefetches to the four-container tablespace mentioned above are performed in 24 extent-sized prefetch requests. 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. If you have high prefetch_wait_time values, the application might be waiting for pages.
DB2_NUM_ CKPW_DAEMONS	0 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. 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 `ps aux | grep db2ckpwd` command and look at the fifth and sixth columns of output. DB2 workaround exists for this problem. Set the following registry variable, and recycle this instance: `db2set DB2_NUM_CKPW_DAEMONS=0` Also see IBM APAR IY78341.
DB2LOCK_TO_RB	STATEMENT

Download from Www.Somanuals.com. All Manuals Search And Download.

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

Download from Www.Somanuals.com. All Manuals Search And Download.

Parameter	Mandatory or Recommended	Recommended Value
MAXAGENTS	Recommended	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. Recommend value: Must be greater than the number of connections needed by Sterling B2B Integrator (that is, the sum of transactional (local) plus NoTrans pools in jdbc.properties file) plus the connections needed by operational management tools. 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. 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.

Estimation Guidelines for Number of Connections

You can roughly estimate the number of concurrent user connections required for Sterling B2B Integrator with this formula:

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)

Download from Www.Somanuals.com. All Manuals Search And Download.

Variable	Description
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.

Mandatory settings for DB CFG parameters

For optimal performance, certain parameters and values are mandatory for DB2 9.x.

Parameter	Mandatory value
Database Code Set	UTF-8

Recommended settings for DB CFG parameters

For optimal performance, follow the recommended settings for DB2 9.*x*.

The parameters and values described in the following table are recommended:

Parameter	Recommended value
SELF_TUNING_MEM	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	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	Allows STMM to dynamically manage memory allocations. Recommended value: AUTOMATIC
MAXLOCKS	Allows STMM to dynamically manage memory allocations. Recommended values: AUTOMATIC If AUTOMATIC is not supported: 100
PCKCACHESZ	Allows STMM to dynamically manage memory allocations. Recommended value: AUTOMATIC
SHEAPTHRES_SHR	
SORTHEAP	

Download from Www.Somanuals.com. All Manuals Search And Download.

Parameter	Recommended value
NUM_IOCLEANERS	Recommended value: AUTOMATIC
NUM_IOSERVERS	Recommended value: AUTOMATIC
DFT_PREFETCH_SZ	Recommended value: AUTOMATIC
MAXAPPLS	Recommended value: AUTOMATIC
APPLHEAPSZ	Recommended value: AUTOMATIC
APPL_MEMORY	Recommended value: AUTOMATIC
APP_CTL_HEAP_SZ	512 **Note:** APP_CTL_HEAP_SZ is deprecated in DB2 9.5.
DBHEAP	Amount required depends on the amount of memory available and the traffic volume. Recommended value: AUTOMATIC **Note:** The default DBHEAP value is inadequate.
LOGFILSIZ	Refer to the topic Log Configuration. Recommended value: 65536 if configuring 20 transaction logs of 256 MB (65536 4 K pages)
LOGPRIMARY	Number of primary transaction logs. Refer to the topic Log Configuration. Recommended value: 40 or more
LOGSECOND	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	Refer to the topic Log Configuration. Recommended value: LOGPRIMARY - Safety buffer
DFT_DEGREE	Sets the default degree of parallelism for intrapartition parallelism. In general, online transactional applications such as Sterling B2B 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 value: STATEMENT

Download from Www.Somanuals.com. All Manuals Search And Download.

Recommended Settings for DB2 9.7

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

Feature	Description	Value
Currently Committed Semantics	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. 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. 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. 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. 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.	Recommended value: ON

Download from Www.Somanuals.com. All Manuals Search And Download.

Feature	Description	Value
	Limitations The following list provides limitations that apply to the currently committed semantics: • 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 does not use currently committed semantics: – Catalog table scans – Referential integrity constraint enforcement scans – LONG VARCHAR or LONG VARGRAPHIC column reference scans – Range-clustered table (RCT) scans – Spatial or extended index scans	
Statement concentrator	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. 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.	You can enable statement concentrator by running the following SQL statement: **db2 update db cfg for <db-alias> using stmt_conc literals**

Download from Www.Somanuals.com. All Manuals Search And Download.

Feature	Description	Value
Inline LOBs	Sterling B2B Integrator extensively uses large objects (LOBs). These LOBs are usually a few kilobytes in size. 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. 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. 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. 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.	Enable LOB inlining through the INLINE LENGTH option on the CREATE TABLE statement or the ALTER TABLE statement.

Recommended Settings for DB2 for Linux on System z (5.2.4.0 or later)

IBM recommends special settings for DB2 when using Sterling B2B Integrator 5.2.4.0 or later on the Linux operating system on System z.

When you are running Sterling B2B Integrator 5.2.4.0 or later using the DB2 database on the Linux operating system on System z (or zLinux), IBM recommends the following settings:

- SCSI over FCP (fiber channel protocol) disks for data and transaction logs.
- RAID 5 for data logs and RAID 10 or solid state drives (SSDs) for transaction logs. All volumes should be striped across as many disks as possible. With this setup, DB2 can be configured to use direct input and output.

An alternative setup is the use of separate extended count key data (ECKD) disks with PAV (parallel access volumes) or HyperPAV for data and transaction logs. The PAV and HyperPAV features are available on the IBM DS8000 series of IBM SANs. This setup tries to overcome the limitations of the typical ECKD disk setup. All ECKD volumes should be striped.

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 B2B Integrator installation creates tables and indexes. The tables use different page sizes: 4K, 8K, and 16K. You should have a tablespace to

Download from Www.Somanuals.com. All Manuals Search And Download.

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.

Download from Www.Somanuals.com. All Manuals Search And Download.

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:

  `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.

Download from Www.Somanuals.com. All Manuals Search And Download.

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.

I/O Sub-System Response Time

Sterling B2B Integrator is an On-Line Transaction Processing (OLTP) application. As an OLTP application, database response time to the I/O sub-system needs to be in the single digit range even during the peak periods. The database response time to the I/O sub-system should be less than:

- 5 ms for logs. 1ms or better is recommended.
- 10 ms or better for data. 5ms is recommended.

Log Configuration

This topic provides information about the following logs:

- LOGFILSIZ, LOGPRIMARY, LOGSECOND

Download from Www.Somanuals.com. All Manuals Search And Download.

- NUM_LOG_SPAN

LOGFILSIZ, LOGPRIMARY, LOGSECOND

At a minimum, configure 20 transaction logs (LOGPRIMARY=20) of 256 MB (LOGFILSIZ=65536 4K-pages) for Sterling B2B 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 IBM Sterling 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

Download from Www.Somanuals.com. All Manuals Search And Download.

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**

View snapshot data	Switches must be turned on to view data.
	GET SNAPSHOT FOR options
	Options:
	`bufferpools on database_name`
	`locks on database_name`
	`dynamic sql on database_name`
	`tables on database_name`
	`applications on database_name`
	`tablespaces on database_name`
	`database on database_name`

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.

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B 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	Run this command:
	
	`db2 -v create event monitor $MON for deadlocks with details write to file $OUTDIR buffersize 64 nonblocked`
	
	`db2 -v set event monitor $MON state = 1`
	
	Options:
	
	`MON` – Monitor name (for example, `DLMON`)
	
	`OUTDIR` – directory to store deadlock information

Determine that a deadlock has occurred	View data in the db2diag.log or in the Sterling B2B Integrator logs

Release a deadlock by flushing buffers	Run the following command:
	
	`db2 flush event monitor $MON`
	
	`db2evmon -path $OUTDIR`
	
	The flush ensures that the deadlock records in the buffers are written out. The `db2evmon` command formats the deadlock information.

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.

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B 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 section provides information about configuring and monitoring Microsoft® SQL Server® 2005, Microsoft® SQL Server® 2008, and Microsoft® SQL Server® 2012.

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.

Mandatory settings for Microsoft SQL Server

The default collation of Microsoft SQL Server must match the collation for the Sterling B2B Integrator database to prevent collation conversions.

Download from Www.Somanuals.com. All Manuals Search And Download.

The *tempdb* database that is used by Microsoft SQL Server must be created with the same collation as the default collation of Microsoft SQL Server. The Microsoft SQL Server uses the tempdb database for results that are too large to fit in memory.

If the collations of the tempdb database and the Sterling B2B Integrator database differ, the database engine must convert from the Sterling B2B Integrator collation to the tempdb collation, and then back again before it sends the results to the Sterling B2B Integrator server. These conversions might lead to severe performance issues.

The collation that is required for the Sterling B2B Integrator database is a collation that most closely matches the character set used by Java. By using this collation, you can avoid character data conversions before the data is stored in the database tables. Use the mandatory parameter that is described in the following table when you configure the collation setting:

Parameter	Value
Database Collation	SQL_Latin1_General_CP850_Bin

Additionally, you must perform these tasks:

- Allow Microsoft SQL Server to manage memory dynamically (default).
- Disable any antivirus software that is running on the Microsoft SQL Server data, transaction log, and binary files directory.

Recommended instance-specific settings for Microsoft SQL Server

The use of the Microsoft SQL Server database with Sterling B2B Integrator includes some recommended instance-specific settings for the database.

Parameter	Value
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).
Min server memory (MB)	0

Recommended database-specific settings for Microsoft SQL Server

The use of the Microsoft SQL Server database with Sterling B2B Integrator includes some recommended database-specific settings.

Parameter	Value	Notes®
IsAutoCreateStatistics	True	This parameter can be set to False if you have a maintenance plan.
IsAutoUpdateStatistics	True	This parameter can be set to False if you have a maintenance plan.

Download from Www.Somanuals.com. All Manuals Search And Download.

Parameter	Value	Notes®
Page_verify_option	Checksum (Microsoft SQL Server default)	
READ_COMMITTED_ SNAPSHOT	On	

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:

- http://msdn.microsoft.com/en-us/library/ms190673(SQL.90).aspx (Microsoft SQL Server 2005)
- https://technet.microsoft.com/en-us/library/ms190673%28v=sql.105%29.aspx (Microsoft SQL Server 2008)

This feature was removed in Microsoft SQL Server 2012. For more information, see https://support.microsoft.com/en-us/kb/2663912.

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

This feature is available in Microsoft SQL Server 2005 and later versions,

This feature 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.

Recommended Settings for Microsoft SQL Server 2008 and 2012

Recommended settings for Microsoft SQL Server 2008 and Microsoft SQL Server 2012 are provided.

These are shown in the following table:

Download from Www.Somanuals.com. All Manuals Search And Download.

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 consolidates this information in detail graphically by running the recently executed processes. For more information about Activity Monitor, refer to the Microsoft Developer Network Web site: • https://technet.microsoft.com/en-us/library/cc879320%28v=sql.105%29.aspx (Microsoft SQL Server 2008) • https://msdn.microsoft.com/en-us/library/hh212951%28v=sql.110%29.aspx (Microsoft SQL Server 2012)

Download from Www.Somanuals.com. All Manuals Search And Download.

Feature	Description
Data Compression	The following list describes the two type of data compression supported by Microsoft SQL Server: • **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. The amount of compression achieved is dependent on the data types and the data contained in the database. Compression, row or page, can be applied to a table or an index in an online mode without interrupting the availability of Sterling B2B 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. Run the following SQL query to determine how compressing an object may affect its size: `sp_estimate_data_compression_savings` The following Sterling B2B Integrator tables may be the most likely candidates for compression: • DATA_TABLE • TRANS_DATA • CORRELATION_SET • WORKFLOW_CONTEXT • ARCHIVE_INFO For more information about implementing row and page compression, refer to the Microsoft Developer Network Web site: • https://msdn.microsoft.com/en-us/library/cc280576%28v=sql.105%29.aspx (Row compression) • https://msdn.microsoft.com/en-us/library/cc280464%28v=sql.105%29.aspx (Page compression) • https://msdn.microsoft.com/en-us/library/cc280449%28v=sql.110%29.aspx (Data compression)

Download from Www.Somanuals.com. All Manuals Search And Download.

Feature	Description
Hot Add CPU and Hot Add Memory	Hot Add CPU enables you to add CPUs dynamically to the servers without shutting down the server or limiting client connections. Hot Add Memory enables you to add physical memory dynamically without restarting the server. For more information about dynamically adding CPU and physical memory, refer to the Microsoft Developer Network Web site, which can be accessed from: • http://msdn.microsoft.com/en-us/library/bb964703.aspx (Hot Add CPU) • http://msdn.microsoft.com/en-us/library/ms175490.aspx (Hot Add Memory) **Note:** The above pages are for Microsoft SQL Server 2008. However, Hot Add CPU and Hot Add Memory are also supported in Microsoft SQL Server 2012. For more information, see the High Availability section at https://msdn.microsoft.com/en-us/library/bb630282%28v=sql.110%29.aspx.
Extended Events	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. For more information about extended events, refer to the Microsoft Developer Network Web site: • http://msdn.microsoft.com/en-us/library/bb630354.aspx (Microsoft SQL Server 2008) • https://msdn.microsoft.com/en-us/library/bb630282%28v=sql.110%29.aspx (Microsoft SQL Server 2012)

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 B2B 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 IBM 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

- 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 B2B Integrator tables, refer to the "Schema Objects" and "Sterling B2B Integrator Database Tables" sections in the topic “Database management for Sterling B2B Integrator” on page 16.

I/O Sub-System Response Time

Sterling B2B Integrator is an On-Line Transaction Processing (OLTP) application. As an OLTP application, database response time to the I/O sub-system needs to be in the single digit range even during the peak periods. The database response time to the I/O sub-system should be less than:

- 5 ms for logs. 1ms or better is recommended.
- 10 ms or better for data. 5ms is recommended.

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.

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

Download from Www.Somanuals.com. All Manuals Search And Download.

- 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: https://technet.microsoft.com/en-us/library/ms181091%28v=sql.105%29.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.

For more information about Troubleshooting Deadlocks in Microsoft SQL Server, 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

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.

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.

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B Integrator database regularly using either the simple model or the full recovery model
- Update the statistics on all Sterling B2B Integrator tables and associated indexes
- Rebuild or reorganize indexes on Sterling B2B Integrator tables
- Run database consistency checks

For more information about MS SQL Maintenance Plans, see the Microsoft Developer Network Web site: https://msdn.microsoft.com/en-us/library/ms187658%28v=sql.105%29.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: https://msdn.microsoft.com/en-us/library/ms191164%28v=sql.105%29.aspx or https://msdn.microsoft.com/en-us/library/ms186289%28v=sql.105%29.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: https://msdn.microsoft.com/en-us/library/ms186289%28v=sql.105%29.aspx.

For more information about Transaction Log Backups, refer to the Microsoft Developer Network Web site, which can be accessed from: https://msdn.microsoft.com/en-us/library/f4a44a35-0f44-4a42-91d5-d73ac658a3b0%28v=sql.105%29 .

Note: The Microsoft Developer Network Web site links provided access pages for Microsoft SQL Server 2008 R2. To view pages for other versions, such as Microsoft SQL Server 2012, use the **Other Versions** feature on the viewed page.

Download from Www.Somanuals.com. All Manuals Search And Download.

Microsoft SQL Server Index, Table Statistics, and Index Rebuilds

Your maintenance plan should include updating 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 B2B 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 B2B 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.

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.

Windows Defragmentation

You must run Windows defragmentation on disks with SQL Server data files and transaction logs once a month.

Download from Www.Somanuals.com. All Manuals Search And Download.

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.

Microsoft SQL Server Tips

Additional tips are provided pertaining to Microsoft SQL Server in the context of Sterling B2B Integrator.

- When using Sterling B2B 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 B2B Integrator and Microsoft SQL Server where they are on separate physical servers. Sterling B2B Integrator uses JDBC connection pool, and idle connections are typical.
- It is recommended that you run Microsoft SQL Server and Sterling B2B 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 B2B 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 trace flags, see https://msdn.microsoft.com/en-us/library/ms188396%28v=sql.105%29.aspx

    For more information about Detecting and Ending Deadlocks in Microsoft SQL Server 2005 and Microsoft SQL Server 2008, refer to the Microsoft Developer Network Web site, which can be accessed from: http://msdn.microsoft.com/en-us/library/ms178104(SQL.90).aspx.

    In Microsoft SQL Server 2012, the System Health session detects deadlocks. See https://technet.microsoft.com/en-us/library/ff877955%28v=sql.110%29.aspx for more information.

Java Virtual Machine

Java Virtual Machine (JVM) is a platform-independent programming language that converts Java bytecode into machine language and executes it.

Download from Www.Somanuals.com. All Manuals Search And Download.

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.

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

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

Note: If your environment handles large concurrent communications traffic, it is recommended to increase the value of the **gmm.maxAllocation** property in the `perimeter.properties` file, along with increasing the maximum heap size of your JVM. In such a scenario, set **gmm.maxAllocation** to a value greater than 384 MB. For more information about the **gmm.maxAllocation** property, refer to the inline comments in the `perimeter.properties` file.

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:

Download from Www.Somanuals.com. All Manuals Search And Download.

```
<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 B2B Integrator on a Sun Solaris operating system running on Sun UltraSPARC processor-based servers and on a Windows operating system.

Download from Www.Somanuals.com. All Manuals Search And Download.

The HP HotSpot JVM is used when you deploy Sterling B2B 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.

To enable JVM verbose garbage collection, refer to the documentation on JVM parameters for the server in “Edit Performance Configuration Settings” on page 123.

To enable JVM verbose garbage collection for the container JVM, refer to the documentation on JVM parameters for container nodes in “Edit Performance Configuration Settings” on page 123.

For information on the default parameters for the HotSpot JVM, refer to “HotSpot JVM Default Parameters” on page 66.

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.

Download from Www.Somanuals.com. All Manuals Search And Download.

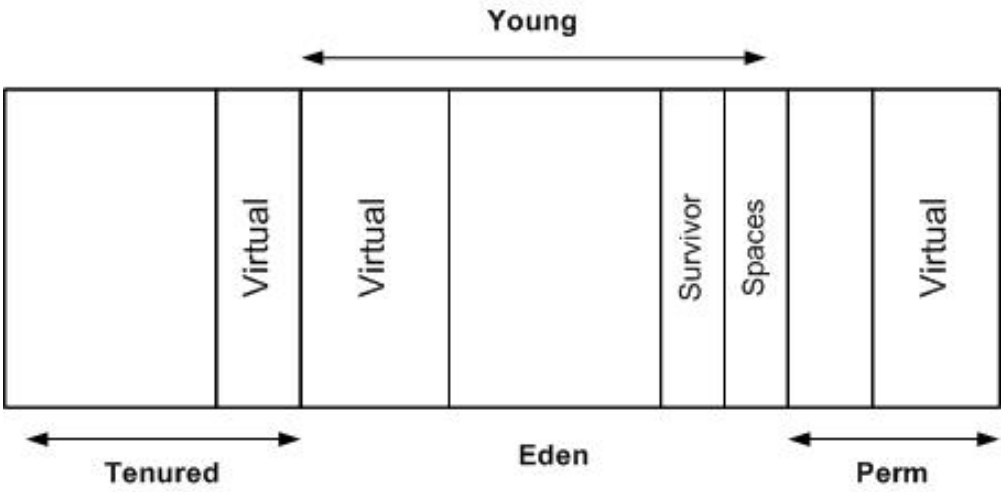

At initialization, a maximum address space is virtually reserved, but not allocated, to physical memory unless it is needed. The complete address space reserved for object memory can be divided into young and tenured (old) generations.

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

Download from Www.Somanuals.com. All Manuals Search And Download.

Tool Name	For Additional information, Refer To
IBM Pattern Modeling and Analysis Tool for Java Garbage Collector	http://www-01.ibm.com/support/docview.wss?uid=swg27015310
visualgc	http://java.sun.com/performance/jvmstat/visualgc.html **Note:** Visaulgc 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 B2B 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 B2B Integrator 5.2 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, on Windows Server 2003, the –d32 option is default as it uses a 32-bit JDK.

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B 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: • Solaris = -server (optimizing compiler) • HP-UX = -server (optimizing compiler)
`-Xmx` If this parameter is tuned correctly, it can: • 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 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.	Refer to "Edit Performance Configuration Settings" on page 123 for the default values for the: • Maximum heap size for the server JVM (MAX_HEAP) • Maximum heap size for the container JVM (MAX_HEAP_CONTAINER)

Download from Www.Somanuals.com. All Manuals Search And Download.

Parameter / Description	Default Value
`-XX:+DisableExplicitGC` Disables the explicit garbage collection calls that are caused by invoking System.gc() inside the application. 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.	For both 32-bit/64-bit noapp and container JVMs: • Windows = -XX:+DisableExplicitGC • Solaris = -XX:+DisableExplicitGC • HP-UX = -XX:+DisableExplicitGC
`-XX:NewSize= and -Xmn<Size>=` –XX:NewSize controls the minimum young generation size in a heap, and the –Xmn sets the size of the young generation heap. 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. • -XX:NewSize = (0.33333* value of –Xms) • -Xmn = (0.33333 * value of -Xmx)	Refer to “Edit Performance Configuration Settings” on page 123 for the default values for the: • Initial new heap size for the server JVM (INIT_AGE) • Initial new heap size for the container JVM (INIT_AGE_CONTAINER) • Maximum new heap size for the server JVM (MAX_AGE) • Maximum new heap size for the container JVM (MAX_AGE_CONTAINER)
`-XX:MaxPermSize` 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 `Java.lang.OutOfMemoryError: PermGen space` error.	Refer to “Edit Performance Configuration Settings” on page 123 for the default values for the: • Maximum permanent generation heap size for the server JVM (MAX_PERM_SIZE) • Maximum permanent generation heap size for the container JVM (MAX_PERM_SIZE_CONTAINER)

Download from Www.Somanuals.com. All Manuals Search And Download.

Parameter / Description	Default Value
`-Xss` Determines the stack size for each thread in the JVM. 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 `java.lang.StackOverflowError` exception.	Refer to "Edit Performance Configuration Settings" on page 123 for the default values for the: • Maximum Java stack size for any thread for the server JVM (STATIC_STACK) • Maximum Java stack size for any thread for the container JVM (STATIC_STACK_CONTAINER)
`-XX:CompileThreshold` Determines the number of method invocations and branches before compiling. Setting a low value for this parameter will trigger the compilations of hot methods sooner.	For both 32-bit/64-bit noapp and container JVMs: • Windows = -XX:CompileThreshold=1000 • Solaris = -XX:CompileThreshold=1000
`-Xnocatch` Disables the Java catch-all signal handler. This option is used to get clean, native code stack traces.	For both noapp and container JVMs: HP-UX = -Xnocatch

Display HotSpot Statistics for HP JVM

To display HotSpot Statistics for the HP JVM, enable one of the following options:

- `-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]
  ```

- `-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 >
  ```

Display HotSpot Statistics for Sun JVM

To display HotSpot Statistics for the Sun JVM, enable the following option:

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


Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B 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 B2B 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)
```


Download from Www.Somanuals.com. All Manuals Search And Download.

```
  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.
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):

Download from Www.Somanuals.com. All Manuals Search And Download.

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. Under **View Dumps**, select the Name of the thread dump file and click **Go!** to view it. 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 B2B 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 B2B 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 B2B 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":
```


Download from Www.Somanuals.com. All Manuals Search And Download.

```
        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

Download from Www.Somanuals.com. All Manuals Search And Download.

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 JNInative method rather than the JVM code. In order to troubleshoot this error, use the OS utilities to find the actual cause.

Download from Www.Somanuals.com. All Manuals Search And Download.

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":
```


Download from Www.Somanuals.com. All Manuals Search And Download.

```
     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$1DisposeAction.run(Window.java:598)
- locked <0xf0c41ec8> (a java.util.Vector)
at java.awt.Window.doDispose(Window.java:617)
at java.awt.Window.dispose(Window.java:574)
at
javax.swing.SwingUtilities$SharedOwnerFrame.dispose(SwingUtilities.java:1743)
at
javax.swing.SwingUtilities$SharedOwnerFrame.windowClosed(SwingUtilities.java:172
2)
at java.awt.Window.processWindowEvent(Window.java:1173)
at javax.swing.JDialog.processWindowEvent(JDialog.java:407)
at java.awt.Window.processEvent(Window.java:1128)
102
at java.awt.Component.dispatchEventImpl(Component.java:3922)
at java.awt.Container.dispatchEventImpl(Container.java:2009)
at java.awt.Window.dispatchEventImpl(Window.java:1746)
at java.awt.Component.dispatchEvent(Component.java:3770)
at java.awt.EventQueue.dispatchEvent(EventQueue.java:463)
at
java.awt.EventDispatchThread.pumpOneEventForHierarchy(EventDispatchThread.java:2
14)
at
java.awt.EventDispatchThread.pumpEventsForHierarchy(EventDispatchThread.java:163
)
at java.awt.EventDispatchThread.pumpEvents(EventDispatchThread.java:157)
at java.awt.EventDispatchThread.pumpEvents(EventDispatchThread.java:149)
at java.awt.EventDispatchThread.run(EventDispatchThread.java:110)
"main":
at java.awt.Window.getOwnedWindows(Window.java:844)
- waiting to lock <0xf0c41ec8> (a java.util.Vector)
at
javax.swing.SwingUtilities$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.concurrency 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".

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B Integrator, HPROF can be invoked by adding the following option as a JVM parameter by using the Sterling B2B Integrator tuning wizard and adding the below option to the JVM argument suffix option for both noapp and container. For more details, please refer to "Edit Performance Configuration Settings" on page 123. Run the **setupfiles.sh** command (UNX/Linux) or the **setupfiles.cmd** command (Windows) and restart ASI and container JVMs.

