

IBM WebSphere Application Server Network Deployment
for IBM i, Version 8.0

Tuning various types of applications

IBM

Note

Before using this information, be sure to read the general information under “Notices” on page 119.

Compilation date: July 16, 2011

© Copyright IBM Corporation 2011.

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

Contents

How to send your comments	vii
Changes to serve you more quickly	ix
Chapter 1. Tuning Application profiling	1
Application profiling performance considerations	1
Chapter 2. Tuning Client applications	5
Adding tracing and logging for stand-alone clients	5
Chapter 3. Tuning Data access resources	7
Tuning data	7
Tuning connection pools	7
Throttling inbound message flow for JCA 1.5 message-driven beans.	10
Database performance tuning	11
DB2 Universal Database performance tips	12
Directory conventions	12
Chapter 4. Tuning EJB applications	15
EJB 2.1 container tuning	15
EJB container tuning	15
Tuning Enterprise JavaBeans applications	20
Tuning EJB cache with trace service	20
EJB method Invocation Queuing	23
Tuning applications that use the Java Persistence API	23
Configuring OpenJPA caching to improve performance.	23
Configuring the WSJPA Object Cache to improve performance.	26
Chapter 5. Tuning Messaging resources	29
Tuning messaging	29
Tuning messaging performance with service integration technologies	29
Configuring MDB throttling for the default messaging provider	32
Chapter 6. Tuning messaging destinations for the WebSphere MQ or V5 default messaging providers	35
Performance for WebSphere MQ queues.	35
Performance for Version 5 default messaging queues	36
Chapter 7. Throttling inbound message flow for JCA 1.5 message-driven beans	37
Chapter 8. Monitoring server session pools for listener ports.	39
Chapter 9. Tuning Object Request Broker (ORB)	41
Tuning Object Request Brokers	41
Object Request Broker tuning guidelines	41
Chapter 10. Tuning Service integration.	45
Tuning messaging engines	45
Setting tuning properties of a messaging engine	45
Controlling the memory buffers used by a messaging engine	46
Tuning the JDBC data source of a messaging engine	47
Setting tuning properties by editing the sib.properties file	47
Tuning messaging performance with service integration technologies	49

Configuring MDB throttling for the default messaging provider	51
Chapter 11. Tuning messaging engine data stores	55
Tuning the JDBC data source of a messaging engine	55
Controlling the memory buffers used by a messaging engine	55
Increasing the number of data store tables to relieve concurrency bottleneck	57
Increasing the number of item tables for a messaging engine when tables are not automatically created	57
Increasing the number of item tables for a messaging engine when tables are automatically created	58
Tuning one-phase commit optimization	58
Tuning the detection of database connection loss.	59
Chapter 12. Setting tuning properties for a mediation	63
Chapter 13. Enabling CMP entity beans and messaging engine data stores to share database connections	65
Chapter 14. Tuning security	67
Tuning, hardening, and maintaining security configurations	67
Tuning security configurations	67
Hardening security configurations	73
Enablement and migration considerations of Security hardening features	74
Securing passwords in files	76
Backing up security configuration files	88
Chapter 15. Tuning Session Initiation Protocol (SIP) applications	91
Tuning your environment to run SIP servlets	91
Chapter 16. Tuning Transactions	97
Chapter 17. Tuning web applications	99
Tuning URL cache	99
Tuning URL invocation cache	99
Tuning sessions	100
Session management tuning	100
Tuning parameter settings	100
Tuning parameter custom settings	101
Best practices for using HTTP sessions	102
Chapter 18. Tuning web services	107
Tuning Web Services Security	107
Tuning Web Services Security for Version 8.0 applications	107
Tuning Web Services Security for Version 5.x applications	109
Tuning web services reliable messaging applications	110
Tuning bus-enabled web services	110
Including SOAP header schemas in the SDO repository	111
Chapter 19. Tuning Work area	115
Work area service performance considerations	115
Work area service performance considerations	115
Appendix. Directory conventions	117
Notices	119
Trademarks and service marks	121

Index 123

How to send your comments

Your feedback is important in helping to provide the most accurate and highest quality information.

- To send comments on articles in the WebSphere Application Server Information Center
 1. Display the article in your Web browser and scroll to the end of the article.
 2. Click on the **Feedback** link at the bottom of the article, and a separate window containing an e-mail form appears.
 3. Fill out the e-mail form as instructed, and click on **Submit feedback** .
- To send comments on PDF books, you can e-mail your comments to: **wasdoc@us.ibm.com** or fax them to 919-254-5250.

Be sure to include the document name and number, the WebSphere Application Server version you are using, and, if applicable, the specific page, table, or figure number on which you are commenting.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Changes to serve you more quickly

Print sections directly from the information center navigation

PDF books are provided as a convenience format for easy printing, reading, and offline use. The information center is the official delivery format for IBM WebSphere Application Server documentation. If you use the PDF books primarily for convenient printing, it is now easier to print various parts of the information center as needed, quickly and directly from the information center navigation tree.

To print a section of the information center navigation:

1. Hover your cursor over an entry in the information center navigation until the **Open Quick Menu** icon is displayed beside the entry.
2. Right-click the icon to display a menu for printing or searching your selected section of the navigation tree.
3. If you select **Print this topic and subtopics** from the menu, the selected section is launched in a separate browser window as one HTML file. The HTML file includes each of the topics in the section, with a table of contents at the top.
4. Print the HTML file.

For performance reasons, the number of topics you can print at one time is limited. You are notified if your selection contains too many topics. If the current limit is too restrictive, use the feedback link to suggest a preferable limit. The feedback link is available at the end of most information center pages.

Under construction!

The Information Development Team for IBM WebSphere Application Server is changing its PDF book delivery strategy to respond better to user needs. The intention is to deliver the content to you in PDF format more frequently. During a temporary transition phase, you might experience broken links. During the transition phase, expect the following link behavior:

- Links to Web addresses beginning with `http://` work
- Links that refer to specific page numbers within the same PDF book work
- The remaining links will *not* work. You receive an error message when you click them

Thanks for your patience, in the short term, to facilitate the transition to more frequent PDF book updates.

Chapter 1. Tuning Application profiling

This page provides a starting point for finding information about application profiling, a WebSphere extension for defining strategies to dynamically control concurrency, prefetch, and read-ahead.

Application profiling and access intent provide a flexible method to fine-tune application performance for enterprise beans without impacting source code. Different enterprise beans, and even different methods in one enterprise bean, can have their own intent to access resources. Profiling the components based on their access intent increases performance in the application server run time.