```
-agentlib:hprof[=options]
```

Or

```
-Xrunhprof[:options]
```

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://bizsupport2.austin.hp.com/bc/docs/support/SupportManual/c02697864/c02697864.pdf

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:

Download from Www.Somanuals.com. All Manuals Search And Download.

http://bizsupport2.austin.hp.com/bc/docs/support/SupportManual/c02697864/c02697864.pdf

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://www.hp.com/go/hpjmeter

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 B2B 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 B2B Integrator using the IBM JVM.

For information about using the tuning wizard to set the IBM JVM parameters, see "Edit Performance Configuration Settings" on page 123.

Before You Begin Tuning Your IBM® JVM

Using the Performance Tuning Utility, Sterling B2B Integrator calculates the recommended settings based on the number of cores and the amount of physical memory that is being made available to Sterling B2B 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 IBM Professional Services. Contact your Sales Representative for more information about this.

IBM® JVM Default Parameters for Sterling B2B Integrator

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

Note: For information about using the tuning wizard to set the IBM JVM parameters, see "Edit Performance Configuration Settings" on page 123.

Download from Www.Somanuals.com. All Manuals Search And Download.

Parameter	Description
-Xmns	• 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 It is recommended to use the following formula to compute –Xmns value when the minimum heap size is modified. -Xmns = (0.33333 * value of -Xms)
-Xmnx	• 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 It is recommended to use the following formula to compute –Xmnx value when the maximum heap size is modified. -Xmnx = (0.33333 * value of -Xmx)
-Xjit:count	• Controls the compilation threshold of the JIT compiler • Value given to count causes the Java method to be compiled after n runs For 64-bit noapp JVM and 64-bit container JVM: • Default Value for Linux: 1000 • Default Value For AIX: 1000
-Xgcpolicy:gencon	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.

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,

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B Integrator also enables users to trigger a thread dump through the Sterling B2B 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 Taken page 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 Dumps download pop-up window, click the **Download** link to download the file to the local machine.

Download from Www.Somanuals.com. All Manuals Search And Download.

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 Dumps download 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/).

Download from Www.Somanuals.com. All Manuals Search And Download.

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.

To enable a heapdump, refer to the documentation on JVM parameters for the server in "Edit Performance Configuration Settings" on page 123.

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

Monitoring Operations

The Operations functions enable you to monitor the operational status of Sterling B2B Integrator, its components and services, current threads and system logging, and to troubleshoot system problems.

Managing System Logs

Sterling B2B Integrator comprises multiple components, including software applications such as archive tools, Web servers, and database servers. To monitor the activities of each component, Sterling B2B 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 B2B Integrator generates a log file only when a component runs. This in turn improves the performance of Sterling B2B Integrator.

Each open log file is closed once every 24 hours (usually at midnight), and a new file is created. When moving into or out of Daylight Savings Time (DST), you must stop and restart both Sterling B2B Integrator and its database to ensure that log 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 B2B 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 B2B 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

- `logs` refers to the primary log directory.
- `directory` refers to the subdirectory created when you start Sterling B2B 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 B2B 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 B2B Integrator, are written to the *install_dir/logs* directory.

Viewing Log File Contents

Users of Sterling B2B 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 B2B 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 B2B Integrator is installed. Open the log file in read-only mode using a text editor.

To view the current log file contents in Sterling B2B 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 B2B Integrator:

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B Integrator (such as WORKFLOW, OPS, or UTIL), the Sterling B2B Integrator

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B 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 B2B Integrator (Services Controller)
- What occurred (INFO_sdi_getName)
- Error code (0297)

Changing Log Settings

Sterling B2B 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 B2B Integrator Property Files* documentation.

Note: The customer override property file is not a part of the initial Sterling B2B 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: • true – Allow new logs to be created (Default) • false – Do not allow new logs to be created Example: `logService.newloggers=true`

Download from Www.Somanuals.com. All Manuals Search And Download.

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

The following table provides the log type name for each log file name defined by default in the log.properties file and extension files. 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 customer_overrides.properties file)
alerterlogger.log	alerterlogger
archive.log	archivelogger
Authentication.log	AuthenticationLogger
cdinterop.log	cdinteroplogger
cdinterop_cdjava.log	cdinteropcdjavalogger
ceuinterop.log	ceulogger
common3splogger.log	common3splogger
delete.log	deletelogger
ebXML.log	ebXMLlogger
EDIINT.log	EDIINTLogger
event.log	event
ftp.log	ftplogger
ftpclient.log	psftpclientlogger
http.log	httplogger
httpclient.log	httpclientlogger
jetty.log	jettylogger
lifecycle.log	lifecycleLogger
mailbox.log	mailboxlogger
mgmtdash.log	neo
noapp.log	noapplogger
ocsp.log	ocsplogger
oftp.log	oftplogger
ops_exe.log	opslogger
Perimeter.log	PSLogger
pipeline.log	pipelinelogger

Download from Www.Somanuals.com. All Manuals Search And Download.

File Name of Log (As Shown on the System Logs Screen)	Log Type (For the customer_overrides.properties file)
report.log	reportlogger
resourcemonitor.log	resourcemonitorlogger
mif.log	rnlogger
sap.log	saplogger
schedule.log	schedulelogger
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 B2B 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:

Download from Www.Somanuals.com. All Manuals Search And Download.

Option	Description
Location	Specifies the absolute path for the log file.
Rollover Interval	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: • 50000 lines • 100000 lines • 150000 lines • 200000 lines **Note:** Here, the term "lines" refers to logical lines (entries) and not physical lines. For example, following is a "line": `[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]`
Logging Level	Specifies the amount of details to log. Select one of the following values: • 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) **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.

4. Click **Save**.

Changing Log File Location

Sterling B2B 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

However, the following log files created by the Sterling B2B 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 B2B 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.
IBM Sterling Gentran:Server® for UNIX		
Data Adapter		Logs activities of the Sterling Gentran:Server for UNIX adapter. You cannot turn logging on or off for Sterling Gentran:Server for UNIX data adapter. **Note:** The Sterling Gentran:Server for UNIX logs are displayed only if you have Sterling B2B Integrator configured for Sterling Gentran:Server for UNIX.
Application Logs		
Adapter Server	servicesctl.log	Used by the service controller component.
AFT Routing	aftrouting.log	

Download from Www.Somanuals.com. All Manuals Search And Download.

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.
IBM Sterling Connect:Direct® Secure Perimeter Adapter	cdsp.log	
Sterling Connect:Direct Server Adapter Protocol Layer	cdinterop_cdjava.log	Used by the Sterling Connect:Direct Server adapter.
Sterling Connect:Direct Server and Requester Adapter and Services	cdinterop.log	Used by the Sterling Connect:Direct Server and Requester Adapter and related services
IBM Sterling Connect:Enterprise® Server Adapter and Services	ceuinterop.log	Used by the Sterling Connect:Enterprise Server adapter.
Crypto	crypto.log	
CSP2 FTP Adapter	cspftp.log	
CSP2 Http Adapter	csphttp.log	
Dashboard	mgmtdash.log	Used by the Dashboard component.
Delete Resources	delete.log	Logs information about resources that have been deleted from Sterling B2B 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	
EDIINT AS1 and AS2	EDIINT.log	Used by the EDIINT components.
Embedded Engine	embeddedEngine.log	

Download from Www.Somanuals.com. All Manuals Search And Download.

Log Type	Log Name	Description
Event Framework	event.log	Logs event framework activities for events completed in Sterling B2B 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.
Sterling Gentran:Server for UNIX Lifecycle	lifecycle.log	Used by the Sterling Gentran:Server for UNIX Lifecycle components when loading lifecycle records.
Sterling Gentran:Server for UNIX Lifecycle Purge Service	system.log	Used by the Sterling 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 B2B 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 B2B Integrator.
OCSP	ocsp.log	Used by the Online Certificate Status Protocol.
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 B2B 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	
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	
Security	security.log	Used by the security components. Indicates problems with startup and component licensing.

Download from Www.Somanuals.com. All Manuals Search And Download.

Log Type	Log Name	Description
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.
IBM Sterling Secure Proxy	secureproxy.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 B2B Integrator.
Web Services Security	wssec.log	
WebDAV Server	webdav.log	
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 B2B 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	wsrm.log	

The following table describes the log files pertaining to the Sterling B2B Integrator Windows service:

Log Name	Description
ScheduleBackup.log	Temporary file that is created when Sterling B2B Integrator Windows service stops.
Backuplogs.log	Temporary file that is created when Sterling B2B Integrator Windows service stops.

Download from Www.Somanuals.com. All Manuals Search And Download.

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

Auditing

In Sterling B2B Integrator, you can find information about the creation, modification, and deletion of a Sterling B2B 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.
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

Download from Www.Somanuals.com. All Manuals Search And Download.

- 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 B2B Integrator using the Thread Monitor. The Thread Monitor lists all the threads that

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B 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 IBM 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 B2B Integrator to use in other business processes.

Stopping a Thread

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

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B 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 B2B 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 B2B Integrator.

It is recommended that you keep the following considerations in mind when monitoring messages within Sterling B2B Integrator:

- The Message Monitor pages of the Sterling B2B 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 B2B 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

- Unsuccessful messages remain in Sterling B2B 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 B2B 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 B2B Integrator has outstanding unsuccessful messages.
     – If there are no unsuccessful messages in Sterling B2B 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 B2B 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 B2B Integrator.

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B 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 the whole Sterling B2B Integrator cluster using the softstop script. After that, you still need to call the hardstop script to stop the whole system. Compare with the **Soft Stop** command.
Select Node	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. 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. 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

Area	Information
Host Information	The Host Information displays the following information: • Start time • Uptime • Host • Location • State • Memory available • Active threads
Sterling Gentran:Server	This area is displayed only if you have configured your system for Sterling Gentran:Server Data Manager support. The Sterling Gentran:Server area displays the following information: • Host • Location • State • Controllers • Adapters
Classpath	Displays the Sterling B2B Integrator classpath.
JNDI Tree	Displays the JNDI tree in Sterling B2B Integrator.
Soft Stop	Stops a node of Sterling B2B Integrator using the softstop script interactively through the UI. Compare with the **Stop the System** command.
Database Usage	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.
Business Process Queue Usage	Displays business process queue usage statistics such as cache disk usage, cache memory usage, queue statistics, and cache statistics.
Business Process Usage	Displays count of business processes by its state.
Cache Usage	Displays size and hit rate for object caches.
Threads	Displays active processes at a thread level.
Clean-Up Processes Monitor	Displays the time since the archive, index, purge, and recovery tasks were completed.
Controllers	The state and name of each controller or server in the Sterling B2B Integrator installation.

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B Integrator. The Perimeter Servers area displays the following information: • 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 B2B Integrator from the System Troubleshooting Page

You can stop Sterling B2B Integrator using the System Troubleshooting page, using either a soft stop (which allows all the business processes to complete) or a hard stop (which stops the system and all the processing immediately). After a soft stop, you must still perform a hard stop to stop the system. For more information about the soft stop process, refer to "Soft Stop of Sterling B2B Integrator" on page 110.

Important: Using the **Soft Stop** option stops only the Sterling B2B Integrator interface immediately, while all the business processes that are in progress run until complete. After all the current business processes are completed, you will need to perform a hard stop to fully stop Sterling B2B Integrator. To stop the system and all the processing immediately, use the **Stop the System** option. All processes that have not been completed will stop and have to be restarted.

To stop Sterling B2B 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 **Soft Stop** (for an interactive soft stop) or **Stop the System** (for an immediate soft stop).

Important: A soft stop does not stop all JVMs and system processing. To fully stop the system, you must perform a hard stop after the soft stop is complete. For instructions, see "Hard Stop of Sterling B2B Integrator" on page 120.

Viewing the Host Information

Host Information displays the current operational status of the processing environment for your installation of Sterling B2B Integrator. The System Troubleshooting page displays separate information for each installation.

Note: If you are working in a clustered installation of Sterling B2B 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**.

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B 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 B2B 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B 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
     - *database*ArchivePool
     - *database*Pool
     - *database*Pool_local
     - *database*Pool_NoTrans
     - *database*Pool_Select
     - *database*UIPool

Download from Www.Somanuals.com. All Manuals Search And Download.

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 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 business process execution cycle time based on priority. The average business process execution cycle time may include the execution times of several steps. It captures the average time that business processes 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

Download from Www.Somanuals.com. All Manuals Search And Download.

Viewing Cache Usage Information

Sterling B2B 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 "Performance Tuning Utility" on page 120. 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 B2B 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

- 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 IBM 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 windows 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

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 wan 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

Heading	Description
Token Node	Indicates whether this is a token node. • 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	Displays the status of a node. • 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. **Notes:** • If it is indicated that a node has gone down, in the *install_dir*/bin directory of the node, run the hardstop.sh (for UNIX) or the hardstop.cmd (for Windows) script, and then run the run.sh (for UNIX) or the startWindowsService.cmd (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	Indicates whether this is a token node. • 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

Heading	Description
Status	Displays the status of a node. • 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. **Notes:** • If it is indicated that a node has gone down, in the *install_dir*/bin directory of the node, run the hardstop.sh (for UNIX) or the hardstop.cmd (for Windows) script, and then run the run.sh (for UNIX) or the startWindowsService.cmd (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.
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 B2B Integrator Version	Displays the version of Sterling B2B 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 queueWatcher=true in the *install_dir*/properties/noapp.properties_platform_ifcresources_ext file.	
Multicast Info	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. Clustering uses the JGroups communication toolkit, which supports the following communication protocols, which can be configured during deployment: • UDP (User Datagram Protocol) (IP multicast) • TCP
Engine Queue Content	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. Use the drop-down menu to select the queue you want to review
Distribution Size	Provides a link to the Distribution Size page. This page contains the following information: • Scheduling policy details • Soft reference cache details • Queue details (for each queue)

Download from Www.Somanuals.com. All Manuals Search And Download.

Monitoring Deprecated Resources

Sterling B2B 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 B2B Integrator.

Note: Resources that have been identified for deprecation will continue to be supported by Sterling B2B Integrator.

Manage the deprecated resources as follows:

1. Identify the deprecated resources using the Sterling B2B 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 alphabet 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 B2B 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: • 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: • Replaced by FTP Client Adapter • Use the DeenvelopeEDIFACT service • Use the DeenvelopeX12 service

Download from Www.Somanuals.com. All Manuals Search And Download.

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.

Soft Stop of Sterling B2B Integrator

A soft stop of Sterling B2B Integrator allows existing processes the opportunity to complete or be saved for resumption after restart.

This can save you time by eliminating the need to manually resume processes after a restart. You can run a soft stop from either the UI or from the command line.

A soft stop provides the ability to interactively step through the shutdown process. It allows the following to happen during a shutdown:

- The transfer of running business processes to other cluster nodes or to the database for recovery.
- The ability to control the completion of queued business processes.
- The ability to show the status of adapter activity and open sessions.

The individual steps of a shutdown are as follows. The soft stop allows you to control the execution of the individual steps required for graceful shutdown.

1. A node is marked unavailable so that it does not accept jobs from other nodes.
2. The unavailable node information is broadcast to other cluster nodes so that each node knows which nodes are in the process of shutting down.
3. Adapters are signaled information about the shutdown so that they can stop accepting new requests.
4. The resource monitor is shut down.
5. The scheduler is shut down. For clustered installations, schedules are moved to another live cluster node.
6. All new live business processes are blocked from their queues and are re-directed to other live cluster nodes or are saved to the database with a status of HALTED_SOFTSTOP for autorecovery.
7. The business process drain policy and completion policy are used to redistribute business processes to other live cluster nodes or save them to the database with a status of HALTED_SOFTSTOP for autorecovery if there is no other live node or if it is a standalone installation.
   - The completion policy determines if you finish the whole business process; the execution cycle of the business process; or the step of the business process.
   - The drain policy determines what you need to do with the business processes in the business process queue (drain all business processes; drain not started business processes; drain not started business processes that have no open servers session against adapters; or drain no business processes).
8. The business processes remaining on the queue are allowed to complete.
9. The system waits for all business processes to be completed.
10. The user can watch the adapter activity/sessions and manually control when to shut down adapters. Some adapters will not stop immediately. These adapters will defer shutdown to allow existing transfers to complete. The status of these adapters can be seen via the **!** icon in the adapter list.

Download from Www.Somanuals.com. All Manuals Search And Download.

Note: You can also shut down adapters using the graceful adapter shutdown tool. For more information, see Graceful Adapter Shutdown.

11. The JVM is exited.

    Note: When you run the interactive soft stop process, it should continue until the JVM is exited. You should not try to undo any of the soft stop actions to bring back the system.

    Important: The soft stop only stops the noapp and ops JVMs. To stop the remaining Sterling B2B Integrator JVMs and stop the system processing, you must manually run a hard stop after the soft stop is complete.

12. After a cluster node restarts, processes that are saved in the database as HALTED_SOFTSTOP are started/resumed automatically.

Completing a Soft Stop of Sterling B2B Integrator from the System Troubleshooter

You can perform a soft stop from the system troubleshooter in the UI.

About this task

To perform this procedure from the command line, see "Completing a Soft Stop of Sterling B2B Integrator from the Command Line" on page 115.

Important: A soft stop of Sterling B2B Integrator using the system troubleshooter is not subject to timeout values and the soft stop will run as long as necessary for all current process to finish processing.

To complete a soft stop of Sterling B2B Integrator from the system troubleshooter:

Procedure

1. Log in to the system troubleshooter (**Operations** > **System** > **Troubleshooter**).
2. If running a cluster, select a node to shut down at the **Select Node** field. If you need to shut down the container node along with the ASI node, you must first shut down the container node and then the ASI node, because shutting down the ASI node first will disable the soft stop UI for the container node.

   In a clustered installation, you also can soft stop the container node by going to the Troubleshooter UI of the other live ASI node in the cluster and soft stopping the container node from there.
3. Click **Soft Stop**.

   The soft stop page opens in a new popup window.

Download from Www.Somanuals.com. All Manuals Search And Download.

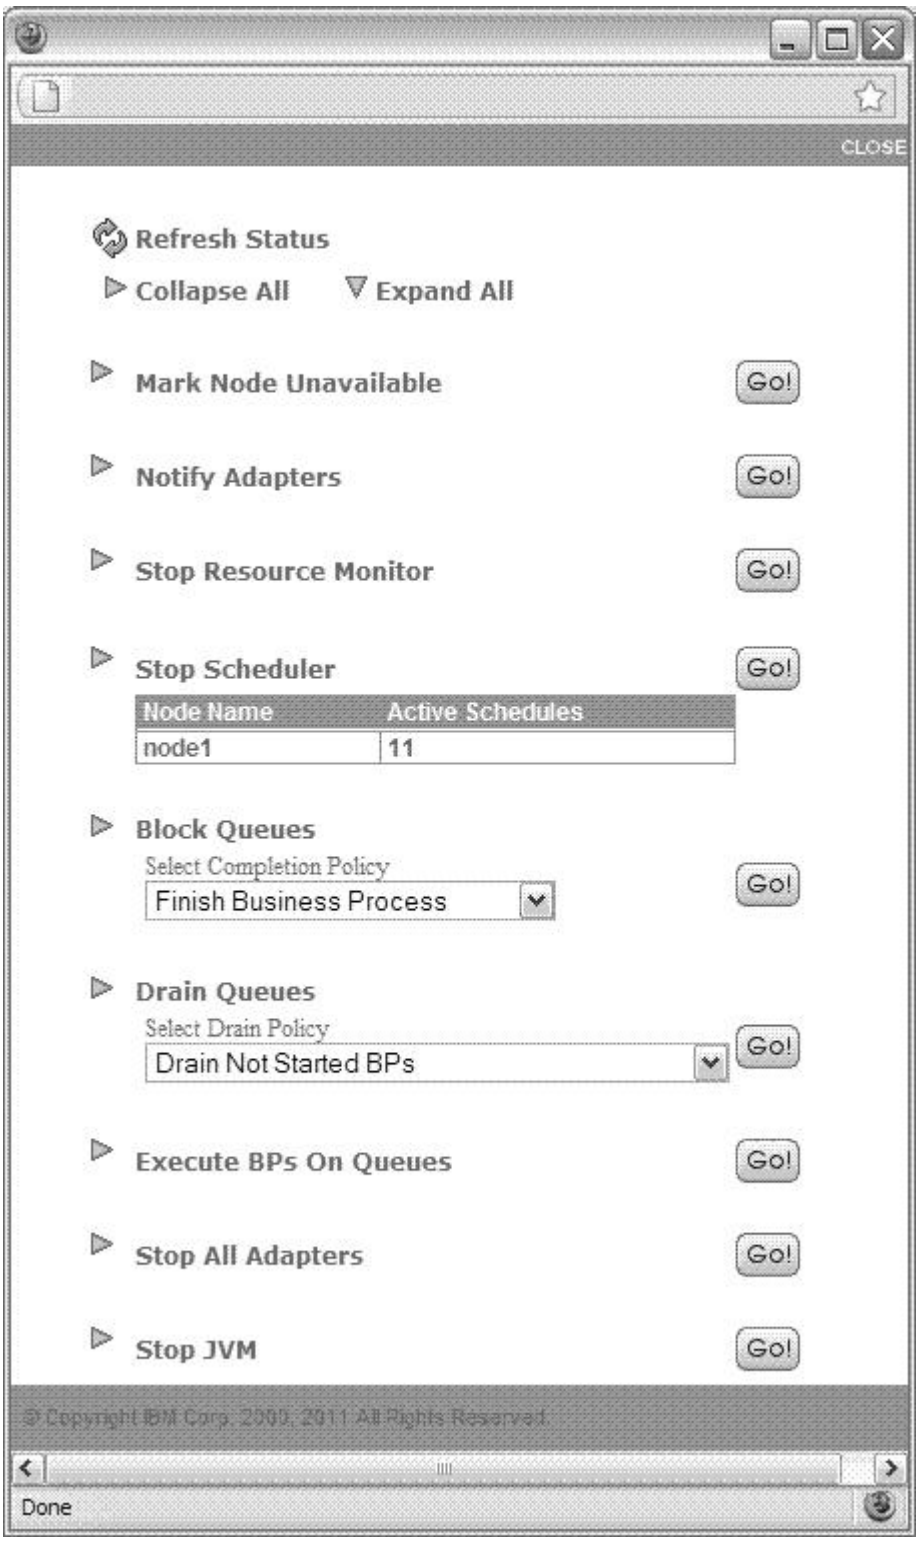

4. On the soft stop page, perform the following actions in the lettered sequence. When you run the interactive soft stop process, it should continue until the JVM is stopped. You should not try to undo any of the soft stop actions to bring back the system.

Download from Www.Somanuals.com. All Manuals Search And Download.

a. **Mark Node Unavailable** In cluster installations, this is used to prevent business processes from coming from other nodes to the node that is marked unavailable.

b. **Notify Adapters** This is used to notify all adapters that the server is shutting down so that adapters will stop taking new external requests.

   Note: This does not stop the adapters, which happens at the **Stop All Adapters** step later in this process.

   For the container node, you will see the adapter session status with two lists:

   - A list of the business processes with a client adapter in session.
   - A list of the running and queued business processes invoked by server adapters.

   For the ASI node, the list of business processes with a client adapter session will show after you click the button **Execute BPs On Queues**. If a business process is a sub business process of the business process with a client adapter session, it will not inherit the client adapter session information from its parent, so it will not show up in the list. It is recommended that you use in-line invoking to invoke a sync BP.

   The **Transport Sessions** section lists the business processes with active transport sessions. Those business processes are invoked by the adapters on the node that is being shut down. When the business process finishes, it will disappear from the list.

c. **Stop Resource Monitor** This is used to stop the resource monitor.

d. **Stop Scheduler** This is used to stop the scheduler in both clustered and standalone installations. In a clustered installation, this is used to move schedules to another live cluster node. A table lists the active schedules for the node.

e. **Block Queues**

   This is used to stop the execution of business processes on a queue and not let nodes send any new business processes to it.

   When you block a queue, if the business process completion policy is **Finish Current Step**, the executing business process will exit after it finishes the current step.

   If the policy is **Finish Current Execution Cycle**, the executing business process will exit after it finishes the current execution cycle.

   If the policy is **Finish Business Process** (the default value), the executing business process will not exit until it finishes.

   Note: The **Finish Current Step** option does not stop a business process that is in the middle of an open communication session. Instead, the softstop waits until the session is terminated. Following termination, the business process will be marked HALTED_SOFTSTOP or move to another available node.

Download from Www.Somanuals.com. All Manuals Search And Download.

For example, a business process starts a Sterling Connect:Direct Begin Session service, and then a sleep service is active for 5 minutes. During that time, a softstop is triggered with the **Finish Current Step** completion policy. The BP will not be marked HALTED_SOFTSTOP or move to another node until the Sterling Connect:Direct End Session service is executed. This allows BP sessions to complete successfully without interruption.

f. **Drain Queues**

This is used to distribute the business processes in the queue to other live nodes (cluster installations only) or save them to the database for recovery (cluster and noncluster installations).

If the drain policy is **Drain ALL BPs**, then both started and unstarted business processes are distributed or saved. If it is **Drain Not Started BPs** (the default value), then just unstarted business processes are distributed or saved. If it is **Drain No BPs**, then no business processes are distributed or saved. If it is **Drain Not Started BPs and No Open Session BPs**, then the business processes that are unstarted and with no open server session are distributed or saved. Processes that are distributed to other queues use the same queue number.

Note: If the business process completion policy in the Block Queues step is **Finish Business Process** (the default value), then the specified drain policy will be honored. If the business process completion policy is **Finish Current Step** or **Finish Current Execution Cycle**, then the drain policy will be **Drain ALL BPs**.

The following graphic shows an example of the results of selecting the **Drain Not Started BPs** option.

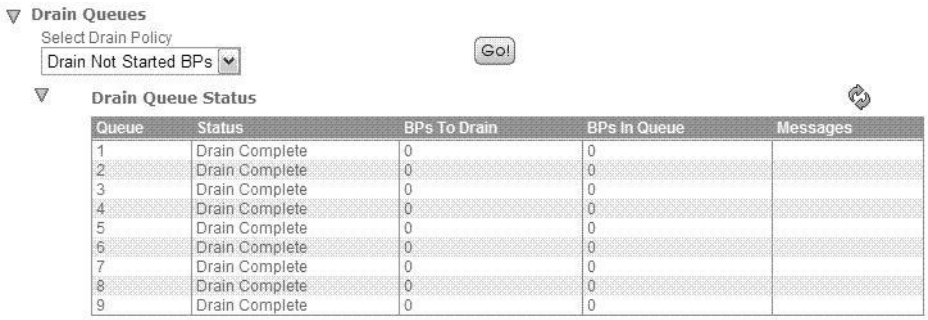

Queue	Status	BPs To Drain	BPs In Queue	Messages
1	Drain Complete	0	0	
2	Drain Complete	0	0	
3	Drain Complete	0	0	
4	Drain Complete	0	0	
5	Drain Complete	0	0	
6	Drain Complete	0	0	
7	Drain Complete	0	0	
8	Drain Complete	0	0	
9	Drain Complete	0	0	

g. **Execute BPs On Queues** This is used to allow the business processes on the queue to complete. The following graphic shows an example of this action.

- The **BPs Execution Summary** section shows the business process activity for all queues.
- In the **BPs In Execution** section, you can do the following:
  - Click **Show Service/Adapter Activity Status** to show all service/adapter activity information. You can access the same information from the Administration UI by clicking **Business Process** > **Monitor** > **Current Activities**.
  - Click **Show business process Execution Thread Status** to show all threads executing business processes. You can access the same information from the Administration UI by clicking **Operations** > **Thread Monitor**.

    If there are business processes executing that you no longer want to wait for, they can be terminated via the **Show business process Execution Thread Status** link, using the **Stop** or **Interrupt** buttons.

Download from Www.Somanuals.com. All Manuals Search And Download.

- The **BPs In Queue** section gives detailed information about business process activity. If there are business processes remaining in the queue that you no longer want to wait for, they can be drained by clicking the **Force Drain BPs** link that appears in the **BPs In Queue** section when business processes are in the queue.
- The **Client Sessions** section lists business processes that have an open client session with adapters. You need to wait for those business processes to complete before you proceed to the next step to stop all adapters.

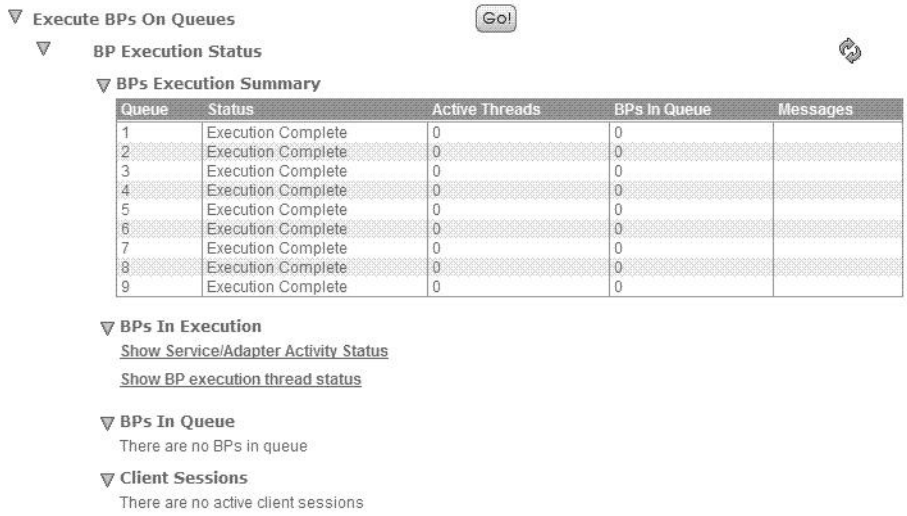

h. **Stop All Adapters**

This is used to shut down all adapters until you reach a timeout (in seconds) that is specified in the **shutdown.timeout_for_adapters** property in the `noapp.properties_platform_ifcresources_ext.in` property file. If the timeout value is zero, the adapters will not shut down gracefully.

Note: You can also shut down adapters using the graceful adapter shutdown tool. For more information, see Graceful Adapter Shutdown.

i. **Stop JVM** This is used to stop the Java Virtual Machine of a node in both clustered and standalone installations.

Important: The soft stop only stops the noapp and ops JVMs. To stop the remaining Sterling B2B Integrator JVMs and stop the system processing, you must manually run a hard stop after the soft stop is complete.

What to do next

Important: A soft stop does not stop all JVMs and system processing. To fully stop the system, you must perform a hard stop after the soft stop is complete. For instructions, see "Hard Stop of Sterling B2B Integrator" on page 120.

Completing a Soft Stop of Sterling B2B Integrator from the Command Line

Use the **softstop** command to stop Sterling B2B Integrator gracefully. Unfinished jobs will be distributed to other nodes or saved to the database to automatically resume after startup.

Download from Www.Somanuals.com. All Manuals Search And Download.

About this task

For individual steps of the graceful soft stop, refer to the documentation on the **opscmd** command.

Important: The individual steps of a soft stop of Sterling B2B Integrator from the command line are subject to timeout values, stored in the **shutdown.*** properties of the noapp.properties_platform_ifcresources_ext.in property file.

Procedure

1. In the installation directory of Sterling B2B Integrator, navigate to the bin subdirectory.
2. Run the softstop.sh (UNIX) or softstop.cmd (Windows) command with the following options:

   Note: The default values for many of these options are stored in the **shutdown.*** properties of the noapp.properties_platform_ifcresources_ext.in property file.

   - node_name

     The cluster node name, like node1 or node2.
   - all

     Shuts down the whole cluster.
   - allLocalNodes

     Shuts down all of the local nodes of an application instance, including container nodes.

     Note: If you shut down all of the nodes and the ActiveMQ process is still running, you need to run hardstop.sh to kill it. This applies to shutting down the whole cluster.
   - [-tq timeout_for_queues]

     Timeout value in seconds for all business processes on the queue to finish.
   - [-ta timeout_for_adapters]

     Timeout value for stopping all adapters.

     Note: You can also shut down adapters using the graceful adapter shutdown tool. For more information, see Graceful Adapter Shutdown.
   - [-dp bp_drain_policy]
     - 1 - Distribute all business processes on the queue to other nodes or save them to the database.
     - 2 - Distribute all not-started business processes on the queue to other nodes or save them to the database.
     - 3 - Let all the business processes on the queue be processed until a timeout for the shutdown.
     - 4 - Distribute all business processes on the queue, that have not started and have no open sessions against adapters, to other nodes or save them to the database.
   - [-cp bp_completion_policy]

     If bp_completion_policy is **1** or **2**, then bp_drain_policy is set to **1**.
     - 1 - Complete the current business process step.
     - 2 - Complete the current business process execution cycle.

Download from Www.Somanuals.com. All Manuals Search And Download.

– 3 - Complete the current business process.

What to do next

Important: A soft stop does not stop all JVMs and system processing. To fully stop the system, you must perform a hard stop after the soft stop is complete. For instructions, see "Hard Stop of Sterling B2B Integrator" on page 120.

Soft Stop opscmd Commands

You can use `opscmd` commands to run commands from the command line for each step of the soft stop process.

The following `opscmd` commands allow you to run commands from the command line for each step of the soft stop process. The chart shows the UNIX command. For Windows, use the `opscmd.cmd` command with the same parameters.

Action	Command
Marking a node unavailable	`opscmd.sh -cMarkNodeUnavailable -pname=node_name`
Marking a node available	`opscmd.sh -cMarkNodeAvailable -pname=node_name`
Block a queue or all queues	`opscmd.sh -cBlockQueue -pname=node_name -pid={queue_id | all} [ -pbpCompletionPolicy={1 | 2 | 3} ]` `queue_id` can be **1-9** or (for all queues).
Unblock a queue or all queues	`opscmd.sh -cUnblockQueue -pname=node_name -pid={queue_id | all}`
Drain a queue or all queues	`opscmd.sh -cDrainQueue -pname=node_name -pid={queue_id | all} [-pbpDrainPolicy={1 | 2 | 3 | 4}]` **Note:** If the business process completion policy is the default value of **Finish Business Process** (pbpCompletionPolicy=1), then the specified drain policy will be honored. If the business process completion policy is **Finish Current Step** (pbpCompletionPolicy=2) or **Finish Current Execution Cycle** (pbpCompletionPolicy=3), then the drain policy will be **Drain ALL BPs**. **Note:** This is an asynchronous command. Use the `cGetDrainQueueStatus` command to see the status of the queue draining.
See the status of the queue draining	`opscmd.sh -cGetDrainQueueStatus -pname=node_name -pid={queue_id | all}`
Move Schedules to another node and stop the scheduler	`opscmd.sh -cStopScheduler -pname=node_name [-pFailoverNodeName=node_name]` **Note:** If the failover node is not specified, the system will automatically pick one for you, if possible.
Signal adapters for softstop	`opscmd.sh -cPrepareStopAllAdapters -pname=node_name`
Stop all adapters **Note:** You can also shut down adapters using the graceful adapter shutdown tool. For more information, see Graceful Adapter Shutdown.	`opscmd.sh -cStopAllAdapters -pname=node_name [-pTimeoutForAdapters=time_out_in_seconds]`

Download from Www.Somanuals.com. All Manuals Search And Download.

Action	Command
List all active adapters.	`opscmd.sh -cListActiveAdapters -pname=node_name` If the adapter type is stateless, this command will list only one instance of the adapter, even though it may have many configurations for each node. If the adapter type is stateless, only one instance (configuration) is applicable for each service type. If the adapter type is stateful, then all the active instances are listed for this command for each adapter.
Stop resource monitor	`opscmd.sh -cStopResourceMonitor -pname=node_name`
List all active business processes invoked by server adapters that have not sent the business process result back yet	`opscmd.sh -cGETTRANSPORTSESSIONBPS [-pname=node_name] -pid={queue_id | all} -padapterNodeName=adapter_node_name` Where `adapter_node_name` is the name of the node that is being shut down. If `-pname` is not specified, all ASI nodes will be scanned. The list of business processes returned by this command have a client session with the adapters on the node.
List all business processes that have not closed sessions with client adapters on the node that is being shut down.	`opscmd.sh -cGETCLIENTSESSIONBPS [-pname=node_name] -padapterNodeName=adapter_node_name` Where `adapter_node_name` is the name of the node that is being shut down. If `-pname` is not specified, all ASI nodes will be scanned.
Execute business processes on a queue or on all queues.	`opscmd.sh -cExecuteBPsOnQueue -pname=node_name -pid={queue_id | all} [-ptimeoutForQueues=timeout_out_in_seconds]` **Note:** This is an asynchronous command. Use the `cGetExecuteBPsOnQueueStatus` command to see the status of BP execution.
See the status of the business process queues execution.	`opscmd.sh -cGetExecuteBPsOnQueueStatus -pname=node_name -pid={queue_id | all}`
Stop the JVM of a node	`opscmd.sh -cSTOPJVM -pname=node_name`
Shut down a node or a whole cluster	`opscmd.sh -cSHUTDOWN [-pname=node_name] [-pBPCompletionPolicy={1 | 2 | 3}] [-pBPDrainPolicy={1 | 2 | 3 | 4}] [-pTimeoutForQueues=time_out_in_seconds] [-pTimeoutForAdapters=time_out_in_seconds]` **Note:** If a node is not specified, the whole cluster is shut down. **Note:** If the optional parameter `BPCompletionPolicy`, `BPDrainPolicy`, `TimeoutForQueues`, or `TimeoutForAdapters` is not specified, the default property value will be used from the `noapp.properties_platform_ifcresources_ext.in` file.

Important: A soft stop does not stop all JVMs and system processing. To fully stop the system, you must perform a hard stop after the soft stop is complete. For instructions, see "Hard Stop of Sterling B2B Integrator" on page 120.

Graceful Adapter Shutdown

Sterling B2B Integrator supports clean adapter shutdown. That is, when stop adapter is issued, the system allows the adapter to clean up its own session before it gets removed from the system.

Only a few stateful adapters support this feature, for example, swiftAdapter.

Download from Www.Somanuals.com. All Manuals Search And Download.

The progress of shutting down the adapter can be viewed via UI pages. The progress information is available in the following pages:

- In the **System Troubleshooting** page:
  - **Soft Stop** link **> Stop All Adapters** section
  - **Adapters** link
  - Clicking on the adapter name
- In the **Services Configuration** page:
  - **List** and **Search** results
  - Clicking on the adapter name

Graceful Adapter Shutdown through the User Interface

The following steps can be performed in both the service configuration and adapters pages.

- To get to the adapter page, click **Operations** > **System** > **Troubleshooter** > **Adapters**.
- To get to the service configuration page, click **Deployment** > **Services** > **Configuration**, and then search for the adapter.

List the running stateful adapters by clicking the **Adapter** link.

1. Shut down an adapter by de-selecting the checkbox under the **Enable** column.

   The following message appears in the **Advanced State** column:

   `Stopping`
2. Click on the icon in the **Advanced State** column.

   A popup window appears with two sections:

   - **System Audit Information**

     This shows the status changes along with a timestamp of the whole shutdown process.
   - **Adapter Advanced State Information**

     When the adapter finishes its current session, and shutdown is completed, the following message appears:

     `Adapter adapter_name is currently unavailable`

Graceful Adapter Shutdown through the Command Line

1. Navigate to the `bin` subdirectory of your installation directory.
2. Shut down the adapter by running the command:

   `opscmd.sh -cSTOPADAPTER -nnode1 -pid=adapter_name`
3. Check on the status of the shutdown by running the command:

   `opscmd.sh -cGETADAPTERADVSTATE -nnode1 -pid=adapter_name` Run this command until the status message changes from `Adapter is preparing to shutdown to Adapter adapter_name is currently unavailable`.
4. Confirm that the adapter has been shut down by running the command:

   `opscmd.sh -cGETADAPTERADVSTATE -nnode1 -pid=adapter_name`

   The following message should appear:

   `Adapter adapter_name is currently unavailable`

   The adapter session is finished and shutdown is complete.

Download from Www.Somanuals.com. All Manuals Search And Download.

Hard Stop of Sterling B2B Integrator

A hard stop of Sterling B2B Integrator immediately stops the system without waiting for business processes to finish.

About this task

Hard stops may result in loss of data in unfinished processes.

To perform a hard stop:

Procedure

1. Navigate to `/install_dir/install/bin`.
2. Enter `./hardstop.sh` (UNIX) or `hardstop.cmd` (Windows).
3. Enter your passphrase.

Performance Tuning

Sterling B2B Integrator provides a utility to help with performance tuning. You can also perform manual performance tuning.

Performance Tuning Utility

The performance tuning utility provided with Sterling B2B Integrator helps calculate the tuning parameter values and configure Sterling B2B Integrator for you.

Using the Performance Tuning utility, you can tune and configure Sterling B2B Integrator according to your business requirements. You can change the queue settings, memory allocations, JVM settings, and pool settings. However, note that you cannot use the utility to tune your business processes. For more information about tuning business processes, refer to the *Business Process Modeling Guide*.

In the future, if you apply patches to Sterling B2B Integrator, it will not overwrite the performance tuning changes made by the Performance Tuning utility.

Note: The suggested performance configuration settings are based on standard calculations. You should test the performance configuration settings in a test environment with appropriate load and then apply the settings in the production environment.

There are performance trade-offs that may occur based on your performance configuration settings. However, you can decide to increase or decrease the values based on your specific requirements. Following is a list of such scenarios that may occur, but not limited to:

- If you increase memory for server nodes and container nodes, the system will allocate maximum memory for the JVMs thereby reducing the amount of memory allocated to other applications. It may even lead to JVMs holding excess amount of memory, which could have been used by other applications.
- If you increase cache settings, it may increase the speed of disk read operations, but might affect other resources waiting for memory.
- If you increase the values in queue settings, business processes may be executed swiftly, but it might affect the read or write operations to the database

Download from Www.Somanuals.com. All Manuals Search And Download.

- If you increase the threads, the number of business processes that can run concurrently may increase, but it might lead to performance degradation as it might consume excess memory.

This section provides information about viewing, restoring, and editing performance configuration settings through performance tuning utility.

View and Restore Performance Configuration Settings

You can view current performance configuration settings to verify if a particular performance configuration setting is correct and restore previous performance configuration settings.

View Performance Configuration Settings

You can view the current performance configuration settings to verify if they are correct. The Performance Tuning page displays all settings that are currently set and suggested changes. The suggested changes displayed in this page are also displayed if you are tuning the performance using the performance tuning utility. For more information about using the performance tuning utility, refer to the topic "Edit Performance Configuration Settings" on page 123.

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, which include information on:
   - 1) System
   - 2) JVM Parameters for Server
   - 3) JVM Parameters Container Nodes
   - 4) BP Queue
   - 5) Cache
   - 6) Purge Settings
   - 7) Database Connection Pool
   - 8) Queue Settings
   - 9) Memory
4. Click **Return** to return to the Performance Tuning page.

Restore Performance Configuration Settings

You can manually restore performance configuration settings to return Sterling B2B Integrator to the previous performance state. In the Restore Tuning page, the properties that will change if you restore the previous performance configuration settings are displayed in blue color text and bold format.

The User ID, along with the user name, and the time when the last performance configuration settings were last edited are displayed under the Restore tab in Performance Tuning page.

Download from Www.Somanuals.com. All Manuals Search And Download.

You can restore your performance configuration settings only once. If you try restoring the performance settings more than once, there will be no changes in the settings. You can restore the performance configuration settings under the following conditions, but not limited to:

- Previous performance setting was more efficient than the current setting.
- When you are testing performance by setting optimum values, restoring to the previous setting is convenient.
- Testing under varying load conditions to benchmark performance configuration settings.

To restore your performance configuration settings:

1. From the **Administration** menu, select **Operations** > **System** > **Performance** > **Tuning**.
2. In the Performance Tuning page, under **Restore**, next to **View/Restore Previous 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 Restore Tuning Configuration page that is displayed, the properties that will change if you restore the previous performance configuration settings are displayed in blue color text and bold format. Review the performance configuration settings and perform one of the following tasks:
   - To abandon restoring the previous performance configuration settings, click **Cancel**.
   - To restore the previous performance configuration settings, click **Finish**. The Restore Tuning Properties page appears after the restoration is complete. The following message appears: `The system update completed successfully.`

     If you restore the previous configuration settings, click **Return** to return to the Performance Tuning page.
5. From the **Administration** menu, select **Operations** > **System** > **Troubleshooter**.
6. From the System Troubleshooting page, click **Soft Stop**. This runs the softstop script, allowing all active business processes to complete on a node before stopping the system. The script is run through a series of interactive steps in the UI. For more information about the softstop process, refer to "Completing a Soft Stop of Sterling B2B Integrator from the System Troubleshooter" on page 111.
7. 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.
8. From the *install_dir*/bin directory, start Sterling B2B Integrator by running one of the following commands:
   - (UNIX or Linux) **run.sh**
   - (Windows) **startWindowsService.cmd**

   Sterling B2B Integrator starts and the previous performance configuration settings are applied. The properties in the tuning.properties file will have the same values as the corresponding properties in the

Download from Www.Somanuals.com. All Manuals Search And Download.

tuningPreviousSetting.properties file. In addition, the tuningPreviousSetting.properties file will have these properties:

- jdbcService.[database]UIPool.buffersize
- jdbcService.[database]UIPool.initsize
- jdbcService.[database]UIPool.maxsize
- jdbcService.[database]Pool_NoTrans.buffersize
- jdbcService.[database]Pool_NoTrans.initsize
- jdbcService.[database]Pool_NoTrans.maxsize
- jdbcService.[database]Pool_local.buffersize
- jdbcService.[database]Pool_local.initsize
- jdbcService.[database]Pool_local.maxsize
- jdbcService.[database]ArchivePool.buffersize
- jdbcService.[database]ArchivePool.initsize
- jdbcService.[database]ArchivePool.maxsize
- jdbcService.[database]Pool_Select.buffersize
- jdbcService.[database]Pool_Select.initsize
- jdbcService.[database]Pool_Select.maxsize
- jdbcService.[database]Pool.buffersize
- jdbcService.[database]Pool.initsize
- jdbcService.[database]Pool.maxsize

Edit Performance Configuration Settings

You can edit performance configuration settings using a tuning utility.

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 B2B Integrator:

- System memory
- JVM parameters for server
- JVM parameters for container nodes
- BP Queue
- Cache
- Purge settings
- Database connection pools
- Queue settings

You can use the Performance Tuning Utility to edit the 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.

The settings modified using the performance tuning utility are stored in the `tuning.properties` file, which updates other properties files with the modified values. However, if you change the settings in other properties files, they are not updated in the `tuning.properties` file. If you are manually changing the performance configuration settings, ensure that you make the changes only in the `tuning.properties` file.

Download from Www.Somanuals.com. All Manuals Search And Download.

If you are running in a cluster environment, you should apply performance configuration settings for each node separately by running the performance tuning utility in all the nodes.

After editing the `tuning.properties` file, you need to do the following to pick up the new property values:

1. In the `tuningFormulas.properties` file, change the TUNING_PROPS_UPDATED property to **TRUE**.
2. Run the `setupfiles.sh` or `setupfiles.cmd` command.
3. Run the `run.sh` or `startWindowsService.cmd` command.

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.
   ```