Application profiling performance considerations

Application profiling enables assembly configuration techniques that improve your application run time, performance and scalability. You can configure tasks that identify incoming requests, identify access intents determining concurrency and other data access characteristics, and profiles that map the tasks to the access intents.

The capability to configure the application server can improve performance, efficiency and scalability, while reducing development and maintenance costs. The application profiling service has no tuning parameters, other than a checkbox for disabling the service if the service is not necessary. However, the overhead for the application profile service is small and should not be disabled, or unpredictable results can occur.

Access intents enable you to specify data access characteristics. The WebSphere runtime environment uses these hints to optimize the access to the data, by setting the appropriate isolation level and concurrency. Various access intent hints can be grouped together in an access intent policy.

In the product, it is recommended that you configure bean level access intent for loading a given bean. Application profiling enables you to configure multiple access intent policies on the entity bean, if desired. Some callers can load a bean with the intent to read data, while others can load the bean for update. The capability to configure the application server can improve performance, efficiency, and scalability, while reducing development and maintenance costs.

Access intents enable the EJB container to be configured providing optimal performance based on the specific type of enterprise bean used. Various access intent hints can be specified declaratively at deployment time to indicate to WebSphere resources, such as the container and persistence manager, to provide the appropriate access intent services for every EJB request.

The application profiling service improves overall entity bean performance and throughput by fine tuning the run time behavior. The application profiling service enables EJB optimizations to be customized for multiple user access patterns without resorting to "worst case" choices, such as pessimistic update on a bean accessed with the `findByPrimaryKey` method, regardless of whether the client needs it for read or for an update.

Application profiling provides the capability to define the following hierarchy: **Container-Managed Tasks > Application Profiles > Access Intent Policies > Access Intent Overrides**. Container-managed tasks identify units of work (UOW) and are associated with a method or a set of methods. When a method associated with the task is invoked, the task name is propagated with the request. For example, a UOW refers to a unique path within the application that can correspond to a transaction or `ActivitySession`. The name of the task is assigned declaratively to a Java EE client or servlet, or to the method of an enterprise bean. The task name identifies the starting point of a call graph or subgraph; the task name flows from the starting point of the graph downstream on all subsequent I/O requests, identifying each subsequent invocation along the graph as belonging to the configured task. As a best practice, wherever a UOW starts, for example, a transaction or an `ActivitySession`, assign a task to that starting point.

The application profile service associates the propagated tasks with access intent policies. When a bean is loaded and data is retrieved, the characteristics used for the retrieval of the data are dictated by the application profile. The application profile configures the access intent policy and the overrides that should be used to access data for a specific task.

Access intent policies determine how beans are loaded for specific tasks and how data is accessed during the transaction. The access intent policy is a named group of access intent hints. The hints can be used, depending on the characteristics of the database and resource manager. Various access intent hints applied to the data access operation govern data integrity. The general rule is, the more data integrity, the more overhead. More overhead causes lower throughput and the opportunity for simultaneous data access from multiple clients.

If specified, access intent overrides provide further configuration for the access intent policy.

Best practices

Application profiling is effective in a variety of different scenarios including:

- **The same bean is loaded with different data access patterns**

The same bean or set of beans can be reused across applications, but each of those applications has differing requirements for the bean or for beans within the invocation graph. One application can require that beans be loaded for update, while another application requires beans be loaded for read only. Application profiling enables deploy time configuration for beans to distinguish between EJB loading requirements.

- **Different clients have different data access requirements**

The same bean or set of beans can be used for different types of client requests. When those clients have different requirements for the bean, or for beans within the invocation graph, application profiling can be used to tailor the bean loading characteristics to the requirements of the client. One client can require beans be loaded for update, while another client requires beans be loaded for read only. Application profiling enables deploy time configuration for beans to distinguish between EJB loading requirements.

Monitoring tools

You can use the Tivoli Performance Viewer, database and logs as monitoring tools.

You can use the Tivoli Performance Viewer to monitor various metrics associated with beans in an application profiling configuration. The following sections describe at a high level the Tivoli Performance Viewer metrics that reflect changes when access intents and application profiling are used:

- **Collection scope**

The enterprise beans group contains EJB life cycle information, either a cumulative value for a group of beans, or for specific beans. You can monitor this information to determine the difference between using the `ActivitySession` scope versus the transaction scope. For the transaction scope, depending on how the container transactions are defined, `activates` and `passivates` can be associated with method invocations. The application could use the `ActivitySession` scope to reduce the frequency of `activates` and `passivates`. For more information, see the topic, `Using the ActivitySession service`.

- **Collection increment**

The enterprise beans group contains EJB life cycle information, either a cumulative value for a group of beans, or for specific beans. You can monitor `Num Activates` to watch the number of enterprise beans activated for a particular `findByPrimaryKey` operation. For example, if the collection increment is set to 10, rather than the default 25, the `Num Activates` value shows 25 for the initial `findByPrimaryKey`, before any result set iterator runs. If the number of `activates` rarely exceeds the collection increment, consider reducing the collection increment setting.

- **Resource manager prefetch increment**

The resource manager prefetch increment is a hint acted upon by the database engine to depend upon the database. The Tivoli Performance Viewer does not have a metric available to show the effect of the resource manager prefetch increment setting.

- **Read ahead hint**

The enterprise beans group contains EJB life cycle information, either a cumulative value for a group of beans, or for specific beans. You can monitor *Num Activates* to watch the number of enterprise beans activated for a particular request. If a read ahead association is not in use, the *Num Activates* value shows a lower initial number. If a read ahead association is in use, the *Num Activates* value represents the number of activates for the entire call graph.

Database tools are helpful in monitoring the different bean loading characteristics that introduce contention and concurrency issues. These issues can be solved by application profiling, or can be made worse by the misapplication of access intent policies.

Database tools are useful for monitoring locking and contention characteristics, such as locks, deadlocks and connections open. For example, for locks the DB2 Snapshot Monitor can show statistics for lock waits, lock time-outs and lock escalations. If excessive lock waits and time-outs are occurring, application profiling can define specific client tasks that require a more string level of locking, and other client tasks that do not require locking. Or, a different access intent policy with less restrictive locking could be applied. After applying this configuration change, the snapshot monitor shows less locking behavior. Refer to information about the database you are using on how to monitor for locking and contention.

The **application server logs** can be monitored for information about rollbacks, deadlocks, and other data access or transaction characteristics that can degrade performance or cause the application to fail.