4. 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.
5. In the System page that is displayed, decide whether you want to use the current settings.
   - If yes, click **Next**.
   - If no, enter information in the following fields and click **Next**.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Physical memory (MB) allocated to Platform Amount of memory allocated for use in processing operations. It is recommended that you use a minimum value of 512-1024 MB for this field, based on your system. Although there is no minimum value for this field, a value lower than the recommended setting could harm performance. **tuning.properties Property:** MEMORY	Not applicable

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Number of Server Nodes The number of server nodes installed. **CAUTION:** **This cannot be changed from the default value of 1, regardless of the number of nodes in the installation.** **Note:** The **Number of Server Nodes** field is not displayed in iSeries environment. **tuning.properties Property:** NUM_ASI_NODES	Not applicable
Number of Container Nodes The number of container nodes installed. Ensure that you enter the correct number of container nodes installed. If you change the number of container nodes, the performance tuning utility will not install additional container nodes. The performance tuning utility will suggest changes to the number of nodes entered in this field instead of the actual number of nodes installed, which if accepted will lead to unnecessary resource allocation. If you add a container to the system after the initial installation, you will have to ensure that the proper amount of memory has been allocated to the container nodes, using the **Physical Memory Allocated to Container (MB)** field on the JVM Parameters for Container Nodes page. **Note:** Number of Container Nodes field is not displayed in iSeries environment. **tuning.properties Property:** NUM_CONTAINER_NODE	Not applicable
Allocate memory for BI Listeners Select **True** from the drop-down list if you want to allocate memory for BI listeners from the physical memory allocated to Sterling B2B Integrator. By default, 256 MB memory is allocated for BI listeners. If you select **False** from the drop-down list, the memory for BI listeners will not be allocated from the physical memory allocated to Sterling B2B Integrator which may lead to improper memory management. **tuning.properties Property:** IS_DEDUCT_MEM_BI	Default value: False

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Allocate memory for Active MQ Select **True** from the drop-down list if you want to allocate memory for Active MQ messaging service from the physical memory allocated to Sterling B2B Integrator. By default, 512 MB memory is allocated for the Active MQ messaging service. If you select **False** from the drop-down list, the memory for Active MQ messaging service will not be allocated from the physical memory allocated to Sterling B2B Integrator, which may lead to improper memory management. **tuning.properties Property:** IS_DEDUCT_MEM_ACTIVEMQ	Default value: False
Allocate memory for SAP Adapter Select **True** from the drop-down list if you want to allocate memory for the SAP Adapter from the physical memory allocated to Sterling B2B Integrator. By default, 256 MB memory is allocated for the SAP Adapter. If you select **False** from the drop-down list, memory for the SAP Adapter will not be allocated from the physical memory allocated to Sterling B2B Integrator, which may lead to improper memory management. **tuning.properties Property:** IS_DEDUCT_MEM_SAPADAPTER	Default value: False
Allocate memory for Command Line Adapter Select **True** from the drop-down list if you want to allocate memory for Command Line Adapter from the physical memory allocated to Sterling B2B Integrator. By default, 128 MB memory is allocated for the Command Line Adapter. If you select **False** from the drop-down list, the memory for Command Line Adapter will not be allocated from the physical memory allocated to Sterling B2B Integrator, which may lead to improper memory management. **tuning.properties Property:** IS_DEDUCT_MEM_CLADAPTER	Default value: False

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Memory will be allocated for Ops Server The memory will be automatically allocated to the ops server. You cannot edit the memory allocated to the ops server. By default, 128 MB memory is allocated for the ops server. **tuning.properties Property:** MEMORY_OPS	Not applicable

Click **Next**.