Chapter 2. Tuning Client applications

This page provides a starting point for finding information about application clients and client applications. Application clients provide a framework on which application code runs, so that your client applications can access information on the application server.

For example, an insurance company can use application clients to help offload work on the server and to perform specific tasks. Suppose an insurance agent wants to access and compile daily reports. The reports are based on insurance rates that are located on the server. The agent can use application clients to access the application server where the insurance rates are located.

Adding tracing and logging for stand-alone clients

You can add tracing and logging to help analyze performance and diagnose problems with stand-alone clients.

About this task

This information applies to the following WebSphere® Application Server stand-alone clients:

- Thin Client for JMS with WebSphere Application Server
- Thin Client for EJB with WebSphere Application Server
- Thin Client for JAX-WS with WebSphere Application Server
- Thin Client for JAX-RPC with WebSphere Application Server

Procedure

- To enable trace, use either a long form or short form system property.

Note: Trace settings are determined from the system property values the first time that a WebSphere Application Server client is called. The trace settings are then fixed. Therefore, any subsequent changes to the system property settings do not change the trace settings that the WebSphere Application Server client uses.

For further information, see the information on the trace user interface for stand-alone clients.

- To enable First Failure Data Capture (FFDC), use either a long or short form system property.

Note: FFDC settings are determined from the system property values the first time that a WebSphere Application Server client performs an FFDC. The FFDC settings are then fixed. Therefore, any subsequent changes to the system property settings do not change the FFDC settings that the WebSphere Application Server client uses.

For further information, see the information on the First Failure Data Capture user interface for stand-alone clients.

Chapter 3. Tuning Data access resources

This page provides a starting point for finding information about data access. Various enterprise information systems (EIS) use different methods for storing data. These backend data stores might be relational databases, procedural transaction programs, or object-oriented databases.

The flexible IBM® WebSphere Application Server provides several options for accessing an information system backend data store:

- Programming directly to the database through the JDBC 4.0 API, JDBC 3.0 API, or JDBC 2.0 optional package API.
- Programming to the procedural backend transaction through various J2EE Connector Architecture (JCA) 1.0 or 1.5 compliant connectors.
- Programming in the bean-managed persistence (BMP) bean or servlets indirectly accessing the backend store through either the JDBC API or JCA-compliant connectors.
- Using container-managed persistence (CMP) beans.
- Using the IBM data access beans, which also use the JDBC API, but give you a rich set of features and function that hide much of the complexity associated with accessing relational databases.

Service Data Objects (SDO) simplify the programmer experience with a universal abstraction for messages and data, whether the programmer thinks of data in terms of XML documents or Java objects. For programmers, SDOs eliminate the complexity of the underlying data access technology, such as, JDBC, RMI/IIOP, JAX-RPC and JMS, and message transport technology, such as, (java.io.Serializable, DOM Objects, SOAP and JMS).

Tuning data

Tuning connection pools

Using connection pools helps to both alleviate connection management overhead and decrease development tasks for data access. Each time an application attempts to access a backend store (such as a database), it requires resources to create, maintain, and release a connection to that datastore. To mitigate the strain this process can place on overall application resources, the application server enables administrators to establish a pool of backend connections that applications can share on an application server. Connection pooling spreads the connection overhead across several user requests, thereby conserving application resources for future requests.

About this task

Connection pooling can improve the response time of any application that requires connections, especially Web-based applications. When you make a request over the web to a resource, the resource accesses a data source. Because you connect and disconnect frequently with applications on the Internet, the application requests for data access can surge to considerable volume. Consequently, the total overhead for a datastore can quickly become high for Web-based applications, causing performance to deteriorate. When connection pooling capabilities are used, however, web applications can realize performance improvements of up to 20 times the normal results.

Procedure

- **Prevent a connection deadlock.** Deadlock can occur if the application requires more than one concurrent connection per thread, and the database connection pool is not large enough for the number of threads. Suppose each of the application threads requires two concurrent database connections and the number of threads is equal to the maximum connection pool size. Deadlock can occur when both of the following are true:
 - Each thread has its first database connection, and all are in use.

- Each thread is waiting for a second database connection, and none would become available since all threads are blocked.

To prevent the deadlock in this case, increase the maximum connections value for the database connection pool by at least one. This ensures that at least one of the waiting threads obtains a second database connection and avoids a deadlock scenario.

For general prevention of connection deadlock, code your applications to use only one connection per thread. If you code the application to require C concurrent database connections per thread, the connection pool must support at least the following number of connections, where T is the maximum number of threads:

$$T * (C - 1) + 1$$

The connection pool settings are directly related to the number of connections that the database server is configured to support. If you increase the maximum number of connections in the pool and the corresponding settings in the database are not increased accordingly, the application might fail. The resulting SQL exception errors are displayed in the following locations:

- the `stderr.log` file

One of the most common causes of connection deadlock is the use of the same connection pool by both servlets and by Enterprise JavaBeans (EJBs), and where the servlet directly or indirectly invokes the bean. For example, a servlet that obtains a JMS connection from the connection pool, sends a message to a Message Driven Bean (MDB) and waits for a reply. The MDB is configured to use the same connection pool as the servlet, therefore, another connection from the pool is required for the MDB to send a reply to the servlet. Servlets and enterprise beans do not share the same connection pool. This is a classic case of concurrent (C) threads, where $C=2$ and T is the maximum size of the servlet and EJB thread pools.

- **Disable connection pooling.**

- For relational resource adapters (RRAs), add the `disableWASConnectionPooling` custom property for your data sources.

1. Click **JDBC > Data sources**.
2. Click on the name of the data source that you want to configure.
3. Click **Custom properties** under the **Additional Properties** heading.
4. Click **New**.
5. Complete the required fields with the following information:
 - **Name:** `disableWASConnectionPooling`
 - **Value:** `true`

- For other resource adapters, consult with the binding specifications for that resource adapter to configure your applications to disable connection pooling.

1. Programmatically disable connection pooling through the resource adapter.
2. The application server leverages the following code to detect the `javax.resource.NotSupportedException` exception and disable connection pooling:

```
_managedFactory.matchManagedConnections(s,subject,cri); // 169059 174269    }
catch(javax.resource.NotSupportedException e){
```

- **Enable deferred enlistment.** In the application server environment, deferred enlistment refers to the technique in which the application server waits until the connection is used before the connection is enlisted in the application's unit of work (UOW) scope.

Consider the following illustration of deferred enlistment:

- An application component that uses deferred enlistment calls the `getConnection` method from within a global transaction.
- The application component does not immediately use the connection.
- When the application issues the call for initial use of the connection, the transaction manager intercepts the call.