6. In the JVM Parameters for Server page that is displayed, decide whether you want to use the current settings or the suggested settings.

   Note: Click **Fill Current** displayed in all pages to retain the current settings.

   - If yes, click **Fill Suggested**.
   - If no, enter information in the following fields:

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Physical Memory Allocated to Server (MB) Amount of physical memory allocated to each server node. **Note:** Ensure that the physical memory allocated to server nodes, container nodes, and maximum permanent generation size for both server nodes and container nodes should not exceed the physical memory allocated to Sterling B2B Integrator. **tuning.properties Property:** MEMORY_ASI	Default value (MB): 1536 A minimum of 1536 MB is allocated to each server node out of the total physical memory allocated to Sterling B2B Integrator. However, if there is additional physical memory available after allocating minimum memory to all the server nodes and the container nodes, the memory is split between server nodes and container nodes. By default, the server nodes are allocated with three-fourths out of the available memory and the container nodes are allocated with one-fourths out of the available memory. The additional memory is calculated using the following formula: ((((Physical memory allocated to Platform - (Number of server nodes * Minimum memory allocated to the server nodes)) - (Number of server nodes * Maximum permanent generation size allocated to server nodes)) - (Number of container nodes * Maximum permanent generation size allocated to container nodes)) – Memory allocated to the ops server) For example, let us consider that two server nodes and one container node is present. After allocating memory to the two server nodes and one container node, 200 MB additional memory is available, each server node is allocated 75 MB memory and the container node is allocated 50 MB memory. You can also change the additional memory allocation ratio. The additional memory allocation ratio is calculated using the following formula: (((Additional memory available - Physical memory allocated to container nodes) * 0.75)/Number of server nodes)

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)

Initial Heap Size (MB)
 	
Controls the initial size of a Java heap allocated to each server JVM.
 	
If this is tuned correctly, it can reduce the overhead associated with garbage collection by improving the server response time and throughput.
 	
If you see a large number of minor garbage collections, the default value set to this parameter might be too low. Try increasing it.
 	
Based on your operating system and architecture, the parameters are provided in the following list:
• HP-UX 32-bit: (HP-UX.INIT_HEAP)
• HP-UX 64-bit: (HP-UX-64.INIT_HEAP)
• IBM AIX 32-bit: (AIX.INIT_HEAP)
• IBM AIX 64-bit: (AIX-64.INIT_HEAP)
• Sun Solaris 32-bit: (SunOS.INIT_HEAP)
• Sun Solaris 64-bit: (SunOS-64.INIT_HEAP)
• Linux 32-bit: (Linux.INIT_HEAP)
• Linux 64-bit: (Linux-64.INIT_HEAP)
• Windows 32-bit: (Windows.INIT_HEAP)
• Windows 64-bit: (Windows-64.INIT_HEAP)
• IBM iSeries OS/400®: (OS400.INIT_HEAP)
• IBM OS/390®: (OS390.INIT_HEAP)
 	
tuning.properties Property:
 	
[operating system].INIT_HEAP
	Based on your operating system and architecture, the formulae are provided in the following list:
	• HP-UX 32-bit:
	If memory allocated to the server node is greater than or equal to 3068 MB, it allocates 3068 MB. Else, if the memory allocated to server node is lesser than or equal to 512 MB, it allocates 512 MB. If the memory allocated is greater than 512 MB, but lesser than 3068 MB, it allocates the memory allocated to each server node.
	• HP-UX 64-bit:
	if (MEMORY_ASI >= 3068, MEMORY_ASI, if (MEMORY_ASI <= 512, 512, MEMORY_ASI))
	• IBM AIX 32-bit:
	(0.5 * Physical memory allocated to server node)
	• IBM AIX 64-bit:
	if (MEMORY_ASI >= 3328, MEMORY_ASI, if (MEMORY_ASI <= 1024, 512, round(0.5 * MEMORY_ASI)))
	• Sun Solaris 32-bit:
	If memory allocated to the server node is greater than or equal to 3900 MB, it allocates 3900 MB. Else, if the memory allocated to server node is lesser than or equal to 512 MB, it allocates 512 MB. If the memory allocated is greater than 512 MB, but lesser than 3900 MB, it allocates the memory allocated to each server node.
	• Sun Solaris 64-bit:
	if (MEMORY_ASI >= 3900, MEMORY_ASI, if (MEMORY_ASI <= 512, 512, MEMORY_ASI))

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)

	• Linux 32-bi:
	(0.5 * Physical memory allocated to server node)
	• Linux 64-bit:
	=if (MEMORY_ASI >= 2036, MEMORY_ASI, if (MEMORY_ASI <= 1024, 512, round(0.5 * MEMORY_ASI)))
	• Windows 32-bit:
	if (MEMORY_ASI >= 1536, 768, if (MEMORY_ASI <= 1024, 512, round(0.5 * MEMORY_ASI)))
	• Windows 64-bit:
	if (MEMORY_ASI >= 1536, MEMORY_ASI, if (MEMORY_ASI <= 1024, 512, round(0.5 * MEMORY_ASI)))
	• IBM iSeries OS/400:
	If memory allocated to the server node is greater than or equal to 2048 MB, it allocates 2048 MB. Else, if the memory allocated to server node is lesser than or equal to 512 MB, it allocates 512 MB. If the memory allocated is greater than 512 MB, but lesser than 2048 MB, it allocates the memory allocated to each server node.
	• IBM OS/390:
	If memory allocated to the server node is greater than or equal to 2048 MB, it allocates 2048 MB. Else, if the memory allocated to server node is lesser than or equal to 512 MB, it allocates 512 MB. If the memory allocated is greater than 512 MB, but lesser than 2048 MB, it allocates the memory allocated to each server node.

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)

Maximum Heap Size (MB)
 	
Controls the maximum size of a Java heap allocated to each server JVM.
 	
If this parameter is tuned correctly, it can:
• 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
 	
If you see a large number of garbage collections, try increasing the value. Based on the Operation System (OS), and whether it is a 32-bit JVM or 64-bit JVM, and the available memory, the maximum value that you can go up to will change.
 	
Based on your operating system and architecture, the parameters are provided in the following list:
• HP-UX 32-bit: (HP-UX.MAX_HEAP)
• HP-UX 64-bit: (HP-UX-64.MAX_HEAP)
• IBM AIX 32-bit: (AIX.MAX_HEAP)
• IBM AIX 64-bit: (AIX-64.MAX_HEAP)
• Sun Solaris 32-bit: (SunOS.MAX_HEAP)
• Sun Solaris 64-bit: (SunOS-64.MAX_HEAP)
• Linux 32-bit: (Linux.MAX_HEAP)
• Linux 64-bit: (Linux-64.MAX_HEAP)
• Windows 32-bit: (Windows.MAX_HEAP)
• Windows 64-bit: (Windows-64.MAX_HEAP)
• IBM iSeries OS/400: (OS400.MAX_HEAP)
• IBM OS/390: (OS390.MAX_HEAP)
 	
tuning.properties Property:
 	
[operating system].MAX_HEAP
	Based on your operating system and architecture, the formulae are provided in the following list:
	• HP-UX 32-bit:
	If memory allocated to the server node is greater than or equal to 3068 MB, it allocates 3068 MB. Else, if the memory allocated to server node is lesser than or equal to 512 MB, it allocates 512 MB. If the memory allocated is greater than 512 MB, but lesser than 3068 MB, it allocates the memory allocated to each server node.
	• HP-UX 64-bit:
	if (MEMORY_ASI >= 3068, MEMORY_ASI, if (MEMORY_ASI <= 512, 512, MEMORY_ASI))
	• IBM AIX 32-bit:
	If memory allocated to the server node is greater than or equal to 3328 MB, it allocates 3328 MB. Else, if the memory allocated to server node is lesser than or equal to 512 MB, it allocates 512 MB. If the memory allocated is greater than 512 MB, but lesser than 3328 MB, it allocates the memory allocated to each server node.
	• IBM AIX 64-bit:
	if (MEMORY_ASI >= 3328, MEMORY_ASI, if (MEMORY_ASI <= 512, 512, MEMORY_ASI))

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)

	• Sun Solaris 32-bit:
	If memory allocated to the server node is greater than or equal to 3900 MB, it allocates 3900 MB. Else, if the memory allocated to server node is lesser than or equal to 512 MB, it allocates 512 MB.
	• Sun Solaris 64-bit:
	if (MEMORY_ASI >= 3900, MEMORY_ASI, if (MEMORY_ASI <= 512, 512, can set a maximum heap limit of 4 GB for MEMORY_ASI))
	• Linux 32-bit:
	If memory allocated to the server node is greater than or equal to 2036 MB, it allocates 2036 MB. Else, if the memory allocated to server node is lesser than or equal to 512 MB, it allocates 512 MB. If the memory allocated is greater than 512 MB, but lesser than 2036 MB, it allocates the memory allocated to each server node.
	• Linux 64-bit:
	if (MEMORY_ASI >= 2036, MEMORY_ASI, if (MEMORY_ASI <= 512, 512, MEMORY_ASI))

	• Windows 32-bit:
	if (MEMORY_ASI >= 1280, 1280, if (MEMORY_ASI <= 512, 512, MEMORY_ASI))
	• Windows 64-bit:
	if (MEMORY_ASI >= 1536, MEMORY_ASI, if (MEMORY_ASI <= 512, 512, MEMORY_ASI))
	• IBM iSeries OS/400:
	If memory allocated to the server node is greater than or equal to 2048 MB, it allocates 2048 MB. Else, if the memory allocated to server node is lesser than or equal to 512 MB, it allocates 512 MB. If the memory allocated is greater than 512 MB, but lesser than 2048 MB, it allocates the memory allocated to each server node.
	• IBM OS/390:
	If memory allocated to the server node is greater than or equal to 2048 MB, it allocates 2048 MB. Else, if the memory allocated to server node is lesser than or equal to 512 MB, it allocates 512 MB. If the memory allocated is greater than 512 MB, but lesser than 2048 MB, it allocates the memory allocated to each server node.

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Initial Size of New (Nursery) Heap (MB) Controls the initial size of the new heap area allocated to each server JVM. It is used with the throughput/parallel collector instead of the JVM built-in tuning the parameter provides. Based on your operating system and architecture, the parameters are provided in the following list: • HP-UX 32-bit: (HP-UX.INIT_AGE) • HP-UX 64-bit: (HP-UX-64.INIT_AGE) • IBM AIX 32-bit: (AIX.INIT_AGE) • IBM AIX 64-bit: (AIX-64.INIT_AGE) • Sun Solaris 32-bit: (SunOS.INIT_AGE) • Sun Solaris 64-bit: (SunOS-64.INIT_AGE) • Linux 32-bit: (Linux.INIT_AGE) • Linux 64-bit: (Linux-64.INIT_AGE) • Windows 32-bit: (Windows.INIT_AGE) • Windows 64-bit: (Windows-64.INIT_AGE) • IBM iSeries OS/400: (OS400.INIT_AGE) • IBM OS/390: (OS390.INIT_AGE) **tuning.properties Property:** [operating system].INIT_AGE	(0.33333 * Initial size of heap memory allocated to the server node) AIX 64-bit: AIX-64.INIT_AGE=256 Linux 64-bit: Linux-64.INIT_AGE=256 Windows 32-bit: Windows.INIT_AGE=round(0.33333 * Windows.INIT_HEAP) HP 64-bit: HP-UX-64.INIT_AGE=round(0.33333 * HP-UX-64.INIT_HEAP) Sun 64--bit: SunOS-64.INIT_AGE=round(0.33333 * SunOS-64.INIT_HEAP) Windows 64-bit: Windows-64.INIT_AGE=round(0.33333 * Windows-64.INIT_HEAP)

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Maximum Size of New (Nursery) Heap (MB) Controls the maximum size of the new heap area allocated to each server JVM. Based on your operating system and architecture, the parameters are provided in the following list: • HP-UX 32-bit: (HP-UX.MAX_AGE) • HP-UX 64-bit: (HP-UX-64.MAX_AGE) • IBM AIX 32-bit: (AIX.MAX_AGE) • IBM AIX 64-bit: (AIX-64.MAX_AGE) • Sun Solaris 32-bit: (SunOS.MAX_AGE) • Sun Solaris 64-bit: (SunOS-64.MAX_AGE) • Linux 32-bit: (Linux.MAX_AGE) • Linux 64-bit: (Linux-64.MAX_AGE) • Windows 32-bit: (Windows.MAX_AGE) • Windows 64-bit: (Windows-64.MAX_AGE) • IBM iSeries OS/400: (OS400.MAX_AGE) • IBM OS/390: (OS390.MAX_AGE) **tuning.properties Property:** [operating system].MAX_AGE	(0.5 * Maximum size of heap memory allocated to the server node) AIX 64-bit: AIX-64.MAX_AGE=512 Linux 64-bit: Linux-64.MAX_AGE=512 Windows 32-bit: Windows.MAX_AGE=round(0.33333 * Windows.MAX_HEAP) HP 64-bit: HP-UX-64.MAX_AGE=round(0.33333 * HP-UX-64.MAX_HEAP) Sun 64-bit: SunOS-64.MAX_AGE=round(0.33333 * SunOS-64.MAX_HEAP) Windows 64-bit: Windows-64.MAX_AGE=round(0.33333 * Windows-64.MAX_HEAP)

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Maximum Permanent Generation Size (MB) Controls the maximum size of the permanent generation heap allocated to the server JVM. It 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 error. Based on your operating system and architecture, the parameters are provided in the following list: • HP-UX 32-bit: (HP-UX.MAX_PERM_SIZE) • HP-UX 64-bit: (HP-UX-64.MAX_PERM_SIZE) • IBM AIX 32-bit: (AIX.MAX_PERM_SIZE) • IBM AIX 64-bit: (AIX-64.MAX_PERM_SIZE) • Sun Solaris 32-bit: (SunOS.MAX_PERM_SIZE) • Sun Solaris 64-bit: (SunOS-64.MAX_PERM_SIZE) • Linux 32-bit: (Linux.MAX_PERM_SIZE) • Linux 64-bit: (Linux-64.MAX_PERM_SIZE) • Windows 32-bit: (Windows.MAX_PERM_SIZE) • Windows 64-bit: (Windows-64.MAX_PERM_SIZE) • IBM iSeries OS/400: (OS400.MAX_PERM_SIZE) • IBM OS/390: (OS390.MAX_PERM_SIZE) **tuning.properties Property:** [operating system].MAX_PERM_SIZE	Based on your operating system and architecture, Sterling B2B Integrator allocates the following values (in MB): • HP-UX (32-bit and 64-bit): 512 • IBM AIX (32-bit and 64-bit): 384 • Sun Solaris (32-bit and 64-bit): 384 • Linux (32-bit and 64-bit): 384 • Windows 32-bit: 256 • Windows 64-bit: 384 • IBM iSeries OS/400: 384 • IBM OS/390: 384

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Thread Stack Size (KB) Controls the maximum Java stack size for any thread. **tuning.properties Property:** [operating system].STATIC_STACK	Based on your operating system and architecture (in KB): • HP-UX 32-bit: 384 • HP-UX 64-bit: 512 • IBM AIX (32-bit and 64-bit): 256 • Sun Solaris (32-bit and 64-bit): 256 • Linux (32-bit and 64-bit): 256 • IBM iSeries OS/400 (32-bit and 64-bit): 256 • IBM OS/390 (32-bit and 64-bit): 256 • Windows (32-bit and 64-bit): 256
Enable Heap Dump on Out of Memory Select **True** from the drop-down list if you want to generate a heapdump when the server JVM encounters an out of memory error. 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, the dump file is generated from where the JVM was started (the *install_dir*/bin directory). If you select **False** from the drop-down list, heapdumps are not generated when the server JVM encounters an out of memory error. **tuning.properties Property:** ENABLE_HEAP_DUMP_OOB	Default value: True

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Enable Heap Dump on Sigquit/Ctrl-break **Note:** This field is not supported in the 1.6 version of the Sun JDK. It will display if you are using the IBM JDK.Select **True** from the drop-down list if you want to request a heapdump for the server JVM using **`Sigquit`** command or pressing **Ctrl+Break** for server JVM. By default, the dump file is generated from where the JVM was started (the *install_dir*/bin directory). If you select **False** from the drop-down list, heapdumps are not generated if you run **`Sigquit`** command or press **Ctrl+Break**. **tuning.properties Property:** ENABLE_HEAP_DUMP_SGTQCRTLBRK	Default value: True
Enable VerboseGC Select **True** from the drop-down list if you want to generate verbose garbage collection data for server JVM. JVM Verbose garbage collection (GC) statistics are critical. It is recommended to enable verbose garbage collection when you are debugging your instance. These statistics can be used to understand the behavior of the JVM heap management and the efficiency of the JVM. If you select **False** from the drop-down list, verbose garbage collection will be disabled. **tuning.properties Property:** ENABLE_VERBOSE_GC	Default value: False
JVM Arguments Prefix Determines the arguments set as JVM parameters before other JVM flags are set for server JVM. Ensure that you enter correct values in this field. Else, the server JVM might fail to start. For example, **`-XX:SurvivorRatio`** parameter optimizes the object creation and object preservation ratio. **tuning.properties Property:** JVM_ARGS_PREFIX	Default value: BLANK

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
JVM Arguments Suffix Determines the arguments set as JVM parameters after other JVM flags are set for server JVM. Ensure that you enter correct values in this field. Else, the server JVM might fail to start. For example, **-Xnoclassgc** parameter disables dynamic class unloading. **tuning.properties Property:** JVM_ARGS_SUFFIX	Default value: BLANK

Click **Next**.

7. In the JVM Parameters for Container Nodes page that is displayed, decide whether you want to use the current settings or the suggested settings.
   - If yes, click **Fill Suggested**.
   - If no, enter information in the following fields:

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Physical Memory Allocated to Container (MB) Amount of physical memory allocated to each container node. Ensure that the physical memory allocated to server nodes, container nodes, and maximum permanent generation size for both server nodes and container nodes should not exceed the physical memory allocated to Sterling B2B Integrator. If you add a container to the system after the initial installation, you will have to ensure that the proper amount of memory has been allocated to the container nodes, using this field. **tuning.properties Property:** MEMORY_CONTAINER	Default value (MB): 512 A minimum of 512 MB is allocated to each container node out of the total physical memory allocated to Sterling B2B Integrator. However, if there is additional physical memory available after allocating minimum memory to all the server nodes and the container nodes, the memory is split between server nodes and container nodes. By default, the server nodes are allocated with three-fourths out of the available memory and the container nodes are allocated with one-fourths out of the available memory. The additional memory is calculated using the following formula: ((((Physical memory allocated to Platform - (Number of server nodes * Minimum memory allocated to the server nodes)) - (Number of server nodes * Maximum permanent generation size allocated to server nodes)) - (Number of container nodes * Maximum permanent generation size allocated to container nodes)) – Memory allocated to the ops server) For example, let us consider that two server nodes and one container node is present. After allocating memory to the two server nodes and one container node, 200 MB additional memory is available, each server node is allocated 75 MB memory and the container node is allocated 50 MB memory. You can also change the additional memory allocation ratio. The additional memory allocation ratio is calculated using the following formula: (((Additional memory available - Physical memory allocated to container nodes) * 0.25)/Number of container nodes)

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)

Initial Heap Size (MB)
 	
Controls the initial size of a Java heap allocated to the container JVM.
 	
Based on your operating system and architecture, the parameters are provided in the following list:
• HP-UX 32-bit: (HP-UX.INIT_HEAP_CONTAINER)
• HP-UX 64-bit: (HP-UX-64.INIT_HEAP_CONTAINER)
• IBM AIX 32-bit: (AIX.INIT_HEAP_CONTAINER)
• IBM AIX 64-bit: (AIX-64.INIT_HEAP_CONTAINER)
• Sun Solaris 32-bit: (SunOS.INIT_HEAP_CONTAINER)
• Sun Solaris 64-bit: (SunOS-64.INIT_HEAP_CONTAINER)
• Linux 32-bit: (Linux.INIT_HEAP_CONTAINER)
• Linux 64-bit: (Linux-64.INIT_HEAP_CONTAINER)
• Windows 32-bit: (Windows.INIT_HEAP_CONTAINER)
• Windows 64-bit: (Windows-64.INIT_HEAP_CONTAINER)
• IBM iSeries OS/400: (OS400.INIT_HEAP_CONTAINER)
• IBM OS/390: (OS390.INIT_HEAP_CONTAINER)
 	
tuning.properties Property:
 	
[operating system].INIT_HEAP_CONTAINER
	Based on your operating system and architecture, the formulae are provided in the following list:
	• HP-UX 32-bit:
	If memory allocated to the container node is greater than or equal to 3068 MB, it allocates 3068 MB. Else, if the memory allocated to container node is lesser than or equal to 512 MB, it allocates the memory assigned to the container node.
	• HP-UX 64-bit:
	if (MEMORY_CONTAINER >= 3068, 3068, if (MEMORY_CONTAINER <= 512, 512, MEMORY_CONTAINER))
	• IBM AIX 32-bit:
	if (MEMORY_CONTAINER >= 3068, MEMORY_CONTAINER, if (MEMORY_CONTAINER <= 512, 512, MEMORY_CONTAINER))
	• IBM AIX 64-bit:
	if (MEMORY_CONTAINER >= 3328, MEMORY_CONTAINER, if (MEMORY_CONTAINER <= 512, 512, MEMORY_CONTAINER))
	• Sun Solaris 32-bit:
	if (MEMORY_CONTAINER >= 3900, 3900, if (MEMORY_CONTAINER <= 512, 512, MEMORY_CONTAINER))
	• Sun Solaris 64-bit:
	if (MEMORY_CONTAINER >= 1536, MEMORY_CONTAINER, if (MEMORY_CONTAINER <= 1024, 1024, MEMORY_CONTAINER))

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)

	• Linux 32-bit:
	(0.5 * Physical memory allocated to container node)
	• Linux 64-bit:
	if (MEMORY_CONTAINER >= 2036, MEMORY_CONTAINER, if (MEMORY_CONTAINER <= 512, 512, MEMORY_CONTAINER))
	• Windows 32-bit:
	(0.5 * Physical memory allocated to container node)
	• Windows 64-bit:
	if (MEMORY_CONTAINER >= 1536, 768, if (MEMORY_CONTAINER <= 512, MEMORY_CONTAINER, round(0.5 * MEMORY_CONTAINER)))
	• IBM iSeries OS/400:
	If memory allocated to the container node is greater than or equal to 2048 MB, it allocates 2048 MB. Else, if the memory allocated to container node is lesser than or equal to 512 MB, it allocates the memory assigned to the container node.
	• IBM OS/390:
	If memory allocated to the container node is greater than or equal to 2048 MB, it allocates 2048 MB. Else, if the memory allocated to container node is lesser than or equal to 512 MB, it allocates the memory assigned to the container node.

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)

Maximum Heap Size (MB)
 	
Controls the maximum size of a Java heap allocated to the container JVM.
 	