- The transaction manager enlists the XA resource for the connection and calls the `XAResource.start` method.
- The connection manager associated with the XA resource sends the call to the database.

Given the same scenario, but the application component does not use deferred enlistment, the component container immediately enlists the connection in the transaction. Thus the application server incurs, for no purpose, an additional load of all of the overhead associated with that transaction. For XA connections, this overhead includes the two phase commit (2PC) protocol to the resource manager.

Deferred enlistment offers better performance in the case where a connection is obtained, but not used, within the UOW scope. The technique saves the cost of transaction participation until the UOW in which participation must occur.

Check with your resource adapter provider if you need to know if the resource adapter provides this functionality. The application server relational resource adapter automatically supports deferred enlistment.

Incorporating deferred enlistment in your code:

The Java Platform, Enterprise Edition (Java EE) Connector Architecture (JCA) Version 1.5 and later specification calls the deferred enlistment technique lazy transaction enlistment optimization. This support comes through a marker interface (`LazyEnlistableManagedConnection`) and a method on the connection manager (`LazyEnlistableConnectionManager()`):

```
package javax.resource.spi; import javax.resource.ResourceException; import
  javax.transaction.xa.Xid; interface LazyEnlistableConnectionManager { // application server void
  lazyEnlist(ManagedConnection) throws ResourceException; } interface LazyEnlistableManagedConnection
  { // resource adapter }
```

- **Control connection pool sharing.**

You can use the `defaultConnectionTypeOverride`, or `globalConnectionTypeOverride` connection pool custom property for a particular connection factory or data source to control connection sharing.

The `defaultConnectionTypeOverride` property changes the default sharing value for a connection pool. This property enables you to control connection sharing for direct look-ups. If resource references are configured for this data source or connection factory the resource reference's configurations take precedence over the `defaultConnectionTypeOverride` property settings. For example, if an application is doing direct look-ups and unshared connections are needed, set the `defaultConnectionTypeOverride` property to `unshared`.

The value specified for the `globalConnectionTypeOverride` custom property takes precedence over all of the other connection sharing settings. For example if you set this property to `unshared`, all connection requests are unshared for both direct look-ups and resource reference lookups. This property provides you with a quick way to test the consequences of moving all connections for a particular data source or connection factory to unshared or shared without changing any resource reference setting.

If you specify values for both the `defaultConnectionTypeOverride` and the `globalConnectionTypeOverride` properties, only the values specified for the `globalConnectionTypeOverride` property are used to determine connection sharing type.

The following is an example of how to set these properties for a particular data source:

Create a new custom property on the connection pool belonging to the data source in the administrative console: **JDBC providers > DB2 Universal JDBC Driver Provider > Data sources > DB2 Universal JDBC Driver DataSource > Connection pools > Custom Properties > New**. Then specify `defaultConnectionTypeOverride` or `globalConnectionTypeOverride` in the **Name** field, and `shared` or `unshared` in the **Value** field.

Connection pool custom properties

You can use the custom properties page to define the following connection pool custom properties:

- “`defaultConnectionTypeOverride`” on page 10
- “`globalConnectionTypeOverride`” on page 10

defaultConnectionTypeOverride

You can use the `defaultConnectionTypeOverride` connection pool custom property for a particular connection factory or data source to control connection sharing.

The `defaultConnectionTypeOverride` property changes the default sharing value for a connection pool. This property enables you to control connection sharing for direct look-ups. If resource references are configured for a data source or connection factory they take precedence over this property and the resource reference settings are used. For example, if an application is doing direct look-ups, and you do not want unshared connections, set this property to `unshared`.

Data Type	String
Value	unshared, shared

globalConnectionTypeOverride

You can use the `globalConnectionTypeOverride` connection pool custom property to globally control connection sharing for a particular connection factory or data source.

The value specified for the `globalConnectionTypeOverride` custom property takes precedence over all of the other connection sharing settings. For example, if you set this property to `unshared`, all connection requests are unshared for both direct look-ups and resource reference lookups.

This property provides you with a quick way to test the consequences of moving all connections for a particular data source or connection factory to unshared or shared without changing the resource reference settings.

If you specify values for both the `defaultConnectionTypeOverride` and the `globalConnectionTypeOverride` properties, only the value that is specified for the `globalConnectionTypeOverride` property is used to determine connection sharing type.

Data Type	String
Value	unshared, shared

For more information on how these properties are used, see the topic, [Tuning connection pools](#).

Throttling inbound message flow for JCA 1.5 message-driven beans

This topic describes how to throttle message delivery for message-driven beans (MDB) which are deployed as message endpoints for Java Platform, Enterprise Edition (Java EE) Connector Architecture (JCA) Version 1.5 inbound resource adapters.

Before you begin

The throttling of messages as described in this topic does not apply to the two JCA 1.5-compliant messaging providers that are supplied with WebSphere Application Server:

- The default messaging provider.
- The WebSphere MQ messaging provider.

For the default messaging provider, you configure message throttling as described in the related tasks. For the WebSphere MQ messaging provider, you configure message throttling by setting the **maximum server sessions** property on the WebSphere MQ messaging provider activation specifications panel, or the **maxPoolSize** property when using the `createWMQActivationSpec` or `modifyWMQActivationSpec` `wsadmin` commands.

If you have a third-party JCA 1.5-compliant JMS messaging provider, check with your supplier to see whether the method of message throttling described in this topic is appropriate for their messaging provider.

About this task

For installations that use resource adapters that implement the Java EE Connector Architecture (JCA) Version 1.5 message delivery support, the WebSphere Application Server provides message throttling support to control the delivery of messages to endpoint message-driven beans (MDB). You can use this support to avoid overloading the server with a flood of inbound messages.

Message delivery is throttled on a message-driven bean basis by limiting the maximum number of endpoint instances that can be created by the adapter that the MDB is bound to. When the adapter attempts to create an endpoint instance, a proxy for the MDB instance is created and returned as allowed by the JCA 1.5 architecture. There is a one-to-one correspondence between proxies and MDB instances, and like the MDB instances, the proxies are pooled based on the minimum and maximum pool size values associated with the message-driven bean. Throttling is performed through the management of the proxy pool.

At the time the adapter attempts to create an endpoint, if the number of endpoint proxies currently created is equal to the maximum size of the pool, adapter *createEndPoint* processing returns an `UnavailableException`. When this displays, the adapter does not issue any more `createEndPoint()` requests until it has released at least one endpoint back to the server for reuse. Thus, installations can control the throttling of message delivery to a JCA 1.5 MDB based on the setting of the maximum size of the pool associated with each JCA 1.5 message-driven bean.

You can specify the pool size by using the `com.ibm.websphere.ejbcontainer.poolsize` JVM system property to define the minimum and maximum pool size of stateless, message-driven, and entity beans. For a message-driven bean that supports JCA 1.5, the maximum pool size value specified limits how many message endpoint instances can be created for that message-driven bean. For example, if the installation sets the maximum size of a JCA 1.5 MDB pool to 5, then at most 5 messages can be concurrently delivered to 5 instances of the message-driven bean. This property can be specified using the `wsadmin` scripting tool or by specifying it under the administrative console as an environmental variable.

Procedure

1. Open the administrative console.
2. Select **Servers > Server Types > WebSphere application server > *server_name***.
3. Under **Server Infrastructure**, expand **Java and Process Management > Process Definition**.
4. Under Additional Properties, select **Java Virtual Machine**.
5. Under Additional Properties, select **Custom Properties**.
6. Select **New**. A panel with three **General Properties** fields displays. This is where you set the property.
7. In the **Name** field, enter `com.ibm.websphere.ejbcontainer.poolsize`.
8. To complete the **Value** field, see the EJB container system properties for values.
9. After defining the value of the property, click **OK**. You are now prompted to save the changes you have made.
10. Click **Save**.

Database performance tuning

Database performance tuning can dramatically affect the throughput of your application. For example, if your application requires high concurrency (multiple, simultaneous interactions with backend data), an improperly tuned database can result in a bottleneck. Database access threads accumulate in a backlog when the database is not configured to accept a sufficient number of incoming requests.

Because WebSphere Application Server supports the integration of many different database products, each one with unique tuning configurations, consult your database vendor documentation for comprehensive tuning information. This information center provides introductory material in the following topics:

- DB2 Universal Database performance tips

DB2 Universal Database performance tips

You can easily adjust your system QSQRVR prestart job settings to optimize the process of acquiring connections from DB2 Universal Database™ for IBM i.

On the IBM i platform, QSQRVR jobs process Java Database Connectivity (JDBC) tasks. By default, five QSQRVR jobs are initially active. When fewer than two QSQRVR jobs are unused, the IBM i system creates two more jobs. If you increase the active job values, an application that establishes a large number of database connections over a short period of time might create connections more quickly. Increase the values for the initial number of jobs, threshold, and additional number of jobs by running this command on an IBM i command line:

```
CHGPJE SBS(D(QSYS/QSYSWRK) PGM(QSYS/QSQRVR)
```

Do not start more QSQRVR jobs than your application requires. Active QSQRVR jobs require some overhead, even if these jobs are not used.

Directory conventions

References in product information to *app_server_root*, *profile_root*, and other directories imply specific default directory locations. This topic describes the conventions in use for WebSphere Application Server.

Default product locations - IBM i

These file paths are default locations. You can install the product and other components in any directory where you have write access. You can create profiles in any valid directory where you have write access. Multiple installations of WebSphere Application Server products or components require multiple locations.

app_client_root

The default installation root directory for the Application Client for IBM WebSphere Application Server is the /QIBM/ProdData/WebSphere/AppClient/V8/client directory.

app_client_user_data_root

The default Application Client for IBM WebSphere Application Server user data root is the /QIBM/UserData/WebSphere/AppClient/V8/client directory.

app_client_profile_root

The default Application Client for IBM WebSphere Application Server profile root is the /QIBM/UserData/WebSphere/AppClient/V8/client/profiles/*profile_name* directory.

app_server_root

The default installation root directory for WebSphere Application Server Network Deployment is the /QIBM/ProdData/WebSphere/AppServer/V8/ND directory.

java_home

Table 1. Root directories for supported Java Virtual Machines.

This table shows the root directories for all supported Java Virtual Machines (JVMs).

JVM	Directory
32-bit IBM Technology for Java	/QOpenSys/QIBM/ProdData/JavaVM/jdk60/32bit
64-bit IBM Technology for Java	/QOpenSys/QIBM/ProdData/JavaVM/jdk60/64bit

plugins_profile_root

The default Web Server Plug-ins profile root is the /QIBM/UserData/WebSphere/Plugins/V8/webserver/profiles/*profile_name* directory.

plugins_root

The default installation root directory for Web Server Plug-ins is the /QIBM/ProdData/WebSphere/Plugins/V8/webserver directory.

plugins_user_data_root

The default Web Server Plug-ins user data root is the /QIBM/UserData/WebSphere/Plugins/V8/webserver directory.

product_library

product_lib

This is the product library for the installed product. The product library for each Version 8.0 installation on the system contains the program and service program objects (similar to .exe, .dll, .so objects) for the installed product. The product library name is QWAS8x (where x is A, B, C, and so on). The product library for the first WebSphere Application Server Version 8.0 product installed on the system is QWAS8A. The *app_server_root*/properties/product.properties file contains the value for the product library of the installation, was.install.library, and is located under the *app_server_root* directory.

profile_root

The default directory for a profile named *profile_name* for WebSphere Application Server Network Deployment is the /QIBM/UserData/WebSphere/AppServer/V8/ND/profiles/*profile_name* directory.

shared_product_library

The shared product library, which contains all of the objects shared by all installations on the system, is QWAS8. This library contains objects such as the product definition, the subsystem description, the job description, and the job queue.

user_data_root

The default user data directory for WebSphere Application Server Network Deployment is the /QIBM/UserData/WebSphere/AppServer/V8/ND directory.

The profiles and profileRegistry subdirectories are created under this directory when you install the product.

web_server_root

The default web server path is /www/*web_server_name*.

Chapter 4. Tuning EJB applications

This page provides a starting point for finding information about enterprise beans.

Based on the Enterprise JavaBeans (EJB) specification, enterprise beans are Java components that typically implement the business logic of Java 2 Platform, Enterprise Edition (J2EE) applications as well as access data.

EJB 2.1 container tuning

This page provides a starting point for finding information about tuning the Enterprise JavaBeans 2.1 container.

About this task

The following topic listed here to get started with tuning your EJB 2.1 container.

EJB container tuning

If you use applications that affect the size of the Enterprise JavaBeans (EJB) container cache, it is possible that the performance of your applications can be affected by an incorrect size setting. Container managed persistence (CMP) is discussed in this topic; although it is important to know that entity beans are not supported in an EJB 3.x module.