Based on your operating system and architecture, the parameters are provided in the following list:
• HP-UX 32-bit: (HP-UX.MAX_HEAP_CONTAINER)
• HP-UX 64-bit: (HP-UX-64.MAX_HEAP_CONTAINER)
• IBM AIX 32-bit: (AIX.MAX_HEAP_CONTAINER)
• IBM AIX 64-bit: (AIX-64.MAX_HEAP_CONTAINER)
• Sun Solaris 32-bit: (SunOS.MAX_HEAP_CONTAINER)
• Sun Solaris 64-bit: (SunOS-64.MAX_HEAP_CONTAINER)
• Linux 32-bit: (Linux.MAX_HEAP_CONTAINER)
• Linux 64-bit: (Linux-64.MAX_HEAP_CONTAINER)
• Windows 32-bit: (Windows.MAX_HEAP_CONTAINER)
• Windows 64-bit: (Windows-64.MAX_HEAP_CONTAINER)
• IBM iSeries OS/400: (OS400.MAX_HEAP_CONTAINER)
• IBM OS/390: (OS390.MAX_HEAP_CONTAINER)
 	
tuning.properties Property:
 	
[operating system].MAX_HEAP_CONTAINER
	Based on your operating system and architecture, the formulae are provided in the following list:
	• HP-UX 32-bit:
	If memory allocated to the container node is greater than or equal to 3068 MB, it allocates 3068 MB. Else, if the memory allocated to container node is lesser than or equal to 512 MB, it allocates the memory assigned to the container node.
	• HP-UX 64-bit:
	if (MEMORY_CONTAINER >= 3068, MEMORY_CONTAINER, if (MEMORY_CONTAINER <= 512, 512, MEMORY_CONTAINER))
	• IBM AIX 32-bit:
	If memory allocated to the container node is greater than or equal to 3328 MB, it allocates 3328 MB. Else, if the memory allocated to server node is lesser than or equal to 512 MB, it allocates the memory assigned to the container node.
	• IBM AIX 64-bit:
	if (MEMORY_CONTAINER >= 3328, MEMORY_CONTAINER, if (MEMORY_CONTAINER <= 512, 512, MEMORY_CONTAINER))
	• Sun Solaris 32-bit:
	If memory allocated to the container node is greater than or equal to 3900 MB, it allocates 3900 MB. Else, if the memory allocated to container node is lesser than or equal to 512 MB, it allocates the memory assigned to the container node.
	• Sun Solaris 64-bit:
	if (MEMORY_CONTAINER >= 3900, MEMORY_CONTAINER, if (MEMORY_CONTAINER <= 512, 512, MEMORY_CONTAINER))

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)

	• Linux 32-bit:
	If memory allocated to the container node is greater than or equal to 2036 MB, it allocates 2036 MB. Else, if the memory allocated to container node is lesser than or equal to 512 MB, it allocates the memory assigned to the container node.
	• Linux 64-bit:
	if (MEMORY_CONTAINER >= 2036, MEMORY_CONTAINER, if (MEMORY_CONTAINER <= 512, 512, MEMORY_CONTAINER))
	• Windows 32-bit:
	If memory allocated to the container node is greater than or equal to 1536 MB, it allocates 1536 MB. Else, if the memory allocated to container node is lesser than or equal to 512 MB, it allocates the memory assigned to the container node.
	• Windows 64-bit:
	if (MEMORY_CONTAINER >= 1536, MEMORY_CONTAINER, if (MEMORY_CONTAINER <= 512, 512, MEMORY_CONTAINER))
	• IBM iSeries OS/400:
	If memory allocated to the container node is greater than or equal to 2048 MB, it allocates 2048 MB. Else, if the memory allocated to container node is lesser than or equal to 512 MB, it allocates the memory assigned to the container node.
	• IBM OS/390:
	If memory allocated to the container node is greater than or equal to 2048 MB, it allocates 2048 MB. Else, if the memory allocated to container node is lesser than or equal to 512 MB, it allocates the memory assigned to the container node.

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Initial Size of New (Nursery) Heap (MB) Controls the initial size of the new heap area allocated to the container JVM. Based on your operating system and architecture, the parameters are provided in the following list: • HP-UX 32-bit: (HP-UX.INIT_AGE_CONTAINER) • HP-UX64-bit: (HP-UX-64.INIT_AGE_CONTAINER) • IBM AIX 32-bit: (AIX.INIT_AGE_CONTAINER) • IBM AIX 64-bit: (AIX-64.IN IT_AGE_CONTAINER) • Sun Solaris 32-bit: (SunOS.INIT_AGE_CONTAINER) • Sun Solaris 64-bit: (SunOS-64.INIT_AGE_CONTAINER) • Linux 32-bit: (Linux.INIT_AGE_CONTAINER) • Linux 64-bit: (Linux-64.INIT_AGE_CONTAINER) • Windows 32-bit: (Windows.IN IT_AGE_CONTAINER) • Windows 64-bit: (Windows-64.INIT_AGE_CONTAINER) • IBM iSeries OS/400: (OS400.INIT_AGE_CONTAINER) • IBM OS/ 390(OS390.INIT_AGE_CONTAINER) **tuning.properties Property:** [operating system].INIT_AGE_CONTAINER	(0.33333 * Initial size of heap memory allocated to the container node) • HP 32-bit: HP-UX.INIT_AGE_CONTAINER= round(0.33333 * HP-UX.INIT_HEAP_CONTAINER) • HP 64-bit: HP-UX-64.INIT_AGE_CONTAINER= round(0.33333 * HP-UX-64.INIT_HEAP_CONTAINER) • AIX 32-bit: AIX.INIT_AGE_CONTAINER= round(0.33333 * AIX.INIT_HEAP_CONTAINER) • AIX 64-bit: AIX-64.INIT_AGE_CONTAINER= round(0.33333 * AIX-64.INIT_HEAP_CONTAINER) • Sun 32-bit: SunOS.INIT_AGE_CONTAINER= round(0.33333 * SunOS.INIT_HEAP_CONTAINER) • Sun 64-bit: SunOS-64.INIT_AGE_CONTAINER= round(0.33333 * SunOS-64.INIT_HEAP_CONTAINER) • Linux 32-bit: Linux.INIT_AGE_CONTAINER= round(0.33333 * Linux.INIT_HEAP_CONTAINER) • Linux 64-bit: Linux-64.INIT_AGE_CONTAINER= round(0.33333 * Linux-64.INIT_HEAP_CONTAINER) • Windows 32-bit: Windows.INIT_AGE_CONTAINER= round(0.33333 * Windows.INIT_HEAP_CONTAINER) • Windows 64-bit: Windows-64.INIT_AGE_CONTAINER= round(0.33333 * Windows-64.INIT_HEAP_CONTAINER)

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Maximum Size of New (Nursery) Heap (MB) Controls the maximum size of the new heap area allocated to the container JVM. Based on your operating system and architecture the parameters are provided in the following list: • HP-UX 32-bit: (HP-UX.MAX_AGE_CONTAINER) • HP-UX 64-bit: (HP-UX-64.MAX_AGE_CONTAINER) • IBM AIX 32-bit: (AIX.MAX_AGE_CONTAINER) • IBM AIX 64-bit: (AIX-64.MAX_AGE_CONTAINER) • Sun Solaris 32-bit: (SunOS.MAX_AGE_CONTAINER) • Sun Solaris 64-bit: (SunOS-64.MAX_AGE_CONTAINER) • Linux 32-bit: (Linux.MAX_AGE_CONTAINER) • Linux 64-bit: (Linux-64.MAX_AGE_CONTAINER) • Windows 32-bit: (Windows.MAX_AGE_CONTAINER) • Windows 64-bit: (Windows-64.MAX_AGE_CONTAINER) • IBM iSeries OS/400: (OS400.MAX_AGE_CONTAINER) • IBM OS/390: (OS390.MAX_AGE_CONTAINER) **tuning.properties Property:** [operating system].MAX_AGE_CONTAINER	(0.5 * Maximum size of heap memory allocated to the container node) • HP 32-bit: HP-UX.MAX_AGE_CONTAINER= round(0.33333 * HP-UX.MAX_HEAP_CONTAINER) • HP 64-bit: HP-UX-64.MAX_AGE_CONTAINER= round(0.33333 * HP-UX-64.MAX_HEAP_CONTAINER) • AIX 32-bit: AIX.MAX_AGE_CONTAINER= round(0.33333* AIX.MAX_HEAP_CONTAINER) • AIX 64-bit: AIX-64.MAX_AGE_CONTAINER= round(0.33333* AIX-64.MAX_HEAP_CONTAINER) • Sun 32-bit: SunOS.MAX_AGE_CONTAINER= round(0.33333 * SunOS.MAX_HEAP_CONTAINER) • Sun 64-bit: SunOS-64.MAX_AGE_CONTAINER= round(0.33333 * SunOS-64.MAX_HEAP_CONTAINER) • Linux 32-bit: Linux.MAX_AGE_CONTAINER= round(0.33333 * Linux.MAX_HEAP_CONTAINER) • Linux 64-bit: Linux-64.MAX_AGE_CONTAINER= round(0.33333 * Linux-64.MAX_HEAP_CONTAINER) • Windows 32-bit: Windows.MAX_AGE_CONTAINER= round(0.33333 * Windows.MAX_HEAP_CONTAINER) • Windows 64-bit: Windows-64.MAX_AGE_CONTAINER= round(0.33333 * Windows-64.MAX_HEAP_CONTAINER)

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Maximum Permanent Generation Size (MB) Controls the maximum size of the permanent generation heap (stores all the class code and class-like data) allocated to the container JVM. Based on your operating system and architecture, the parameters are provided in the following list: • HP-UX 32-bit: (HP-UX.MAX_PERM_SIZE_CONTAINER) • HP-UX 64-bit: (HP-UX-64.MAX_PERM_SIZE_CONTAINER) • IBM AIX 32-bit: (AIX.MAX_PERM_SIZE_CONTAINER) • IBM AIX 64-bit: (AIX-64.MAX_PERM_SIZE_CONTAINER) • Sun Solaris 32-bit: (SunOS.MAX_PERM_SIZE_CONTAINER) • Sun Solaris 64-bit: (SunOS-64.MAX_PERM_SIZE_CONTAINER) • Linux 32-bit: (Linux.MAX_PERM_SIZE_CONTAINER) • Linux 64-bit: (Linux-64.MAX_PERM_SIZE_CONTAINER) • Windows 32-bit: (Windows.MAX_PERM_SIZE_ CONTAINER) • Windows 64-bit: (Windows-64.MAX_PERM_SIZE_ CONTAINER) • IBM iSeries OS/400: (OS400.MAX_PERM_SIZE_CONTAINER) • IBM OS/390: (OS390.MAX_PERM_SIZE_CONTAINER) **tuning.properties Property:** [operating system].MAX_PERM_SIZE_CONTAINER	Based on your operating system and architecture, Sterling B2B Integrator allocates the following values (in MB): • HP-UX (32-bit and 64-bit): 192 • IBM AIX (32-bit and 64-bit): 128 • Sun Solaris (32-bit and 64-bit): 128 • Linux (32-bit and 64-bit): 128 • Windows (32-bit and 64-bit): 128 • IBM iSeries OS/400: 128 • IBM OS/390: 128
Thread Stack Size (KB) Controls the maximum Java stack size for any thread. **tuning.properties Property:** [operating system].STATIC_STACK_CONTAINER	Based on your operating system and architecture (in KB): • HP-UX (32-bit and 64-bit): 384 • IBM AIX (32-bit and 64-bit): 256 • Sun Solaris (32-bit and 64-bit): 256 • Linux (32-bit and 64-bit): 256 • IBM iSeries OS/400 (32-bit and 64-bit): 256 • IBM OS/390 (32-bit and 64-bit): 256 • Windows (32-bit and 64-bit): 256

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Enable Heap Dump on Out of Memory Select **True** from the drop-down list if you want to generate a heapdump when the container JVM encounters an out of memory error. 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, the dump file is generated from where the JVM was started (the *install_dir*/bin directory). If you select **False** from the drop-down list, heapdumps are not generated when the container JVM encounters an out of memory error. **tuning.properties Property:** ENABLE_HEAP_DUMP_OOB_CONTAINER	Default value: True
Enable Heap Dump on Sigquit/Ctrl-break **Note:** This field is not supported in the 1.6 version of the Sun JDK. It will display if you are using the IBM JDK.Select **True** from the drop-down list if you want to request a heapdump for the container JVM using **`Sigquit`** command or pressing **Ctrl+Break** for container JVM. By default, the dump file is generated from where the JVM was started (the *install_dir*/bin directory). If you select **False** from the drop-down list, heapdumps are not generated if you run **`Sigquit`** command or press **Ctrl+Break** . **tuning.properties Property:** ENABLE_HEAP_DUMP_ SGTQCRTLBRK_CONTAINER	Default value: True

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Enable VerboseGC Select **True** from the drop-down list if you want to generate verbose garbage collection data for container JVM. JVM Verbose garbage collection (GC) statistics are critical. It is recommended to enable verbose garbage collection when you are debugging your instance. These statistics can be used to understand the behavior of the JVM heap management and the efficiency of the JVM. If you select **False** from the drop-down list, verbose garbage collection will be disabled. **tuning.properties Property:** ENABLE_VERBOSE_GC_CONTAINER	Default value: False
JVM Arguments Prefix Determines the arguments set as JVM parameters before other JVM flags are set for container JVM. Ensure that you enter correct values in this field. Else, the container JVM might fail to start. For example, **-Xalwaysclassgc** parameter enables dynamic class unloading on every instance of garbage collection. **tuning.properties Property:** JVM_ARGS_PREFIX_CONTAINER	Default value: BLANK
JVM Arguments Suffix Determines the arguments set as JVM parameters after other JVM flags are set for container JVM. Ensure that you enter correct values in this field. Else, the container JVM might fail to start. For example, **-Xlp** parameter enables large page support. **tuning.properties Property:** JVM_ARGS_SUFFIX_CONTAINER	Default value: BLANK

Click **Next**.

8. In the BP Queue page that is displayed, decide whether you want to use the current settings or the suggested settings.

Download from Www.Somanuals.com. All Manuals Search And Download.

- If yes, click **Fill Suggested**.
- If no, enter information in the following fields:

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Number of Processor Cores Number of processor cores in the system, where each core is an independent processing engine. **tuning.properties Property:** PROCESSORS	Not Applicable
Desired Global Threads Number of active business processes that can run concurrently. This field changes the JMS_LISTNERS property in the tuning.properties file in the *install_dir*/properties directory. 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. **tuning.properties Property:** JMS_LISTNERS	JMS_LISTNERS = (PROCESSORS * 4)
In memory cache size (MB) for small contexts Amount of memory allocated to the business process queue. After Sterling B2B 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 *install_dir*/properties directory. Increasing this value too much reduces the amount of memory available for the other components of Sterling B2B Integrator, such as caches and general processing. **tuning.properties Property:** NOAPP.JMS_PAGING_HIGH	NOAPP.JMS_PAGING_HIGH = 100

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Disk cache size (MB) Maximum Maximum amount of disk space used to store the business process context when it is being moved from memory. This limit ensures that Sterling B2B Integrator does not consume all the available disk space. This field changes the JMS_PAGING_MAX property in the tuning.properties file in the *install_dir*/properties directory. **tuning.properties Property:** NOAPP.JMS_PAGING_MAX	NOAPP.JMS_PAGING_MAX = round(2 * Maximum Java heap size allocated to the server JVM)
Memory Cache Threshold (MB) Amount of threshold memory size for caching a context in the in-memory cache versus the disk cache. The distribution of context sizes tends to look vaguely like an "M" with one cluster of small contexts and another cluster of larger contexts. If this value is set large, the value for memory cache size should also be large. A moderate multiple of the disk block size seems to work very well in many cases. **tuning.properties Property:** BP_SIZE	Default value: 16384
Use Shadow Cache Select **True** from the drop-down list if you want to enable shadow cache. It is used to retain business process objects in the memory as long as JVM permits. When the thread is available for the business process to run, the object is retrieved from the memory instead of the disk. Use Shadow Cache plays an important role during load balancing in a cluster environment. It uses soft reference for caching workflow context objects. If you select **False** from the drop-down list, shadow cache will be disabled and space will not be allocated to retain business process objects in the memory. **tuning.properties Property:** USE_SHADOW_CACHE	Default value: True

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Shadow Cache Size (MB) Maximum Amount of soft reference in memory to retain business process objects in the memory as long as JVM permits. When the thread is available for the business process to run, the object is retrieved from the memory instead of the disk. You can set maximum shadow cache size, but it is used only if you enable shadow cache. **tuning.properties Property:** MAX_SHADOW_CACHE_SIZE	Default value: 1
Distribution Cache Maximum Determines the maximum number of business processes in queue on node2. If the number of business processes on node2 queue exceeds this number, node1 will not transfer any more business processes to node2. **tuning.properties Property:** DISTRIBUTION_CACHE_MAX	Default value: 5000
Distribution Cache Minimum Determines the number of business processes that must be in queue before this node will make any processes eligible for distribution. **tuning.properties Property:** DISTRIBUTION_CACHE_MIN	DISTRIBUTION_CACHE_MIN = (PROCESSORS * 4)

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Broadcast Interval (ms) Determines how often (in milleseconds) a multicast message is sent out by a node about the current state of the node. The multicast message provides load factor and current status information about a node to all other nodes in the cluster. Clustering uses the JGroups communication toolkit, which supports the following communication protocols, which can be configured during deployment: • UDP (User Datagram Protocol) (IP multicast) • TCP **tuning.properties Property:** MULTICAST_INTERVAL	Default value: 6000
Switch Interval (ms) Determines the time interval (in milliseconds) that a business process must be queued after distribution to a node before the business process is eligible for distribution to another node again. This is for load balancing in a cluster environment. **tuning.properties Property:** SWITCH_INTERVAL	Default value: 180000

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Workflow in transit list timeout (ms) Allows the business process to transfer from one node to another without the Recover process interrupting the business process. In a cluster, when a business process is switched to another node due to a heavy load, the workflow id will be "seen" by the Recover process as still in the queue on this node for the amount of time (in milliseconds) specified by this parameter. This allows the business process to finish transfering to the other node without the Recover process interrupting the business process. Increase this value if the Recover process is interrupting business processes during a heavy load in a clustered environment. This may make the UI list the business process as 'in queue' but the business process can actually be running. This will not interfere with the running result. It may only delay displaying the business process state on the UI. Setting this value to **0** will prevent this functionality from being used and will keep an extra thread that performs part of this functionality from being started. **tuning.properties Property:** WF_IN_TRANSIT_TIMEOUT_LIST	Default value: 20000

Click **Next**.

9. In the Cache page that is displayed, decide whether you want to use the current settings or the suggested settings.
   - If yes, click **Fill Suggested**.
   - If no, enter information in the following fields:

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Translation Maps The number of translation maps that can exist in the cache. This field changes the MAP_CACHE property in the tuning.properties file in the *install_dir*/properties directory. **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. **tuning.properties Property:** MAP_CACHE	If memory allocated to the server node is lesser than or equal to 768 MB, 10 translation maps can exist in the cache. Else, if the memory allocated to server node is lesser than or equal to 1024 MB but greater than 768 MB, 50 translation maps can exist in the cache. If the memory allocated to the server node is greater than 1024 MB, then 100 translation maps can exist in the cache.
Envelopes The number of envelopes that can exist in the cache. This field changes the ENVELOPE_CACHE property in the tuning.properties file in the *install_dir*/properties directory. **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. **tuning.properties Property:** ENVELOPE_CACHE	If memory allocated to the server node is lesser than or equal to 768 MB, 50 envelopes can exist in the cache. Else, if the memory allocated to server node is lesser than or equal to 1024 MB but greater than 768 MB, 250 envelopes can exist in the cache. If the memory allocated to the server node is greater than 1024 MB, then 500 envelopes can exist in the cache.
Other EDI The number of EDI objects that can exist in the cache. This field changes the EDI_CACHE property in the tuning.properties file in the *install_dir*/properties directory. **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. **tuning.properties Property:** EDI_CACHE	If memory allocated to the server node is lesser than or equal to 768 MB, 25 EDI objects can exist in the cache. Else, if the memory allocated to server node is lesser than or equal to 1024 MB but greater than 768 MB, 50 EDI objects can exist in the cache. If the memory allocated to the server node is greater than 1024 MB, then 100 EDI objects can exist in the cache.

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Workflow Definition The number of workflow definition objects that can exist in the cache. This field changes the WFD_CACHE property in the tuning.properties file in the *install_dir*/properties directory. **tuning.properties Property:** WFD_CACHE	If memory allocated to the server node is lesser than or equal to 768 MB, 50 workflow definition objects can exist in the cache. Else, if the memory allocated to server node is lesser than or equal to 1024 MB but greater than 768 MB, 100 workflow definition objects can exist in the cache. If the memory allocated to the server node is greater than 1024 MB, then 120 workflow definition objects can exist in the cache.
User Cache The number of user definition objects that can exist in the cache. This field changes the USER_CACHE property in the tuning.properties file in the *install_dir*/properties directory. **tuning.properties Property:** USER_CACHE	If memory allocated to the server node is lesser than or equal to 768 MB, 50 user definition objects can exist in the cache. Else, if the memory allocated to server node is lesser than or equal to 1024 MB but greater than 768 MB, 100 user definition objects can exist in the cache. If the memory allocated to the server node is greater than 1024 MB, then 120 user definition objects can exist in the cache.
Group Cache The number of user group definition objects that can exist in the cache. This field changes the GROUP_CACHE property in the tuning.properties file in the *install_dir*/properties directory. **tuning.properties Property:** GROUP_CACHE	If memory allocated to the server node is lesser than or equal to 768 MB, 50 user group definition objects can exist in the cache. Else, if the memory allocated to server node is lesser than or equal to 1024 MB but greater than 768 MB, 100 user group definition objects can exist in the cache. If the memory allocated to the server node is greater than 1024 MB, then 120 user group definition objects can exist in the cache.
Permission Cache The number of permission definition objects that can exist in the cache. Permission cache controls the number of cache protectors and stores information on user permissions. This field changes the PERM_CACHE property in the tuning.properties file in the *install_dir*/properties directory. **tuning.properties Property:** PERM_CACHE	If memory allocated to the server node is lesser than or equal to 768 MB, 1000 permission definition objects can exist in the cache. Else, if the memory allocated to server node is lesser than or equal to 1024 MB but greater than 768 MB, 2000 permission definition objects can exist in the cache. If the memory allocated to the server node is greater than 1024 MB, then 2400 permission definition objects can exist in the cache.

Click **Next**.

10. In the Purge Settings page that is displayed, decide whether you want to use the current settings or the suggested settings.
    - If yes, click **Fill Suggested**.
    - If no, enter information in the following fields:

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Number Of Purge Deadlock Retries Controls the number of repeated attempts to start the purge process before aborting. **tuning.properties Property:** PURGE.DEADLOCK_RETRIES	Default value: 30
Maximum Purge Threads Controls the maximum number of threads that can be allocated to the purge process. **tuning.properties Property:** PURGE.MAX_PURGE_THREADS	Default value: 1
Purge Documents on Disk Select **True** from the drop-down list to purge the documents on file system disk space, along with database disk space. It is recommended to purge documents on the disk, else the disk space may get filled up. For more information about purging, refer to *Purge Service* documentation. If you select **False** from the drop-down list, the documents are purged only in the database disk space. **tuning.properties Property:** PURGE.PURGE_DOCS_ON_DISK	Default value: True

Click **Next**.

11. In the Database Connection Pool page that is displayed, decide whether you want to use the current settings or the suggested settings.

    Note: Sterling B2B Integrator restart is not required if you modify settings in this page.

    - If yes, click **Fill Suggested**.
    - 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

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.

Database connection pool values are initially displayed by obtaining the values from jdbc.properties file and jdbc_customer.properties file. When you set the connection pool properties, such as initial connections, maximum connections, and buffer size, the performance tuning utility updates system_overrrides.properties file, which will be created if it does not exist.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Default Document Storage Determines the storage type for document objects constructed without a storageType property. Valid values in the **New** field at the top of the table: • FS - Stores the document data on the local file system • DB - (Default) Stores the document data in the database **tuning.properties Property:** jdbcService.defaultDocumentStorageType	jdbcService.defaultDocumentStorageType = DB

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Connection Pool Name Determines the database connection pool name. A database connection pool is created to manage multiple connections where each connection connects to the same database. You can create custom connection pools apart from those present in Sterling B2B Integrator. If you create your own custom connection pools, you should add the connection pool name in language properties file. The following connection pool types can be created: • UI • Archive • Default • Select • Non-transactional - Number of connections that are reserved and available for general data requests (for example, retrieving data) when Sterling B2B Integrator starts. If Sterling B2B Integrator uses all the connections concurrently, it creates a new connection until it reaches the maximum number of connections. When Sterling B2B 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. When a connection is no longer in use, the connection is returned to the available connection pool and can be reused. • Local	Not applicable

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Initial Connections Determines the initial number of connections to be reserved in the database connection pool. Based on the connection pool, the parameters are provided in the following list: • Local: (jdbcService.*db*Pool_local.initsize) • Select: (jdbcService.*db*Pool_Select.initsize) • Non-transactional: (jdbcService.*db*Pool_NoTrans.initsize) • Default: (jdbcService.*db*Pool.initsize) • Archive: jdbcService.*db*ArchivePool.initsize) • UI: (jdbcService.*db*UIPool.initsize)	Depending on the database, replace the *db* with the database name. For example, if you are using Oracle as your database, enter **oracle**. Default values are provided in the following list: • Local: jdbcService.*db*Pool_local.initsize = 1 • Select: jdbcService.*db*Pool_Select.initsize = 0 • Non-transactional: jdbcService.*db*Pool_NoTrans.initsize = 1 • Default: jdbcService.*db*Pool.initsize = 1 • Archive: jdbcService.*db*ArchivePool.initsize = 1 • UI: jdbcService.*db*UIPool.initsize = 4
Maximum Connections Determines the maximum number of connections to be reserved in the database connection pool. Based on the connection pool, the parameters are provided in the following list: • Local: (jdbcService.*db*Pool_local.maxsize) • Select: (jdbcService.*db*Pool_Select.maxsize) • Non-transactional: (jdbcService.*db*Pool_NoTrans.maxsize) • Default: (jdbcService.*db*Pool.maxsize) • Archive: (jdbcService.*db*ArchivePool.maxsize) • UI: (jdbcService.*db*UIPool.maxsize)	Depending on the database, replace the *db* with the database name. For example, if you are using Oracle as your database, enter **oracle**. Default values are provided in the following list: • Local: jdbcService.*db*Pool_local.maxsize = 30 • Select: jdbcService.*db*Pool_Select.maxsize = 20 • Non-transactional: jdbcService.*db*Pool_NoTrans.maxsize = 30 • Default: jdbcService.*db*Pool.maxsize = 30 • Archive: jdbcService.*db*ArchivePool.maxsize = 20 • UI: jdbcService.*db*UIPool.maxsize = 20

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Buffer Size Determines the number of additional connections reserved by the database connection pool. These additional connections are released when there is an unexpected demand for more connections. The buffer connections are created in addition to the maximum connections reserved. Based on the connection pool, the parameters are provided in the following list: • Local: (jdbcService.*db*Pool_local.buffersize = 500) • Select: (jdbcService.*db*Pool_Select.buffersize = 500) • Non-transactional: (jdbcService.*db*Pool_NoTrans.buffersize = 500) • Default: (jdbcService.*db*Pool.buffersize = 500) • Archive: (jdbcService.*db*ArchivePool.buffersize = 500) • UI: (jdbcService.*db*UIPool.buffersize = 500)	Depending on the database, replace the *db* with the database name. For example, if you are using Oracle as your database, enter **oracle**. Default values are provided in the following list: • Local: jdbcService.*db*Pool_local.buffersize = 500 • Select: jdbcService.*db*Pool_Select.buffersize = 500 • Non-transactional: jdbcService.*db*Pool_NoTrans.buffersize = 500 • Default: jdbcService.*db*Pool.buffersize = 500 • Archive: jdbcService.*db*ArchivePool.buffersize = 500 • UI: jdbcService.*db*UIPool.buffersize = 500

Click **Next**.