EJB cache size

Monitor Tivoli® Performance Viewer to diagnose whether the EJB container cache size setting is tuned correctly for your application.

If the application has filled the cache causing evictions to occur, Tivoli Performance Viewer shows a high rate of the `ejbStores` method being called and probably a lower than expected processor utilization on the workstation machine.

All applications that use enterprise beans must have this setting adjusted from the default value, if the following formula result equals more than 2000:

$$\begin{aligned} \text{EJB_Cache_Size} = & (\text{Largest number of Option B or C entity beans enlisted in a} \\ & \text{transaction * maximum number of concurrent transactions}) + \\ & (\text{Largest number of unique Option A entity beans expected to be accessed during} \\ & \text{typical application workload}) + \\ & (\text{Number of stateful session beans active during typical workload}) + \\ & (\text{Number of stateless session bean types used during typical workload}) \end{aligned}$$

Where:

Option B and C entity beans are only held in the EJB cache during the lifetime of the transaction they are enlisted in. Therefore, the first term in the formula computes the average EJB cache requirements for these types of beans.

Option A entity beans are held in the EJB cache indefinitely, and are only removed from the cache if there starts to become more beans in the cache than the cache size has been set to. If your application uses Read Only Beans, consider them to be Option A beans for this tuning calculation.

Stateful session beans are held in the EJB cache until they are removed by the application, or until their session timeout value is reached.

Only a single stateless session bean instance for each EJB type is held in the cache during the time any methods are running on that stateless session bean. If two or more methods are being implemented simultaneously on the same stateless session bean type, each method runs on its own bean instance, but only one cache location is used for all of these instances.

This formula calculates the upper bound on the maximum number of enterprise beans active at one time inside the application server. Because the cache of the EJB container is built to contain all these beans for performance optimizations, best performance can be achieved by setting this cache size to be larger than the number resulting from the previous formula.

You can set the EJB cache size in the administrative console under **Servers > Application Servers > *server_name* > EJB Container > EJB Cache Settings**.

Also, while adjusting the EJB cache size, you can tune the EJB container management thread parameter to meet the needs of the application. The management thread is controlled through the **Clean Up Interval** setting. This setting controls how frequently a daemon thread inside of the product attempts to remove bean instances from the cache that have not been used recently, attempting to keep the number of bean instances at or below the cache size. This behavior ensures that the EJB container places and looks up items in the cache quickly. Leave this interval set to the default; however, in some cases, it might be worthwhile to see if there is a benefit to reducing this interval.

For information about tuning the EJB cache using the EJB cache trace service, read about tuning the EJB cache and using the trace service.

EJB stateful session bean tuning

Stateful session bean timeout is configured in different ways, with different scopes, depending on the version of WebSphere Application Server.

WebSphere Application Server version 6.1 and earlier supports configuration of stateful session bean timeout, per bean, using the `ibm-ejb-jar-ext.xmi` file.

WebSphere Application Server version 7.0 supports configuration of stateful session bean timeout, per bean, using the `ibm-ejb-jar-ext.xmi` file (for EJB 2.x modules), and the `ibm-ejb-jar-ext.xml` file (for EJB 3.x modules).

WebSphere Application Server version 8 supports configuration of stateful session bean timeout, per bean, using the `ibm-ejb-jar-ext.xmi` (for EJB 2.x modules) and the `ibm-ejb-jar-ext.xml` (for EJB 3.x modules) files, and the stateful-timeout XML element and the `@StatefulTimeout` annotation. Additionally, you can configure a server-wide (global) stateful session timeout value using the `com.ibm.websphere.ejbcontainer.defaultStatefulSessionTimeout` system property.

Note: For IBM extension and binding files, the `.xmi` or `.xml` file name extension is different depending on whether you are using a pre-Java EE 5 application or module or a Java EE 5 or later application or module. An IBM extension or binding file is named `ibm-*-ext.xmi` or `ibm-*-bnd.xmi` where `*` is the type of extension or binding file such as `app`, `application`, `ejb-jar`, or `web`. The following conditions apply:

- For an application or module that uses a Java EE version prior to version 5, the file extension must be `.xmi`.
- For an application or module that uses Java EE 5 or later, the file extension must be `.xml`. If `.xmi` files are included with the application or module, the product ignores the `.xmi` files.

However, a Java EE 5 or later module can exist within an application that includes pre-Java EE 5 files and uses the `.xmi` file name extension.

The `ibm-webservices-ext.xmi`, `ibm-webservices-bnd.xmi`, `ibm-webservicesclient-bnd.xmi`, `ibm-webservicesclient-ext.xmi`, and `ibm-portlet-ext.xmi` files continue to use the `.xmi` file extensions.

Stateless session beans do not have a timeout value because they have no conversational state and are not dedicated to any specific client.

You can use Rational® Application Developer to update the `ibm-ejb-jar-ext.xml` file, which is used to configure the stateful session timeout value for beans in an EJB 2.x module. For more information, read about defining session timeout settings for a bean in the Rational Application Developer Information Center.

For example, the generated XML code to set a stateful session timeout value of 15 minutes is:

```
<ejbExtensions xmi:type="ejbext:SessionExtension" xmi:id="SessionExtension_1"
  timeout="15">
```

You can modify the `ibm-ejb-jar-ext.xml` file to set the stateful session timeout for beans in an EJB 3.x module. For example, the code to set a stateful session timeout value to 15 minutes for the `myBean` bean is:

```
<ejbExtensions xmi:type="ejbext:SessionExtension" xmi:id="SessionExtension"
  timeout="15">
  <enterpriseBean xmi:type="ejb:Session" href="META-INF/ejb-jar.xml#MyBean"/>
</ejbExtensions>
```

You can configure stateful session bean timeout using the `@StatefulTimeout` annotation. The `@StatefulTimeout` annotation takes a required value parameter representing the duration of the timeout, and an optional unit parameter. If the optional unit parameter is not specified, the default unit is minutes. The `@StatefulTimeout` annotation is introduced as part of EJB 3.1.

For example, use the `@StatefulTimeout` annotation to declare a timeout value of 100 seconds:

```
@StatefulTimeout(value=100 unit=java.util.concurrent.TimeUnit.SECONDS)
```

You can configure stateful session bean timeout using the `stateful-timeout` XML element in the `ejb-jar.xml` deployment descriptor. The `stateful-timeout` element is introduced as part of EJB 3.1.