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.

12. In the Queue Settings page that is displayed, decide whether you want to use the current settings or the suggested settings.

    Note: Sterling B2B Integrator restart is not required if you modify settings in this page.

    - If yes, click **Fill Suggested**.
    - If no, enter information in the fields in the following columns:

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Queue Name Name of the queue. Nine queues (Queue 1 to Queue 9) enable flexible management of your workload.	Not applicable

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Initial Steps in the first Execution Cycle 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 *install_dir*/properties directory. The parameters for each queue are provided as follows: • Queue [queue number]: (NOAPP.INITIAL_CYCLES_[queue number]) **tuning.properties Property:** NOAPP.INITIAL_CYCLES_[queue number]	The default values for each queue are provided in the following list: • Queue 1: NOAPP.INITIAL_CYCLES_1 = ASYNC_BP • Queue 2: NOAPP.INITIAL_CYCLES_2 = ASYNC_BP • Queue 3: NOAPP.INITIAL_CYCLES_3 = ASYNC_BP • Queue 4: NOAPP.INITIAL_CYCLES_4 = ASYNC_BP • Queue 5: NOAPP.INITIAL_CYCLES_5 = ASYNC_BP • Queue 6: NOAPP.INITIAL_CYCLES_6 = ASYNC_BP • Queue 7: NOAPP.INITIAL_CYCLES_7 = ASYNC_BP • Queue 8: NOAPP.INITIAL_CYCLES_8 = 10 • Queue 9: NOAPP.INITIAL_CYCLES_9 = ASYNC_BP

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Number of Steps in the Execution Cycle 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 *install_dir*/properties directory. 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. 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 B2B Integrator. The parameters for each queue are provided as follows: • Queue [queue number]: (NOAPP.EXEC_CYCLE_[queue number]) **tuning.properties Property:** NOAPP.EXEC_CYCLE_[queue number]	The default values for each queue are provided in the following list: • Queue 1: NOAPP.EXEC_CYCLE_1 = 1000 • Queue 2: NOAPP.EXEC_CYCLE_2 = 100 • Queue 3: NOAPP.EXEC_CYCLE_3 = 100 • Queue 4: NOAPP.EXEC_CYCLE_4 = BP_STEPS • Queue 5: NOAPP.EXEC_CYCLE_5 = BP_STEPS • Queue 6: NOAPP.EXEC_CYCLE_6 = 10 • Queue 7: NOAPP.EXEC_CYCLE_7 = 5 • Queue 8: NOAPP.EXEC_CYCLE_8 = 2 • Queue 9: NOAPP.EXEC_CYCLE_9 = 50

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Minimum Threads 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. The parameters for each queue are provided as follows: • Queue [queue number]: (NOAPP.MIN_POOL_SIZE_[queue number]) **tuning.properties Property:** NOAPP.MIN_POOL_SIZE_[queue number]	The default values for each queue are provided in the following list: • Queue 1: NOAPP.MIN_POOL_SIZE_1 = 0 • Queue 2: NOAPP.MIN_POOL_SIZE_2 = 0 • Queue 3: NOAPP.MIN_POOL_SIZE_3 = 0 • Queue 4: NOAPP.MIN_POOL_SIZE_4 = 1 • Queue 5: NOAPP.MIN_POOL_SIZE_5 = 1 • Queue 6: NOAPP.MIN_POOL_SIZE_6 = JMS_LISTNERS • Queue 7: NOAPP.MIN_POOL_SIZE_7 = (2 * JMS_LISTNERS) • Queue 8: NOAPP.MIN_POOL_SIZE_8 = (2 * JMS_LISTNERS) • Queue 9: NOAPP.MIN_POOL_SIZE_9 = 2

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Maximum Threads Maximum number of threads available for the queue. The parameters for each queue are provided as follows: • Queue [queue number]: (NOAPP.MAX_POOL_SIZE_[queue number]) **tuning.properties Property:** NOAPP.MAX_POOL_SIZE_[queue number]	The default values for each queue are provided in the following list: • Queue 1: NOAPP.MAX_POOL_SIZE_1 = 1 • Queue 2: NOAPP.MAX_POOL_SIZE_2 = (round(JMS_LISTNERS / 4) + 1) • Queue 3: NOAPP.MAX_POOL_SIZE_3 = (round(JMS_LISTNERS / 2) + 1) • Queue 4: NOAPP.MAX_POOL_SIZE_4 = JMS_LISTNERS • Queue 5: NOAPP.MAX_POOL_SIZE_5 = (2 * JMS_LISTNERS) • Queue 6: NOAPP.MAX_POOL_SIZE_6 = JMS_LISTNERS • Queue 7: NOAPP.MAX_POOL_SIZE_7 = (2 * JMS_LISTNERS) • Queue 8: NOAPP.MAX_POOL_SIZE_8 = (2 * JMS_LISTNERS) • Queue 9: NOAPP.MAX_POOL_SIZE_9 = 4

Download from Www.Somanuals.com. All Manuals Search And Download.

Field / Description	Tuning Formula or Default Value (Default value is provided if there is no tuning formula)
Resource Allocation 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. **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. The parameters for each queue are provided as follows: • Queue [queue number]: (NOAPP.RESOURCE_ ALLOCATION_[queue number]) **tuning.properties Property:** NOAPP.RESOURCE_ALLOCATION_[queue number]	The default values for each queue are provided in the following list: • Queue 1: NOAPP.RESOURCE_ALLOCATION_1 = 10 • Queue 2: NOAPP.RESOURCE_ALLOCATION_2 = 20 • Queue 3: NOAPP.RESOURCE_ALLOCATION_3 = 40 • Queue 4: NOAPP.RESOURCE_ALLOCATION_4 = 50 • Queue 5: NOAPP.RESOURCE_ALLOCATION_5 = 50 • Queue 6: NOAPP.RESOURCE_ALLOCATION_6 = 60 • Queue 7: NOAPP.RESOURCE_ALLOCATION_7 = 90 • Queue 8: N OAPP.RESOURCE_ALLOCATION_8 = 120 • Queue 9: NOAPP.RESOURCE_ALLOCATION_9 = 50

Click **Next**.

13. In the Confirm page that is displayed, click **View all settings** to review the changes you made to the performance configuration. The All Settings pop-up window is displayed.

    Click **Close** to close the All Settings pop-up window.

    The Confirm page displays both the properties modified, properties that require restart and properties that do not require restart.

14. In the Confirm page, 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**.

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

Download from Www.Somanuals.com. All Manuals Search And Download.

16. From the System Troubleshooting page, click **Soft Stop**. This runs the softstop script, allowing all active business processes to complete on a node before stopping the system. The script is run through a series of interactive steps in the UI. For more information about the softstop process, refer to "Completing a Soft Stop of Sterling B2B Integrator from the System Troubleshooter" on page 111.
17. In the tuningFormulas.properties file, change the TUNING_PROPS_UPDATED property to **TRUE**.
18. 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.
19. From the *install_dir*/bin directory, start Sterling B2B Integrator by running one of the following commands:
    - (UNIX or Linux) **run.sh**
    - (Windows) **startWindowsService.cmd**

    Sterling B2B 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 and related extension (*.ext) files in Sterling B2B Integrator.

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 IBM Sterling File Gateway, set the number of open file descriptors to at least 4096, using the command **ulimit -n 4096**.

Tuning Your ASI Environment

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

Properties that you may have to tune include:

SchedulingPolicyName

Set of rules that Sterling B2B Integrator uses to make decisions about how to manage the workload. Valid values are:

- BasicSchedulingPolicy
- FairShareSchedulingPolicy

Note: For more information about scheduling policy, refer to the topic "Scheduling Policy" on page 172.

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B 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 that 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

Download from Www.Somanuals.com. All Manuals Search And Download.

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.

CacheThreshold.#

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.

.# indicates the queue number.

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 B2B Integrator System Troubleshooter. For more information about reviewing system, refer to the topic Reviewing System Information.

There is a tradeoff between performance and memory consumption. Setting this value too high can result in one of the following:

- Leave Sterling B2B Integrator with insufficient memory in some circumstances
- Cause some business processes to remain in the queue for too long without being examined for rescheduling

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.

MinPoolSize.#

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.

.# indicates the queue number.

JavaPriority.#

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).

.# indicates the queue number.

EnableDeadlines.#

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:

- false – Disables deadline support for a queue.
- true – Enables deadline support for a queue.

This value can be different for each queue.

.# indicates the queue number.

Download from Www.Somanuals.com. All Manuals Search And Download.

Rescheduling.#

Enables or disables rescheduling support for a queue. Valid values are:

- false – Business processes are not rescheduled in a queue.
- true – Business processes are rescheduled in a queue.

This value can be different for each queue.

.# indicates the queue number.

ReschedulingInterval.#

Time interval, in milliseconds, to wait before the rescheduler passes through the queue to reschedule business processes. This value can be different for each queue.

.# indicates the queue number.

MaxWaitTime.#

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.

.# indicates the queue number.

Example: MaxWaitTime.8=60000 sets the maximum wait time for queue number 8 as 60 seconds.

ResourceAllocation.#

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.

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.

.# indicates the queue number.

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. This value can be different for each queue.

.# indicates the queue number.

StealThreads.#

Enables or disables the ability to steal threads from other queues. Valid values:

Download from Www.Somanuals.com. All Manuals Search And Download.

- 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 only from queues that have the AllowStealing property set to true. This value can be different for each queue.

.# indicates the queue number.

AllowStealing.#

Enables or disables the ability of other queues to steal threads from this queue. Valid values:

- 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) only steal from queues that have the AllowStealing property set to true. This value can be different for each queue.

.# indicates the queue number.

persistence_level

Persistence (storage to the database) level for Sterling B2B Integrator. Set the persistence_level value to PERSISTENCE_NONE during day-to-day processing.

Valid values are:

- PERSISTENCE_DEFAULT
- 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.
- 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 one of these workarounds:

  – Change the INVOKE_MODE of the second business process from INLINE to SYNC.

  – Add a dummy <assign> statement after the </operation> statement as shown:

```
<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>
```


Download from Www.Somanuals.com. All Manuals Search And Download.

```
<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>
```

  - Change the PERSISTENCE_LEVEL of the second business process from NONE (BP Start Stop - Engine May Override) to FULL.

- 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.
- 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 B2B 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.

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).

Download from Www.Somanuals.com. All Manuals Search And Download.

Scheduling Policy
About this task

Sterling B2B Integrator uses scheduling policies to manage workload. A scheduling policy is a set of rules that Sterling B2B 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 B2B Integrator:

Procedure

1. Open the *install_dir*/properties/noapp.properties_platform_ifcresources_ext file using a text editor.
2. 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 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

Sterling B2B 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 B2B 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 "Tuning Queues" on page 173.

Download from Www.Somanuals.com. All Manuals Search And Download.

You can set both a maximum and minimum thread number for each queue you use. Sterling B2B 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:

You can tune the queues in Sterling B2B Integrator.

About this task

To tune queues in Sterling B2B Integrator:

Procedure

1. Open the *install_dir*/noapp.properties_platform_ifcresources_ext file using a text editor.

Download from Www.Somanuals.com. All Manuals Search And Download.

2. Under `# POLICY CONFIGURATION - ALL Policies`, locate the MaxThreads property and change the value to the value you determine is best for your level of processing.

   Note: Setting the MaxThreads value higher does not 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. Under `# 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
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 enables smoother multitasking. Interactive use favors a lower number of steps, while batch processing favors a higher number of steps.
AE_ExecuteCycleTime.#	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 AE_ExecuteCycle. This is intended to ensure that a series of unusually slow steps do not tie up a thread completely.
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 that 10,000 business processes in a queue at the same time.
MaxPoolSize.#	Maximum number of threads executed for a specific queue. The MaxThreads value overrides the MaxPoolSize value if the MaxPoolSize value is set higher than the MaxThreads value. **Note:** Setting all the queues' MaxPoolSize parameters to the maximum threads available does not mean you will have faster processing. Depending on your system resources, setting the MaxPoolSize value too high may cause the queues to back up and degrade performance. A good place to start is 4 times the number of your CPUs. Interactive use tends to favor more threads, while batch or document processing tends to favor less threads.

Download from Www.Somanuals.com. All Manuals Search And Download.

Property	Description
CacheThreshold.#	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. 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 B2B Integrator System Troubleshooter. For more information, refer to the topic Reviewing System Information. There is a trade-off between performance and memory consumption. Setting this value too high can lead to the following: • Leave Sterling B2B Integrator with insufficient memory in some circumstances • Cause some business processes to remain in the queue for too long without being examined for rescheduling. 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.
MinPoolSize.#	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. .# indicates the queue number. **Note:** Setting the MinPoolSize parameter too low may cause queues to back up and degrade performance.
JavaPriority.#	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).
EnableDeadlines.#	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: • false – Disables deadline support for a queue. • true – Enables deadline support for a queue.
Rescheduling.#	Enables or disables rescheduling support for a queue. Valid values are: • false – Business processes are not rescheduled in a queue. • true – Business processes are rescheduled in a queue.
ReschedulingInterval.#	Time interval, in milliseconds, to wait before the rescheduler passes through the queue to reschedule business processes.

Download from Www.Somanuals.com. All Manuals Search And Download.

Property	Description
MaxWaitTime.#	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.

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. This will only steal unused threads from that queue.Valid values: • 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. Only unused allocated threads can be stolen. Valid values: • 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B Integrator.

6. Stop Sterling B2B 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 B2B Integrator. The changes are applied to Sterling B2B 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 B2B 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/queueWatch, where *host:port* is the IP address and port number where Sterling B2B 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 B2B Integrator Administrator group.
2. Type your username and password. The Queue Watcher displays the following information:

Download from Www.Somanuals.com. All Manuals Search And Download.

Heading	Description
View Active Threads for All Queues	Displays a list of all active queue threads. When selected, you can review the following information: QueueName – Displays the queue name. Min – Minimum number of threads available for the queue. The threads will be honored even if they are higher than MaxThreads (global maximum queue threads). The minimum number of threads cannot be higher than the maximum number. The fairness calculation does not apply for minimum threads. Used – Number of business processes currently running on a thread. Calc – Fairshare thread calculation for the queue. Fairshare is based on concurrent activities on all queues and is dynamically updated. Pool – Number of threads in a queue's pool. Threads timeout if they are not used. Max – Maximum number of threads used by the queue. *Calc* determines the maximum concurrent threads that is dynamically calculated. Queue Depth – Number of business processes waiting for a thread in the queue. List of Working Threads – List of business processes currently running on a thread.
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 Parms	Displays the parameters set for all of the queues.
View Active Queue Configuration Parms	Displays the current queue configuration.
View list of Workflow IDs that recover would see in the queue	Displays the workflow ID when it is run or moved to another node in the cluster. Valid values are: Executed Moved to another (cluster) node

Download from Www.Somanuals.com. All Manuals Search And Download.

Heading	Description
View Context Cache Entries	Displays the context cache entries. **Note:** If entries show up as invalid they are still correct and do not indicate an error. 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. In Memory Cache Bytes in use - This memory cache holds the workflow contexts with a size lesser than the configured threshold if it is has space. The workflow contexts are serialized on this queue. Disk Cache Bytes in use - This cache holds workflow contexts larger than the defined threshold. The workflow contexts are serialized on this queue.
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_platform_ifcresources_ext 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	Configure the queue parameters to tune performance. The parameters are not persisted and are reset when Sterling B2B 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

Heading	Description
Reset Queue	Resets the queue to default values. The parameters are not persisted and are reset when Sterling B2B 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 B2B Integrator without having to restart the system for it to take affect.

To enable Queue Watcher without restarting Sterling B2B 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 B2B Integrator without having to restart the system for it to take affect.

About this task

To disable Queue Watcher without restarting Sterling B2B Integrator:

Procedure

1. Access the Queue Watcher tool. For more information about accessing Queue Watcher, refer to the topic "Accessing Queue Watcher" on page 177.
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.

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B Integrator.
4. Stop Sterling B2B 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 B2B Integrator. The changes are applied to Sterling B2B 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

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.

About this task

When you install Sterling B2B 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 **Soft Stop**. This stops Sterling B2B Integrator. For more information about the softstop process, refer to "Completing a Soft Stop of Sterling B2B Integrator from the System Troubleshooter" on page 111
5. From the *install_dir*/bin directory, run one of the following commands:
   - (For UNIX or Linux) **setupfiles.sh**
   - (For Windows) **setupfiles.cmd**

Download from Www.Somanuals.com. All Manuals Search And Download.

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

6. From the install_dir/bin directory, start Sterling B2B Integrator by running one of the following commands:
   - (UNIX or Linux) **run.sh**
   - (Windows) **startWindowsService.cmd**

   Sterling B2B Integrator starts, and the updated performance configuration settings are applied.

Allocate JNDI Contexts

About this task

Several components of Sterling B2B 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 B2B Integrator to enhance performance. In Sterling B2B 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 **Soft Stop**. This runs the softstop script, allowing all active business processes to complete on a node before stopping the system. The script is run through a series of interactive steps in the UI. For more information about the softstop process, refer to "Completing a Soft Stop of Sterling B2B Integrator from the System Troubleshooter" on page 111
3. In the *install_dir*/properties directory, open the tuning.properties file using a text editor.
4. Locate the tune.noapp.jdni.contextpoolsize 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 B2B Integrator by running one of the following commands depending on your operating system:
   - (For UNIX or Linux) **run.sh**
   - (For Windows) **startWindowsService.cmd**

   Sterling B2B Integrator starts, and the updated performance configuration settings are applied.

Changing Persistence Levels

Persistence level is the level of detail written to the database even as the business process runs. Sterling B2B 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 B2B Integrator is Full persistence. You may find that you have to change the persistence level to a lower level in order to enhance performance.

Download from Www.Somanuals.com. All Manuals Search And Download.

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 the 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 B2B Integrator.
4. In the *install_dir*/bin directory, run one the following commands:
   - (For UNIX or Linux) **setupfiles.sh**
   - (For Windows) **setupfiles.cmd**
5. Start Sterling B2B Integrator.

Changing Persistence at the Business Process Level

To edit the persistence at the business process level:

1. From the **Administration** menu, select **Business Processes** > **Manager**.

Download from Www.Somanuals.com. All Manuals Search And Download.

2. Enter either the entire business process name for which you want to edit the persistence level or a part of the business process name in the **Search** box 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 B2B 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 B2B 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 B2B 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.
   - Override None No IC - Does not retain any details while the business process is running or in the current process, and will not retain any details after the business process is run.
8. Click **Next**.

Download from Www.Somanuals.com. All Manuals Search And Download.

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**.

Note: When the workflow has the storage set to file system or document tracking is enabled, regardless of the persistence level setting of the workflow, the documents processed by the workflow will be saved in the storage internally. The payload in the file system can be purged later or the document can be tracked by the end user using the **Current Documents** search page. It is invisible to the end user directly. Because the documents are saved to storage, it will have an impact on the performance.

Changing System Logging

Sterling B2B 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 B2B Integrator has document type definitions (DTDs) and schemas in a repository for XML components. When a DTD or schema is required, Sterling B2B Integrator searches the database for the appropriate file. If Sterling B2B Integrator cannot find the appropriate file, you can configure the *install_dir*/properties/noapp.properties_platform_ifcresources_ext file to have Sterling B2B Integrator go to trusted Web sites and download the appropriate DTD or schema.

To configure Sterling B2B 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_ifcresources_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 B2B Integrator will search through these Web sites to find the appropriate DTD or schema.
3. Save the noapp.properties_platform_ifcresources_ext.in file under the same name and in the same location.
4. Stop Sterling B2B Integrator.
5. In the *install_dir*/bin directory, run one the following commands:
   - (For UNIX or Linux) **setupfiles.sh**

Download from Www.Somanuals.com. All Manuals Search And Download.

- (For Windows) **setupfiles.cmd**

6. Start Sterling B2B 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 B2B 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 B2B Integrator. The changes are applied and you can now view the events as they are appended to the noapp.log file during processing.

Enabling the JMX Event Listener

You can enable the JMX (Java Management Extensions) Event Listener to capture Sterling B2B Integrator events and perform the necessary notification and handling based on your custom JMX event handling requirements.

About this task

You can plug in standard JMX tools to monitor Sterling B2B Integrator.

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 B2B Integrator, 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, open the jmx.properties file.
2. Set the **jmx.start** and **jmx.eventListener.start** properties to true.
3. In the *install_dir*/properties directory, open the noapp.properties_platform_ifcresources_ext file.

Download from Www.Somanuals.com. All Manuals Search And Download.

4. Set the **startup.class.ifcresources** property to `com.sterlingcommerce.woodstock.jmx.event.SPEventJMXAgentStartup` by editing the existing line.

   If the **startup.class.ifcresources** property does not exist, create a new line to set the **startup.class.ifcresources4** property to `com.sterlingcommerce.woodstock.jmx.event.SPEventJMXAgentStartup`.

   Remember: To prevent errors, do not set the **startup.class.ifcresources4** more then once.

   Setting the **startup.class.ifcresources** property starts the JMX Agent when Sterling B2B Integrator starts.
5. Restart Sterling B2B Integrator to pick up these changes.
6. Configure a business process to send events.
7. In the Sterling B2B Integrator **JMX Console**, you can view platform events by using a URL of this form: `service:jmx:rmi://HOST:JMXPORT/jndi/rmi://HOST:JMXRMIPORT/events`. 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 as specified in the `jmx.properties` file. 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 B2B 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

Setting	Description
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 B2B 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 B2B 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 B2B 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

The property files in Sterling B2B Integrator help you set up, customize, and diagnose problems pertaining to the installation and operation of Sterling B2B 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 B2B 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 B2B 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 B2B Integrator Property Files* documentation.

CAUTION:
Property files directly affect the operation of Sterling B2B Integrator. Contact IBM 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

Download from Www.Somanuals.com. All Manuals Search And Download.

Managing System Recovery
About this task

Sterling B2B 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 B2B Integrator is installed to a stable state, and then restart Sterling B2B 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 B2B Integrator stops pulling jobs from queues. The system recovers these jobs when the database becomes available.

Sterling B2B 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	Maximum number of business processes that will be resumed or restarted by the Recovery business process each time it runs. Default is 1000. 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 batchSize property).

Download from Www.Somanuals.com. All Manuals Search And Download.

Property	Description
numberOfThreads	Number of threads that will be used simultaneously to restart or resume business processes. Default is 10. Each thread will start or resume a batch of business processes. The size of the batch is specified by the batchSize 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. If you frequently have very large numbers of business processes running at the same time, a larger numberOfThreads value will enhance recovery performance. **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.
batchSize	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. **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.

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 B2B Integrator and restart it to use the new values. For assistance in determining the proper settings for your business requirements, contact IBM Customer Support.

Changing the Soft Stop Time Limit

You can set the time limit used when completing a soft stop of Sterling B2B Integrator from the command line.

About this task

This may be needed to give the business processes that are running adequate time to finish.