For example, to set a timeout value of 100 seconds:

```
<stateful-timeout>
  <timeout>100</timeout>
  <unit>Seconds</unit>
</stateful-timeout>
```

The `@StatefulTimeout` annotation and the `stateful-timeout` XML element are the specification-defined mechanisms for declaring timeout values per bean, beginning with EJB 3.1. Before EJB 3.1, there is no specification-defined way for declaring stateful session timeout per bean. When using the `stateful-timeout` XML element or `@StatefulTimeout` annotation, a value of `-1` means that the bean never times out, and a value of `0` means that the bean is immediately eligible for removal.

You can configure stateful session bean timeout on a global (server-wide) basis, using the `com.ibm.websphere.ejbcontainer.defaultStatefulSessionTimeout` system property. The unit of time for the `com.ibm.websphere.ejbcontainer.defaultStatefulSessionTimeout` is minutes. The specified value must be `0` or greater, and if an invalid value is specified, the default value of 10 minutes is used instead. The global timeout value cannot be configured using XML or annotations. The global timeout value applies to all stateful session beans running in the server, including stateful session beans in EJB 2.x or EJB 3.x modules.

In WebSphere Application Server version 8, bean-level stateful timeout settings take precedence over the server-wide timeout setting. The server-wide timeout setting takes precedence over the default (unspecified) timeout. The following order of precedence is used to determine the stateful session timeout value for a bean running in WebSphere Application Server version 8:

1. `stateful-timeout` XML element
2. `@StatefulTimeout` annotation
3. `ibm-ejb-jar-ext.xml` or `ibm-ejb-jar-ext.xmi`

4. `com.ibm.websphere.ejbcontainer.defaultStatefulSessionTimeout` system property
5. If no bean-level or server-wide timeout value is explicitly specified, then the default timeout value of 10 minutes is applied.

Dcom.ibm.websphere.ejbcontainer.poolSize

If the application is using most bean instances in the pool, the Tivoli Performance Viewer indicates this. When the majority of the bean instances are used, increase the size of those bean pools that are being exhausted by adding this parameter in the custom properties tag of the JVM; for example:

```
-Dcom.ibm.websphere.ejbcontainer.poolSize=<application_name>#<module_name>#  
<enterprisebean_name>=<minSize>,<maxSize>
```

Where:

The `<application_name>` element is the Java EE application name as defined in the enterprise archive (EAR) file deployment descriptor, for the bean whose pool size is being set.

The `<module_name>` element is the Java archive (JAR) file name of the EJB module, for the bean whose pool size is being set.

The `<bean_name>` element is the Java EE enterprise bean name as defined in the EJB module deployment descriptor, for the bean whose pool size is being set

The `<minSize>` element is the number of bean instances the container maintains in the pool, irrespective of how long the beans have been in the pool (beans greater than this number are cleared from the pool over time to optimize memory usage)

The `<maxSize>` element is the number of bean instances in the pool where no more bean instances are placed in the pool after they are used (that is, when the pool is at this size, any additional beans are discarded rather than added into the pool, which ensures the number of beans in the pool have an upper limit so that memory usage does not grow in an unbounded way).

To keep the number of instances in the pool at a fixed size, the `minSize` and `maxSize` elements can be set to the same number. A separate instance pool for every EJB type running in the application server exists, and that every pool starts out with no instances; the number of instances grow as beans are used and then placed in the pool.

When a bean instance is needed by the container and no beans are available in the pool, the container creates a bean instance, uses it, then places that instance in the pool, unless there are already `maxSize` instances in the pool. For example, the statement -

```
Dcom.ibm.websphere.ejbcontainer.poolSize=ivtApp#ivtEJB.jar#ivtEJBObject=125,1327
```

would set a `minSize` of 125 and a `maxSize` of 1327 on the bean named `ivtEJBObject` within the `ivtEJB.jar` file, in the application `ivtApp`.

The application, `ivtApp`, is replaced by the actual application name, the `ivtEJB.jar` file is replaced by the JAR file that contains the bean that must have its pool size increased, and `ivtEJBObject` is the bean name of the enterprise bean whose pool size must be increased. The minimum number of beans that are held in the pool is 125. The maximum number of beans that are held in the pool is 1327. Set these so that no more evictions occur from the pool. In most cases these must be set equal if memory is plentiful because no growth and shrinkage of the pool occurs.

Dcom.ibm.websphere.ejbcontainer.noPrimaryKeyMutation

You must understand how your application handles the creation of primary key objects for use by CMP beans and bean-managed persistence (BMP) beans inside of the product.

The EJB container uses the primary key of an entity bean as an identifier inside of internal data structures to optimize performance. However, the EJB container must copy these primary key objects upon the first access to the bean to ensure that the objects stored in the internal caches are separate from the ones used in an application. This process occurs to keep the internal structures consistent in case the application changes or mutates the primary key. If the application does not mutate any of the primary keys used to create and access entity beans after they are created, a special flag can be used that ensures that the EJB container skips the copy of the primary key object, saving processor cycles, and increasing performance. This mechanism can be enabled at your own risk by adding the **-D** property to the JVM custom property field.

```
-Dcom.ibm.websphere.ejbcontainer.noPrimaryKeyMutation=true
```

Note: Entity beans are not supported in EJB 3.x and later modules.

The performance benefit of this optimization depends on the application. If the application uses primitive types for the primary keys of enterprise beans, there is no gain because these objects are already immutable and the copy mechanism takes this into account. If, however, the application uses many complex primary keys, that is, an object for a primary key or multiple fields, this parameter can yield significant improvements.

Dcom.ibm.ws.pm.deferredcreate

The persistence manager is used by the EJB container to persist data to the database from CMP entity beans.

When creating entity beans by calling the `ejbCreate` method, by default the persistence manager immediately inserts the empty row with only the primary key in the database. In most cases, after creating the bean you must modify fields in the bean created or in other beans inside of the same transaction. If you want to postpone the insert into the database until the end of the transaction to eliminate one trip to the database, set the **-D** flag inside of the JVM custom properties field. The data is inserted into the database and consistency is maintained.

Note: Entity beans are not supported in EJB 3.x and later modules.

```
-Dcom.ibm.ws.pm.deferredcreate=true
```

The performance benefit of this optimization depends on the application. If the EJB applications transactions are insert intensive, the application can benefit from this optimization. If the application performs few inserts, the benefit of this optimization is less.

Dcom.ibm.ws.pm.batch

When an EJB application accesses multiple CMP beans inside of a single transaction, depending on the operations performed on the beans, such as updates, inserts, and reads, the number of operations issued to the database corresponds directly to the operations performed on the CMP beans. If the database system you are using supports batching of update statements, you can enable this flag and increase performance on all interactions with the database that involve more than two updates in a single transaction.