You should perform a soft stop by doing the following:

1. Click **Operations** > **System** > **Troubleshooter**.
2. In the System Troubleshooting page, click **Soft Stop**.

For more information about the soft stop process, refer to "Completing a Soft Stop of Sterling B2B Integrator from the Command Line" on page 115.

Download from Www.Somanuals.com. All Manuals Search And Download.

To override the default soft stop 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 lines:
   - noapp.shutdown.timeout_for_bps=*new_value*
   - noapp.shutdown.timeout_for_adapters=*new_value*

   Note: These lines override the **shutdown.timeout_for_bps** and **shutdown.timeout_for_adapters** properties in the noapp.properties_platform_ifcresources_ext.in file.
   In these lines, *new_value* refers to the new time limit, in seconds, for a soft stop of Sterling B2B Integrator. For example, assume that you want to set the soft stop time limit for adapters to 600 seconds (10 minutes). To do so, add the following line to the customer_overrides.properties file:
   noapp.shutdown.timeout_for_adapters=600
4. Save and close the customer_overrides.properties file.
5. Stop Sterling B2B 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 B2B 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:

Download from Www.Somanuals.com. All Manuals Search And Download.

```
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`

Download from Www.Somanuals.com. All Manuals Search And Download.

Property	Description
event_input_queue_capacity	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: `dmivisibility.event_input_queue_capacity=2048` **Notes:** • 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 persistent_batching_interval 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 **`ulimit -n 4096`**.
visibility_coverage	Specifies which AFT Visibility components are enabled. Valid values are: • 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. Example: `dmivisibility.visibility_coverage=All` Visibility coverage property settings can be also combined using commas as delimiters, for example: `dmivisibility.visibility_coverage=CommBase,Authentication,AdminAudit`

4. Save and close the customer_overrides.properties file.
5. Stop Sterling B2B 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 B2B Integrator, and shares many of the resources with the latter.

You should, therefore, tune your Sterling B2B 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 B2B Integrator.

Download from Www.Somanuals.com. All Manuals Search And Download.

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** adjacent to the business process you want to edit.
4. Click **edit** adjacent to the corresponding business process in order to lock that 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, 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 B2B 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 B2B 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:

Download from Www.Somanuals.com. All Manuals Search And Download.

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	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. 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: `<assign to="max_business_processes">500000</assign>`

Download from Www.Somanuals.com. All Manuals Search And Download.

Recommendation	Comments
Batch purge	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.propeties file that can be located in *install_dir*/install/properties (*install_dir*\install\properties for Windows) directory. Modify the batch size by changing the variable as shown in the following example: `query.purge_archiveInfoRowLimit.database_name` **Note:** Replace the *database_name* with the database you are using along with Sterling B2B Integrator, for example, mssql. 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.
Document lifespan	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. You can modify the business process lifespan by logging in as a Sterling B2B Integrator administrator. Navigate to **Operations** > **Archive Manager** > **Configure Archive Settings** and perform the changes.
Database compression	The Schedule_IndexBusinessProcessService and the Schedule_PurgeService runs complex queries which results in signification 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: • 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 B2B 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

The following table lists the threads created in Sterling B2B 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
Business process queues	ReschedulingThread	10
B2B	B2B http Servlet Thread	3
	FIFOTaskListener	10
	QueueThread:queue	11
System	RMI	7
	Timer	7
Others	From various components	30
Total		**287**

Controlling the Threads: Several threads created by Sterling B2B Integrator for various purposes may not be required always and they can be controlled wherever required. This will enhance the Sterling B2B Integrator performance considerably.

Following are the concepts described in this topic:

- ActiveMQ Threads
- Jetty Threads
- JGroup Threads
- JetSpeed Threads
- Adapter Threads
- Business Process Queue Threads
- FIFOTaskListener 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. The Sterling B2B 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

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B 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 B2B 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B Integrator.

To use an external ActiveMQ environment in Windows:

1. Shut down Sterling B2B 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 B2B Integrator.

Sterling B2B 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.

Sterling B2B Integrator

Note: Read install/activemq/conf/readme_cluster.txt file before making any changes.

Jetty Threads

Sterling B2B 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 the noapp.properties_platform_ifcresources_ext 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

JGroup Threads

JGroup is a reliable multicast communication toolkit and is used in Sterling B2B Integrator cluster environment. You cannot control the number of threads created by JGroup.

JetSpeed Threads

Jetspeed is the portal engine used in Sterling B2B 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

Download from Www.Somanuals.com. All Manuals Search And Download.

- SWIFTNet Http Server Adapter

Business Process Queue Threads

Sterling B2B 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 B2B Integrator starts. It is not recommended to control these threads as they are necessary for Sterling B2B Integrator to run smoothly.

The following timer threads are created when Sterling B2B Integrator starts:

- Check Sterling B2B 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

WebSphere MQ Parameters

Your processing volumes and the number of WebSphere MQ queue consumers and producers that you expect to start might require you to change the channel and log parameters in the qm.ini or mqs.ini file.

Channel

Each started thread that reads from or writes to the WebSphere MQ queues requires a channel. If you start 20 JVMs with 5 threads each, you will need at least 100 channels (the default value). You may also have to increase the number of channels if you have workloads that open and close the JMS/MQ connections rapidly.

If you experience messages indicating that the maximum number of channels has been reached, do the following:

- Check to see if there is a connection or channel leak. Run the following command to see how many active channels are used. `YOUR_QM_NAME` should be replaced with your real Queue Manager name.

  ```
  echo "dis chs(*)" | runmqsc YOUR_QM_NAME | grep RUNNING | wc -l
  ```

- You may have to run each workload at peak production loads in your test environment to diagnose channel leaks.
- If you suspect that channels are not getting reclaimed fast enough or if your TCP/IP connection is not reliable, you should set the following parameters. The `KeepAlive` parameter tells the queue manager to check the existence of the client. If the client is not there, the queue manager will reclaim the channel. The `MaxChannels` parameter defaults to 100. In production settings, that parameter could grow to a much higher number like 300 or 500.

  ```
  TCP:
    KeepAlive=YES

  Channels:
    MaxChannels=300
    MaxActiveChannels=100
  ```

Log Files

WebSphere MQ uses log files to maintain message integrity in the event of a queue manager restart or a media failure.

The number of log files depends on your configuration, the size of the messages, the logging type, and the message volumes. You should performance test your application at or above peak production loads to see if the default WebSphere MQ log settings are sufficient. If you are using CIRCULAR logging, the following may be reasonable starting values:

```
Log:
  LogPrimaryFiles=4
  LogSecondaryFiles=1
  LogFilePages=65536
  LogType=CIRCULAR
  LogBufferPages=0
```

If you use LINEAR logging (for example, to survive media failure), you will have to set `LogPrimaryFiles` higher.

Download from Www.Somanuals.com. All Manuals Search And Download.

Placement of WebSphere MQ Log and Data Files

If your system has to be able to process a high message rate, you should consider placing your WebSphere MQ log and data files on a fast SAN, preferably configured with a large NVRAM and RAID-10. A single internal disk should have sufficient capacity to allow up to 150K to 200K messages per second. Files on a RAID-10 LUN should be able to get up to around 1.5M to 2.0M messages per hour. Beyond that message rate, you may want to consider implementing multiple queue managers with separate data and log files.

Performance Statistics

The Performance Statistics Report is the key to managing your Sterling B2B Integrator 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 B2B Integrator.
- Tune Sterling B2B Integrator using the information provided in the topic "Manual Performance Tuning" on page 166.
- Complete manual tuning, if any, to the *install_dir*/properties/noapp.properties file and related 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 B2B 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!**.

Download from Www.Somanuals.com. All Manuals Search And Download.

Reading a Performance Statistics Report

Performance Statistics Report provides you with information that can help you maintain your Sterling B2B 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:

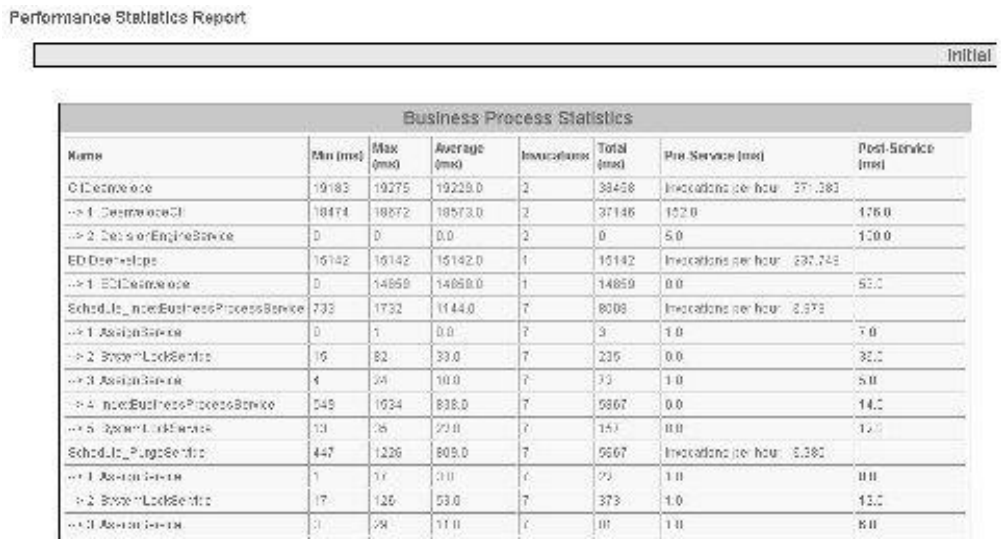
Performance Statistics Report

Initial

Business Process Statistics

Name	Min (ms)	Max (ms)	Average (ms)	Invocations	Total (ms)	Pre-Service (ms)	Post-Service (ms)
EDIEnvelope	19183	19275	19229.0	2	38458	Invocations per hour: 371.383	
--> 1 EnvelopeEDI	18474	18872	18673.0	2	37346	152.0	176.0
--> 2 DecisionEngineService	0	0	0.0	2	0	5.0	120.0
EDIEnvelope	15142	15142	15142.0	1	15142	Invocations per hour: 237.743	
--> 1 EDIEnvelope	0	14859	14859.0	1	14859	0.0	52.0
ScheduleProcessBusinessProcessService	733	1732	1144.0	7	8008	Invocations per hour: 8.373	
--> 1 AssignService	0	1	0.0	7	3	1.0	7.0
--> 2 SystemLockService	16	82	33.0	7	235	0.0	32.0
--> 3 AssignService	4	24	10.0	7	73	1.0	5.0
--> 4 ScheduleProcessBusinessProcessService	549	1934	838.0	7	5867	0.0	14.0
--> 5 SystemLockService	13	26	20.0	7	143	0.0	12.0
SchedulePurgeService	447	1226	809.0	7	5667	Invocations per hour: 8.382	
--> 1 AssignService	1	17	3.0	7	23	1.0	8.0
--> 2 SystemLockService	17	126	53.0	7	373	1.0	13.0
--> 3 AssignService	3	28	11.0	7	81	1.0	8.0

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	

Download from Www.Somanuals.com. All Manuals Search And Download.

Statistic Name	Description
Name	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. For example, in the preceding figure, the business process CIIDeenveloper includes two activities: • --> 1 DeenveloperCII • --> 2 DecisionEngineService
Min (ms)	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. For example, in the figure preceding this table, it can be seen that the CIIDeenvelope 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. 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. For more information about slow systems, refer to the topic A Slow System: Symptoms, Causes, and Resolution. For more information about improving business process execution time, refer to the topic Symptoms and Causes of Poor Business Process Execution Time.
Max (ms)	Maximum time (in milliseconds) it took for a business process or activity took to complete. If there are multiple invocations, this is the highest maximum time taken for all the invocations. For example, in the figure preceding this table, it can be seen that the CIIDeenvelope 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. 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. For more information about slow systems, refer to the topic A Slow System: Symptoms, Causes, and Resolution. For more information about improving business process execution time, refer to the topic Symptoms and Causes of Poor Business Process Execution Time.

Download from Www.Somanuals.com. All Manuals Search And Download.

Statistic Name	Description
Average (ms)	Average processing time taken for a business process or service to be completed. For example, in the figure preceding this table, it can be seen that the CIIDeenvelope 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. 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. For more information about slow systems, refer to the topic A Slow System: Symptoms, Causes, and Resolution. For more information about improving business process execution time, refer to the topic Symptoms and Causes of Poor Business Process Execution Time.
Invocations	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. For example, in the figure preceding this table, it can be seen that the CIIDeenvelope business process had two invocations, and the Deenveloping step (1 DeenveloperCII) had two invocations. 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

Statistic Name	Description
Total (ms)	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. For example, in the figure preceding this table, it can be seen that the CIIDeenvelope 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. 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. **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. For more information about improving business process execution time, refer to the topic Understanding Business Process.
Pre-Service (ms)	Total time taken by the business process engine to complete activities before a business process or service is run. 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 B2B Integrator must complete system service invocations before running the activity. 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. For more information about operating system issues, refer to your vendor's documentation. For more information about slow systems, refer to the topic A Slow System: Symptoms, Causes, and Resolution.

Download from Www.Somanuals.com. All Manuals Search And Download.

Statistic Name	Description
Post-Service (ms)	Total time taken by the business process engine to complete activities after the business process or service is run. 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 B2B Integrator completes some persistence of information to the database. This persistence takes time to complete, and the postservice time includes this persistence time. 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. For more information about operating system issues, refer to your vendor's documentation. For more information about slow systems, refer to the topic A Slow System: Symptoms, Causes, and Resolution.

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 B2B 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:

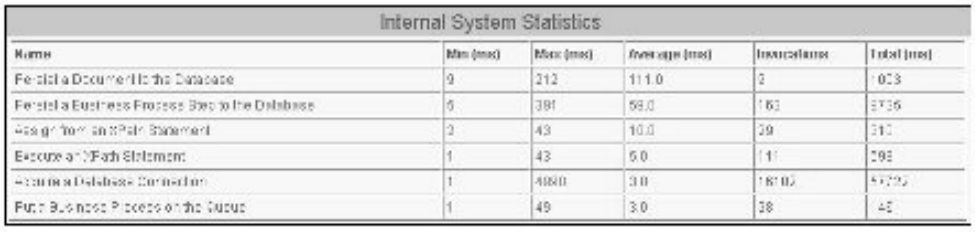

Internal System Statistics

Name	Min (ms)	Max (ms)	Average (ms)	Invocations	Total (ms)
Persist a Document to the Database	9	212	111.0	2	1023
Persist a Business Process Step to the Database	6	391	59.0	163	9735
Assign from an XPath Statement	2	43	10.0	29	310
Execute an XPath Statement	1	43	5.0	111	598
Acquire a Database Connection	1	4880	3.0	16102	57222
Put a Business Process on the Queue	1	49	3.0	38	145

The following table describes the columns in the Internal System Statistics section:

Statistic Name	Description
Internal System Statistics	
Name	Name of the system activity, including the following: • 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

Download from Www.Somanuals.com. All Manuals Search And Download.

Statistic Name	Description
Min (ms)	Minimum time (in milliseconds) the system activity took to complete. If there are multiple invocations, this is the lowest time for all the invocations. 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.
Max (ms)	Maximum time (in milliseconds) the system activity took to complete. If there are multiple invocations, this is the highest time for all the invocations. 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.
Average (ms)	Average processing time taken for the system activity to complete. This is the average time for all invocations. 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.
Invocations	Number of times the system activity was invoked since you last generated the report. For example, in the figure preceding this table, the Persist a Document to the Database activity was invoked nine times.
Total (ms)	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. For example, in the figure preceding this table, the Persist a Document to the Database activity took a total of 1003 ms to complete.

The following table describes each internal system activity, and how to use the statistics to reduce the chances of a performance issue occurring:

Download from Www.Somanuals.com. All Manuals Search And Download.

Statistic Name	Description
Persist a Document to the Database	The length of time and the number of times a document was persisted to the database. For more information about persistence, refer to the topic Changing Persistence Levels. For example, in the figure preceding this table, Persist a Document to the Database had the following statistics: • Min (ms) – 9 • Max (ms) – 212 • Average (ms) – 111.0 • Invocations – 9 • Total (ms) – 1003 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: • 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

Statistic Name	Description
Persist a Business Process Step to the Database	The length of time and the number of times a business process step was persisted to the database. For more information about persistence, refer to the topic Changing Persistence Levels. For example, in the figure preceding this table, Persist a Business Process Step to the Database had the following statistics: • Min (ms) – 5 • Max (ms) – 391 • Average (ms) – 59.0 • Invocations – 163 • Total (ms) – 9735 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: • 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. For more information about slow systems, refer to the topic Changing Persistence Levels. • 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	The length of time and the number of times an assign activity was completed from an XPath statement in a business process. For example, in the figure preceding this table, Assign from an XPath Statement had the following statistics: • Min (ms) – 2 • Max (ms) – 43 • Average (ms) – 10.0 • Invocations – 29 • Total (ms) – 310 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: • 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

Statistic Name	Description
Execute an XPath Statement	The length of time and the number of times an XPath statement ran in a business process. For example, in the figure preceding this table, Assign from an XPath Statement had the following statistics: • Min (ms) – 1 • Max (ms) – 43 • Average (ms) – 5.0 • Invocations – 111 • Total (ms) – 598 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: • 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	The length of time and the number of times a database connection was made. For example, in the figure preceding this table, Acquire a Database Connection had the following statistics: • Min (ms) – 1 • Max (ms) – 4990 • Average (ms) – 3.0 • Invocations – 16102 • Total (ms) – 57722 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: • 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 B2B 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

Statistic Name	Description
Put a Business Process on the Queue	The length of time and the number of times a business process was placed in a queue. For example, in the figure preceding this table, Put a Business Process on the Queue had the following statistics: • Min (ms) – 1 • Max (ms) – 49 • Average (ms) – 3.0 • Invocations – 38 • Total (ms) – 146 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.

Turning On and Turning Off Performance Statistics

You may find that you want to have performance statistics turned on during specific time periods and turned off during other time periods. For example, 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.

Turning On Performance Statistics

You can turn on performance statistics in Sterling B2B Integrator.

Procedure

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**. Performance statistics are turned on, and are ready to be captured in a report.

Turning Off Performance Statistics

You can 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 **Disable Performance Statistics**. Performance statistics are turned off.

Reporting Performance Statistics

Reporting performance statistics includes the following tasks:

- Creating a New Performance Statistics Report
- Deleting a Performance Statistics Report

Download from Www.Somanuals.com. All Manuals Search And Download.

Creating Performance Statistics

After enabling performance statistics, create a Performance Statistics report to view the statistics.

About this task

To create a Performance Statistics report:

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 B2B Integrator.

About this task

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.

Troubleshooting

Tools, tips and guidance are provided to help you troubleshoot problems in Sterling B2B Integrator.

Database Maintenance Check

Sterling B2B 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B 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 **Name** box and click **Go!**.

   Note: The DBStats report can also be found under **List**. However, it cannot be found using the Type "list" under **Search**.
3. Click **source manager**.
4. Select the appropriate report format from the drop-down 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: • 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).

Download from Www.Somanuals.com. All Manuals Search And Download.

Report Section	Description and Fields
Number of BPs by Removal Method	Shows the number of business processes flagged to be archived, purged, or indexed. It contains the following fields: • Archive Flag – Shows the number of business processes in each category: – 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: • Archive Flag – Shows one of the following conditions: – To Be Archived – Business processes flagged to be archived – To Be Purged – Business processes flagged to be purged – Archived, to be Purged – Business processes that have been archived, and are flagged to be purged • Min – Date and time of the first expired business process • Max – Date and time of the last expired business process
Table Row Counts	Shows the number of rows for each table in the active system. Fields are: • 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.
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.

Download from Www.Somanuals.com. All Manuals Search And Download.

Report Section	Description and Fields
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: • 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: • 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 B2B 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. **Note:** The database optimization scripts can do a lot of database updates to rebuild the indexes and analyze the database tables. This could lead to deadlocks if done while Sterling B2B Integrator is running. If you want to run these scripts while Sterling B2B Integrator is running, you should run them during non-peak times. You can also avoid deadlocks by running these scripts while Sterling B2B Integrator is not running, as long as the database is running.
-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.

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B 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 B2B 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 B2B 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B 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 IBM 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_BPRecovery 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 B2B 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

- 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	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. To verify that the index business process is running, review the archive.log file for data that was purged or archived. This log file will state whether each data purge or data archive was successful. If the archive.log file shows no data being archived or purged, your index business process may not be running, causing your database to fill up. If the index business process is running and failing, contact IBM Professional Services for assistance.
Complete archiving activities	If your archival business process is scheduled, but is not running, your database can fill up because data is not being archived and moved off the database tables to another storage location. Review the archival business processes and verify that the schedule is defined and turned on in the business process, and that business processes are archived after a reasonable amount of time. To verify that the archival business process is running, review the archive.log file for data that was purged or archived. This log file will state whether each data purge or data archive was successful. If the archive.log file shows no data being archived or purged, your archival business process may not be running, causing your database to fill up. If the archival business process is running and failing, contact IBM Professional Services for assistance.
Complete purging activities	If your purging business process is scheduled, but is not running, your database can fill up because data is not being removed from the database tables. Review the purging business processes and verify that the schedule is defined and turned on in the business process, and that business processes are purged after a reasonable amount of time. To verify that the purging business process is running, review the archive.log file for data that was purged or archived. This log file will state whether each data purge or data archive was successful. If the archive.log file shows no data being archived or purged, your purging business process may not be running, causing your database to fill up. If the PurgeService business process is running and failing, contact IBM Professional Services for assistance.

Download from Www.Somanuals.com. All Manuals Search And Download.

Approach	Details
Complete Associate Docs to BP activities	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. 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.
Reduce the amount of large or old data stored in the database	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.
Optimizing the Size of Your Database	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: • 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. For more information about compressing your database, refer to your database documentation.
Optimizing the Data Retention Period	You may be retaining information in the database for longer periods of time than necessary. Set the data retention period in the database to lower levels as appropriate. For more information about setting data retention periods in your database, refer to your database documentation.
BPRecovery Business Process is Not Running or is Failing	If your BPRecovery business process is either not running or is failing: 1. Contact IBM Professional Services for assistance with the BPRecovery business process 2. Run the corrected BPRecovery. 3. Restart, resume, or terminate interrupted business processes. You can perform a mass termination of business processes if many are in an interrupted state.

Document Persistence

Sterling B2B 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

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B 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 B2B 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 B2B 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 "View and Restore Performance Configuration Settings" on page 121.

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.

Download from Www.Somanuals.com. All Manuals Search And Download.

2. The System Troubleshooting page is displayed as shown in the following figure. Click **Database Usage**.

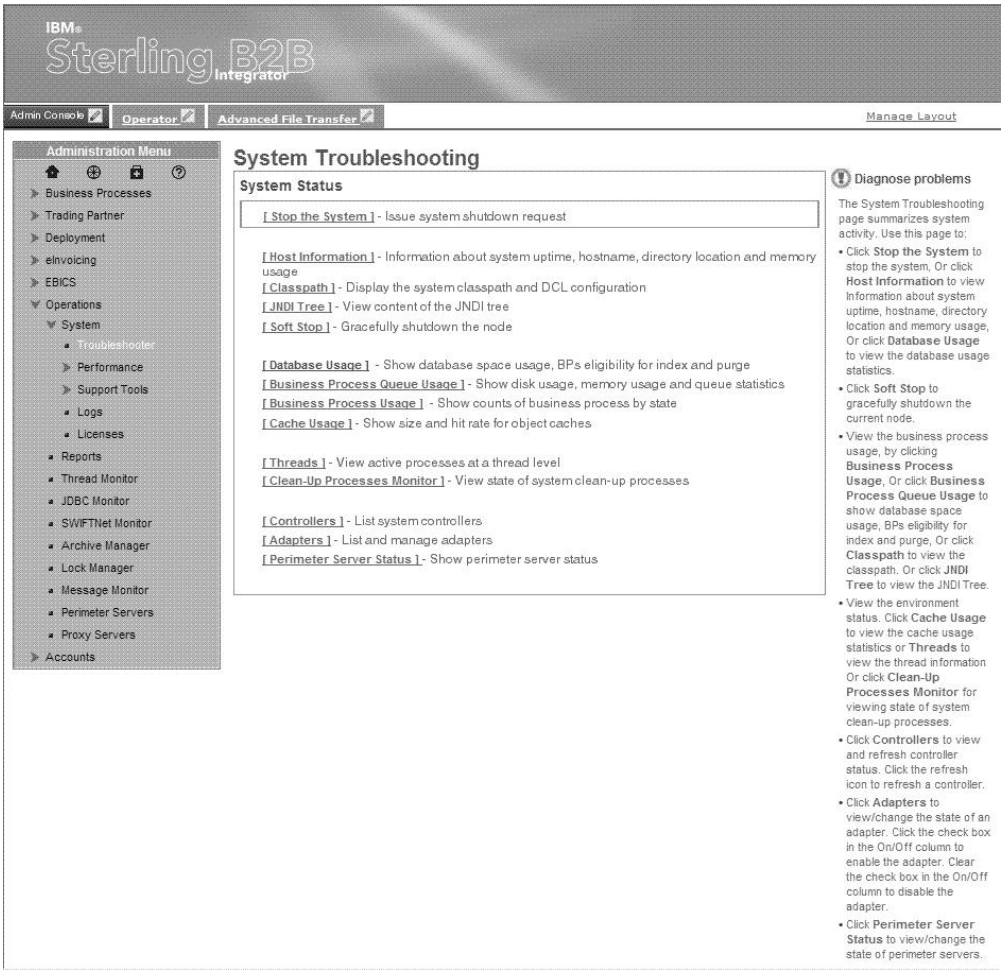

The Database Usage Report (details) is displayed, as shown in the following figure.

Download from Www.Somanuals.com. All Manuals Search And Download.

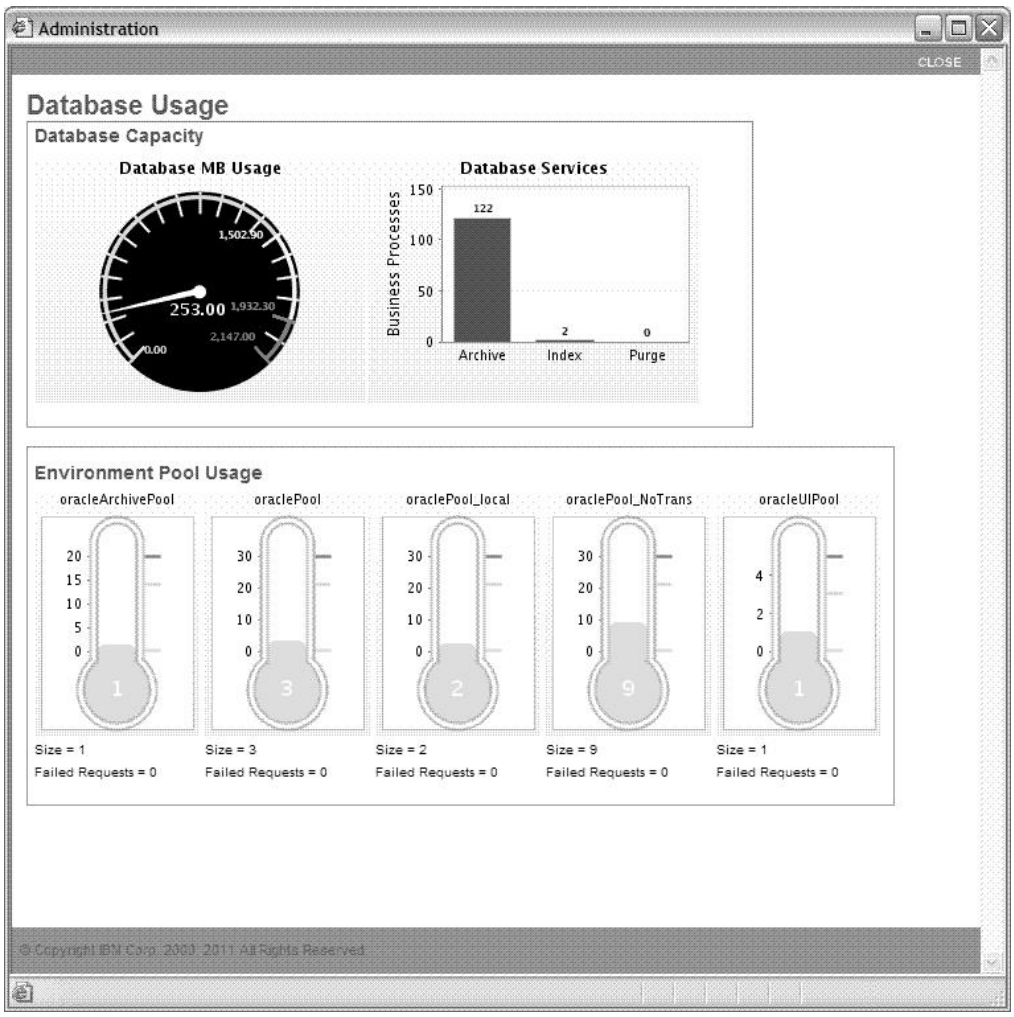

Resolve Database Connection Issues by Increasing the Database Pool Size

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 B2B Integrator, navigate to **Operations** > **System** > **Performance** > **Tuning**.

For more information about increasing the database pool size, refer to "Performance Tuning Utility" on page 120.

Database Down Check

Sterling B2B 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B 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 B2B 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 B2B Integrator to use the new values.

Download from Www.Somanuals.com. All Manuals Search And Download.

For assistance in determining the proper settings to suit your business needs, contact IBM Customer Support.

Tracking JDBC Connections

Sterling B2B 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 IBM 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 B2B 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 B2B Integrator is restarted, you can use the internal **use tracking** parameter.

To control JDBC tracking with the **use tracking** parameter:

1. Stop Sterling B2B 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 B2B 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 or save a JDBC report:

1. From the **Administration** menu, select **Operations** > **JDBC Monitor**.

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B 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 B2B Integrator follows this practice with soft caches. Most soft caches in Sterling B2B 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 B2B Integrator and the number of CPUs, the Sterling B2B 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B 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 B2B 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 B2B 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B Integrator interacts extensively with other systems, sometimes, problems with dropped data occur because of the configuration settings of the other systems.

Download from Www.Somanuals.com. All Manuals Search And Download.

For example, when Sterling B2B 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 B2B 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 B2B 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 more information about allocating physical memory to Sterling B2B Integrator, 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 B2B 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 B2B 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 more information about how to increase the memory available to Sterling B2B Integrator, refer to the topic Viewing or Editing Performance Configuration Settings.

Understanding Business Process

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

There are many symptoms and causes of poor business process execution time.

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**.

Download from Www.Somanuals.com. All Manuals Search And Download.

- 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 B2B 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 other Sterling B2B Integrator 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 other Sterling B2B Integrator business processes or components.
- Improperly written XPath statements, which use resources that can otherwise be used by other Sterling B2B Integrator 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:

Download from Www.Somanuals.com. All Manuals Search And Download.

- 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**. 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 234.

Download from Www.Somanuals.com. All Manuals Search And Download.

- 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 218.
- 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 B2B Integrator stops running. For example, your site experiences a power outage and you must restart Sterling B2B 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 **Operations** > **System** > **Troubleshooting** page. This report shows the number of business processes in the different states. You can click the number next to the state to view detailed information about the process, which may help troubleshoot any problems.
- Review 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 213.

Download from Www.Somanuals.com. All Manuals Search And Download.

- 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 B2B 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 B2B 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 B2B Integrator. Review your operating system or database applications for performance issues.

Download from Www.Somanuals.com. All Manuals Search And Download.

- 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 B2B 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 B2B 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 B2B Integrator. These issues can be averted or resolved by implementing well-designed business processes.

Disabled Service

If a service that Sterling B2B 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 B2B 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

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 “Performance Tuning Utility” on page 120 or Manual Performance Tuning.

Performing a Thread Dump

A thread dump is a snapshot of every thread that is running in Sterling B2B Integrator at the time the thread dump is generated. You can perform a thread dump to determine bottlenecks in Sterling B2B 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:

- Perform a thread dump.
- Download and delete the dumps that are 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 (Version 5.2.4.1 and interim fix 5.2.4.1_2 or later), 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. You can also use the command line to take a thread dump with the shell script.

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 B2B 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

Performing a Thread Dump in UNIX (IBM AIX)

To perform a thread dump in IBM AIX:

1. Access the computer on which Sterling B2B 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 B2B 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 B2B Integrator.
3. Run the **start_si_console.cmd** command from the *install_dir*\bin directory. It will run Sterling B2B Integrator in console mode.

   Running the **start_si_console.cmd** command does not start Sterling B2B Integrator services. It will run Sterling B2B 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 B2B Integrator to start.
4. Perform a thread dump using the following steps:
   - Go to the **start_noapp_console.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 B2B Integrator is installed.

Download from Www.Somanuals.com. All Manuals Search And Download.

2. In the iSeries command line, enter **WRKACTJOB**.
3. Locate the two QP0ZSPWP jobs that are running in Sterling B2B 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)
```


Download from Www.Somanuals.com. All Manuals Search And Download.

```
- locked <0xcc709728> (a java.lang.ref.ReferenceQueue$Lock)
at sun.rmi.transport.DGCClient$EndpointEntry$RenewCleanThread.
run(DGCClient.java:500)
at java.lang.Thread.run(Thread.java:534)
```

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
During Oracle RAC failover, DMI visibility takes time to recover under high volume test when default queue sizes and thread counts are used.	You can increase the number of visibility queues and event input queue capacity in customer_overrides.properties to ensure that adapters are responsive. To increase the default queue size and the thread counts, perform the following tasks: 1. In the customer_overrides.properties file, set the following values: • dmivisibility.number_visibility_threads=6 • dmivisibility.event_input_queue_capacity=3072 2. From the *install_dir*/bin directory, run one of the following commands: • (UNIX or Linux) **setupfiles.sh** • (Windows) **setupfiles.cmd** 3. Restart Sterling B2B Integrator.

Download from Www.Somanuals.com. All Manuals Search And Download.

Situation	Explanation/Resolution

While running Oracle 11*g* in non-RAC or RAC environment, you might encounter index hot spots, which may cause performance degradation.	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.
	• PERF_IDX_1 : WORKFLOW_CONTEXT (BASIC_STATUS, START_TIME, WORKFLOW_ID, WFD_ID)
	• PERF_IDX_2 : WORKFLOW_CONTEXT (START_TIME, WORKFLOW_ID, STEP_ID, ACTIVITYINFO_ID, WFD_ID)
	• PERF_IDX_3 : WORKFLOW_CONTEXT (WFD_ID, START_TIME, WORKFLOW_ID, STEP_ID, ACTIVITYINFO_ID)
	• SCI_IDX_125 : WORKFLOW_CONTEXT (WFD_ID, WFD_VERSION)
	• SCI_IDX_195 : WORKFLOW_CONTEXT (END_TIME, WFD_ID, WFD_VERSION, STEP_ID)
	• SCI_IDX_3 : WORKFLOW_CONTEXT (WORKFLOW_ID)
	• PERF_IDX_4 : WORKFLOW_CONTEXT (WORKFLOW_ID, START_TIME, STEP_ID, ACTIVITYINFO_ID, WFD_ID)
	• SCI_IDX_120 : WORKFLOW_CONTEXT (WORKFLOW_ID, BASIC_STATUS, NEXT_AI_ID, ACTIVITYINFO_ID, ARCHIVE_FLAG)
	• SCI_IDX_109 : WORKFLOW_CONTEXT" (WORKFLOW_ID, BRANCH_ID)PERF_IDX_5 : EVENT (EXPIRES_AT)
	• SCI_IDX_56_2 : TRANS_DATA (WF_ID)
	• SCI_IDX_DMI17 : ACT_SESSION (CON_END_TIME, PROTOCOL, SESSION_ID, ADAPTER_NAME)
	• FG_RTE_MODTS_IDX : FG_ROUTE_EVENT (MODIFYTS)
	• FG_TEA_EVTKEY_IDX : FG_EVENTATTR (EVENT_KEY)
	• FG_RE_DF_IDX : FG_ROUTE_EVENT (DATA_FLOW_ID)
	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.
	The following example explains the procedure to drop an index and create the same in a global partition:
	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);`

Download from Www.Somanuals.com. All Manuals Search And Download.

Situation	Explanation/Resolution
In some cases, perimeter services may fail with the following error. The error will be recorded in the perimeter services log (perimeter.log).	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. The default value of the delayed.event.queue.size property is 30000. You can increase the value based on your requirement.

Performance and Tuning Worksheet

The Performance and Tuning Worksheet can assist you in tuning Sterling B2B 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 B2B Integrator System Requirements* documentation that you used to install Sterling B2B 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.	

Download from Www.Somanuals.com. All Manuals Search And Download.

Request	Description	Your Answer
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.	
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?	Enter the level of persistence (storage to the database) you require for your business process. Choices include: • 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 B2B 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 B2B 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.	

Download from Www.Somanuals.com. All Manuals Search And Download.

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 B2B Integrator. The build date and lower release numbers must also match.	
What are the Sterling B2B Integrator version and patch numbers?	Enter the version and patch number for your copy of Sterling B2B Integrator.	
Who is your relational database vendor?	Enter the name of the relational database vendor Sterling B2B 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 B2B Integrator will process.	
What is the average size of each inbound transaction type?	Enter the average size of your inbound transactions that Sterling B2B Integrator will process. Base your planning on the larger transactions at peak times for better performance.	
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 B2B Integrator will process.	

Download from Www.Somanuals.com. All Manuals Search And Download.

Request	Description	Your Answer
What is the average size of each outbound transaction type?	Enter the average size of each of your outbound transactions that Sterling B2B 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.	

workflowLauncher: Running a Business Process from a Command Line

About this task

Sterling B2B 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 *BPname*	Specifies the name of the business process definition to launch. Mandatory. Example: **workflowLauncher.sh -n testBP** 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

Option	Description
-u *username*	Specifies the user name from which to include the userToken. Optional. Example: **-u** *username* This parameter is used to specify the user token of the business process.
-d *directory*	Specifies the output directory. Optional. Example: **-d** *absolutePath/directory* Used in conjunction with the **-t** 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 -d 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 **-o** parameter.
-o	Used to prevent prefixing the trace file name with the business process ID when generating step XML trace files. Optional.
-v *versionNumber*	Specifies a specific version of the business process definition to launch. Optional. Example: **-v** *business_process_version_number* 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. 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 B2B Integrator Business Process Source Manager.
-f *inputFile*	Specifies a file to use as the primary document when launching the business process. Optional. Example: **-f** *path/filename* 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. An error will be generated if the file is not found.
-e *encoding*	Specifies the encoding of the primary document. Optional. Example: **-e** *encodingName* Used in conjunction with the **-f** parameter. To see what encoding names are supported, see the encodings.properties file in the *install_dir*/properties directory. The full listing is shown in the encodings_large.properties file.

Download from Www.Somanuals.com. All Manuals Search And Download.

Option	Description
-w *count*	Specifies the number of times to sleep. Default is 10. Example: **-w** *count* 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.
-i *time*	Specifies the length of time, in milliseconds, to sleep. Default is 3000 (3 seconds). Example: **-i** *time* Used in conjunction with the -w parameter to define the total timeout value. 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.
-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.

Download from Www.Somanuals.com. All Manuals Search And Download.

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

© Copyright IBM Corp. 2000, 2015

Download from Www.Somanuals.com. All Manuals Search And Download.

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.

Download from Www.Somanuals.com. All Manuals Search And Download.

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 2015. Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. 2015.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and 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.

Download from Www.Somanuals.com. All Manuals Search And Download.

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.

Download from Www.Somanuals.com. All Manuals Search And Download.

Index

A

active threads 101
activities
  monitoring 94
  stopping 95
AE_ExecuteCycle 174
AE_ExecuteCycleTime 174
aftrouting.log 89
alerterlogger.log 90
AllowStealing 176
analyzing log files 84
archive.log 90
ASI 173
  performance tuning 166
  queue tuning 173
  queues 172
ASI queue 173
ASM 46
audit.log 93, 94
auditing 93
Authentication.log 92

B

Backuplogs.log 92
Basic Scheduling policy 172
bizIntel.log 91
BLOB 32
bpdeadline.log 92
business process
  resolving 232
    halted 232
    halting 232
    interrupted 232
    waiting 232
business process execution time
  causes 231
  improving 230
  symptoms
    increasing 230
business process threads 3

C

cache 3
  determining cause of inefficiency 228
  disk 227
  inefficient 228
  memory 227
  soft 227
  symptoms of inefficiency 228
cache contents
  compressing 180
cache usage 104
CacheThreshold 175
capacity requirements 14
  example 15
causes
  slow system 234
cdinterop_cdjava.log 90
cdinterop.log 90
cdsp.log 90
ceuinterop.log 90
checklist
  Sterling File Gateway 9
  Sterling Integrator 5
Clean-Up Processes Monitor 104
cluster environment and JNDI 101
cluster node status 3
common3splogger.log 92
compressing cache contents 180
configuring
  DB resource monitor 225
  DB2 35
  Microsoft SQL Server 51
  Oracle 22
console listener
  enabling 186
Consume business process 96
controllers 99
creating
  performance statistics report 214
crypto.log 90
cspftp.log 90
csphttp.log 90

D

database
  buffer pool statistics 29
  connection issues 222
    failed requests 222
    resolve 224
  db_optimization_tool 217
  down 224
  full database
    causes 219
    issues 218
    resolution 220
    symptoms 219
  instance efficiency 29
  JDBC connections 226
  maintenance 214
  monitoring disk space 14
  optimizing 217
  planning 13
  statistics report 215
  storage 26
    ASM 26
    raw disks 26
    tablespaces 27
DB2
  configuring 35
  connections 39
  event monitors 49
  indexes 51
  logs 47
    LOGFILSIZ 48
    LOGPRIMARY 48
    LOGSECOND 48
    NUM_LOG_SPAN 48
DB2 *(continued)*
  monitoring 35, 48
    snapshots 48
  parameters 38
  performance manager 50
  storage 44
  tablespaces 45
  volatile tables 51
deadlock 50, 53, 61, 72
defrag 61
delete.log 90
deprecated resources
  identifying 109
  monitoring 109
descriptors
  file 191, 193
    file descriptors 166
determining cause
  inefficient cache 228
  slow system 234
directory naming convention 82
disabling
  perimeter server 106
document persistence 221
DTDs
  repositories 185
  trusted domains 185

E

ebXML.log 90
EDI 187
edi.log 90
EDIINT.log 90
embeddedEngine.log 90
EnableDeadlines 175
enabling
  console listener 186
  perimeter server 106
enveloping.properties.in 187
event.log 91

F

Fair Share Scheduling policy 172
file descriptors 191, 193
file system 14
  native 16
ftp.log 91
ftpclient.log 91

G

garbage collection 62
  example 62
  HotSpot 64
  tools 65

© Copyright IBM Corp. 2000, 2015

Download from Www.Somanuals.com. All Manuals Search And Download.

H

heap 70
HotSpot JVM
  garbage collection
    GCViewer 65
    HPjmeter 66
    Java Garbage Collector 66
    visualgc 66
http.log 91
httpclient.log 91

I

improving business process execution time 230
indexes
  DB2 51
  Microsoft SQL Server 60
  Oracle 31
  rebuild 31, 60
inefficient cache
  resolving inefficiency 228
InitialCycles 176
installation path 101
interrupting threads 95

J

JavaPriority 175
JBOD 15
JDBC
  monitor 3
JDBC tracking
  disable 226
  enable 226
jdbc.properties.in 181, 226
jetty.log 91
JMX event listener 186
jmx.log 91
JNDI 101
  contexts 182
JNDI and cluster environment 101
JVM 62
  garbage collection
    example 62
    HotSpot 64
    HotSpot tools 65
    statistics 62
  heap dumps 77
  heap memory 64
  HotSpot 63, 69, 70, 72
    blocked thread 73
    deadlock 72
    tuning 63
  Hprof 77
  IBM 78
    parameters 78
    tuning 78
  profiling 77
  version 62
JVM monitor 3

L

latency 1
lifecycle.log 91
listenerStartup.properties.in 186
load 1
log file
  changing location 88
  naming convention 83
  types 89
logs 3
  analyzing 84
  changing system 185
  DB2 47
  naming conventions 82
  overview 82
  redo 25
  types 89
  viewing contents 83

M

mailbox.log 91
manageability 14
managing
  performance statistics 203
  system logs 82
MaxPoolSize 173, 174
MaxThreads 174
MaxWaitTime 176
memory in use 101
Message monitor 3
messages 3
  monitoring 96
mgmtdash.log 90
Microsoft SQL Server
  autogrow 57
  configuring 51
  dynamic management views 58
  indexes 60
    rebuild 60
  maintenance plan 59
  Management Studio 57
  monitoring 51, 57
  performance monitor 58
  SQL Server Profiler 57
  storage 56
  system monitor 58
  table statistics 60
  tips 61
MinPoolSize 173, 175
modifying log levels 85
monitoring
  activities 94
  AWR report 28
  database usage report 222
  DB2 35, 48
  deprecated resources 109
  disk space 14
  HotSpot JVM heap 70
  HotSpot JVM thread 70
  messages 96
  Microsoft SQL Server 51, 57
  node status 106
  Oracle 22
  threads 94, 95

N

naming conventions 82
noapp.log 92, 186
noapp.properties_platform_ifcresources_ext 173, 176
noapp.properties_platform_ifcresources_ext.in 185
node monitoring status 106

O

ocsp.log 91
OdetteFTP.log 91
oftp.log 91
open file descriptors 166, 191, 193
ops_exe.log 92
ops.log 89
opsSecurity.log 89
opsServer.log 89
Oracle
  BLOB caching 32
  configuring 22
  indexes 31
    rebuild 31
  monitoring 22
  table statistics 31
  undo changes 25

P

performance concepts 1
performance statistics
  creating report 214
  managing 203
  turning on/off 213
Performance statistics report 3
performance tuning 182
  ASI 166
  document body serialization threshold 181
  JNDI contexts 182
  manual 166
  methodology 4
  prerequisites 1
  queues 173
  Sterling e-Invoicing 193
  system logging 185
  worksheet 241
Performance tuning wizard 3
performance. 14
perimeter server
  area on troubleshooting page 106
  disabling 106
  enabling 106
Perimeter.log 91
persistence level 183
pipeline.log 91
planning worksheet
  performance and tuning worksheet 241
platform.log 91
Produce business process 96
PSLogger 93

Download from Www.Somanuals.com. All Manuals Search And Download.

Q

QueueDepth 174
queues
  ASI 172

R

RAID 15
redo
  file size 25
  logs 25
report.log 91
requirements planning worksheet 4
rescheduling 175
ReschedulingInterval 175
resolutions
  cache 228
  slow system 235
ResourceAllocation 176
resourcemonitor.log 91
rnif.log 91
rollback transaction errors 26
run.sh 182

S

sap.log 91
sapxi.log 91
scalability 1, 14
schedule.log 91
ScheduleBackup.log 92
schedulemonitor.log 91
ScheduleStopOps.log 93
scheduling policies 3
scheduling policy
  changing 172
  overview 172
SchedulingPolicyName 172
schemas
  repositories 185
  trusted domains 185
searching messages 97
secureproxy.log 92
security.log 91
segment
  order
    data 194
segment order
  maps 194
server sizing 14
servers 99
servicesctl.log 89
setupfiles.cmd 181, 182, 186, 188
setupfiles.sh 181, 182, 185, 186, 188
sftpclient.log 92
sftpserver.log 92
si_exe.log 92
slow system
  causes 234
  determining cause 234
  resolutions 235
  resolving 229, 234
  symptoms 234
soft stop 110
sql.log 92
start time 99
starting system logs 85
StartupPS.log 93
startWindowsService.cmd 182
state 101
status viewing
  system logs 82
StealThreads 176
stopping
  activities 95
  system logs 85
stopping threads 95
storage
  DB2 44
  Microsoft SQL Server 56
survivability 14
symptoms
  cache 228
  slow system 234
system classpath 101
system information 3
system logs 3
  managing 82
  modifying log level 85
  starting 85
  stopping 85
  turning on and off 85
system recovery 3
System troubleshooting 3
system troubleshooting page 98
system.log 91, 92

T

T-SQL 60
tablespaces
  DB2 45
  DMS 45
  SMS 45
thread 70
  blocked 73
thread dump
  example 70
  IBM iSeries 237
  Linux 237
  performing 71
  reading 238
  tools 72
  UNIX (HP-UX or Sun Solaris) 236
  UNIX (IBM AIX) 237
  Windows 237
Thread monitor 3
Thread Monitor 94
threads 3
  interrupting 95
  monitoring 94, 95
  stopping 95
  viewing 95
throughput 1
tracking.log 90
trusted domains
  DTDs 185
  schemas 185
tuning 173
  HotSpot JVM 63, 66
  IBM JVM 78
tuning.properties 177, 182
tx.log 92
txtrace.log 92

U

ui.log 91
uptime 99

V

viewing
  JDBC report 226
  log files 83
  server status 82
  service status 82
  threads 95
visibility.log 92

W

wait events 30
webdav.log 92
WebSphereMQSuite.log 92
webx.log 92
wf.log 90
wfexception.log 90
wfstatistics.log 90
workflowLauncher 244
workflows.properties file 181
worksheet 241
wsrm.log 92
wssec.log 92

Download from Www.Somanuals.com. All Manuals Search And Download.

Download from Www.Somanuals.com. All Manuals Search And Download.

Download from Www.Somanuals.com. All Manuals Search And Download.

Product Number:

Printed in USA

Download from Www.Somanuals.com. All Manuals Search And Download.

Free Manuals Download Website

http://myh66.com

http://usermanuals.us

http://www.somanuals.com

http://www.4manuals.cc

http://www.manual-lib.com

http://www.404manual.com

http://www.luxmanual.com

http://aubethermostatmanual.com

Golf course search by state

http://golfingnear.com

Email search by domain

http://emailbydomain.com

Auto manuals search

http://auto.somanuals.com

TV manuals search

http://tv.somanuals.com