Note: Entity beans are not supported in EJB 3.x and later modules.

Use this flag, which supports the persistence manager adding all the update statements into one single batch statement that is issued to the database. This process saves round trips to the database, which increases performance. If you know that your application exhibits the behavior of updating multiple CMP beans in a single transaction, and the database supports batch updates, you can set the **-D** flag inside of the JVM custom properties field; for example:

```
-Dcom.ibm.ws.pm.batch=true
```

The performance benefit of this optimization depends on the application. If the application never or infrequently updates CMP beans, or updates only a single bean per transaction, there is no performance gain. If the application updates multiple beans per transaction, this parameter benefits the applications performance.

The following table lists the backend databases that support batch update.

Table 2. Backend databases that support batch update. The following table lists the backend databases that support batch update.

Database	Supports batch update	Supports batch update with Optimistic Concurrency Control
DB2®	yes	no
Oracle	yes	no
DB2 Universal Driver	yes	yes
Informix®	yes	yes
SQLServer	yes	yes
Apache Derby	yes	yes

Note: Batch update with OCC cannot be performed for databases that do not support it, even if specified by the access intent.

com.ibm.ws.pm.useLegacyCache

Specifies the name of the Java class that the product uses to implement the javax.rmi.CORBA.UtilDelegate interface.

Note: Entity beans are not supported in EJB 3.x and later modules.

Persistence manager has two types of caching mechanisms, *legacy cache* and *two-level cache*. Typically, two-level cache is more efficient than legacy cache because of optimizations in this mode. The default is legacy cache, although two-level cache is recommended. Set this configuration through the system property as follows:

```
com.ibm.ws.pm.useLegacyCache=false
```

com.ibm.ws.pm.grouppartialupdate and com.ibm.ws.pm.batch

The partial updates feature enhances the performance of applications with enterprise beans in certain scenarios. Persistence manager has two caching mechanisms available, *legacy cache* and *two-level cache*. Typically, two-level cache performs better than legacy cache because of the optimizations in this mode.

Note: Entity beans are not supported in EJB 3.x and later modules.

In certain applications where you must perform both batch updates and partial updates, you must configure the following system properties to gain the benefits of both:

```
'com.ibm.ws.pm.grouppartialupdate=true' and 'com.ibm.ws.pm.batch=true'
```

Tuning Enterprise JavaBeans applications

Tuning EJB cache with trace service

The size of your Enterprise JavaBeans (EJB) cache can affect the performance of the application server. One of the steps in tuning your EJB container to optimum performance levels is to fine-tune the EJB cache.

Before you begin

Note: This topic references one or more of the application server log files. Beginning in WebSphere Application Server Version 8.0 you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using `SystemOut.log`, `SystemErr.log`, `trace.log`, and `activity.log` files or native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

About this task

The following procedure describes how to use the diagnostic trace service to help determine the best cache size.

Procedure

1. Enable the EJB cache trace. To learn about working with the trace service, see the topic, *Working with trace*. For information about the trace service settings, see the topic, *Diagnostic trace service settings*.

Set up your trace to use this trace string:

```
com.ibm.ejs.util.cache.BackgroundLruEvictionStrategy=all=enabled:com.ibm.ejs.util.cache.CacheElementEnumerator=all=enabled
```

Set **Maximum File Size** to 200MB or more. If you leave the default value of 20MB, you could fill up the single 20 MB trace log and lose some data because of trace wrapping.

Set **Maximum Number of Historical Files** to 5. Five files should be sufficient, but if you see that all five files are full and trace wrapping occurs, increase this value.

2. Stop your server, delete existing logs, then start your server.
3. Run typical scenarios to capture cache trace data. By running a typical scenario with the trace enabled, you get the EJB cache trace data to analyze in the following steps.
4. View and analyze the trace output.
 - a. Open your trace log. Look for either or both of the following trace strings to display:

```
BackgroundLru 3 EJB Cache: Sweep (1,40) - Cache limit not reached : 489/2053
```

```
BackgroundLru > EJB Cache: Sweep (16,40) - Cache limit exceeded : 3997/2053 Entry
```

In the trace strings that include the words **Cache limit** you find a ratio. For example, 3997/2053.

The first number is the number of enterprise beans currently in the EJB cache (called the *capacity*).

The second number is the EJB cache setting (more about this in later steps). Use this ratio, particularly the capacity, in your analysis.

Also look for the statements *Cache limit not reached* and *Cache limit exceeded*.

Cache limit not reached

Your cache is equal to or larger than what is appropriate. If it is larger, you are wasting memory, and should reduce the cache size to a more appropriate value.

Cache limit exceeded

The number of beans currently being used is greater than the capacity you have specified, indicating that your cache is not properly tuned. The capacity can exceed the EJB Cache setting because the setting is not a hard limit. The EJB Container does not stop adding beans to the cache when the limit is reached. Doing so could mean that when the cache is full, a request for a bean would not be fulfilled, or would at least be delayed until the cache fell below the limit. Instead, the cache limit can be exceeded, but the EJB Container attempts to clean up the cache and keep it below the EJB Cache size.

In the case where the cache limit is exceeded, you might see a trace point similar to this:

```
BackgroundLru < EJB Cache: Sweep (64,38) - Evicted = 50 : 3589/2053 Exit
```

Notice the *Evicted =* string. If you see this string, you are using either Stateful Session Beans or Entity Beans configured for Option A or B caching. Evicted objects mean that you are not taking full advantage of the caching option that you have chosen. Your first step is to try increasing the EJB Cache size. If continued running of your application results in more evictions, it means that the application is accessing or creating more new beans between EJB Cache sweeps than the cache can hold, and *NOT* reusing existing beans.

You might want to consider using Option C caching for your entity beans, or checking your application to see if it is not removing Stateful Session Beans after they are no longer needed.

Note: Entity beans configured with Option C caching are only in the cache while in a transaction, and are required to be held in the cache for the entire transaction. Therefore, they are never evicted during a cache sweep, but are removed from the cache when the transaction completes. In addition, if you are using only Stateless Session Beans or Entity Beans with Option C caching (or both), then you might want to increase your EJB Cache *cleanup interval* to a larger number. The cleanup interval can be set as described in EJB cache settings. Stateless Session Beans are **NOT** in the EJB Cache, and since Entity Beans using Option C caching are never evicted by the caching (LRU) strategy, there is really no need to sweep often. When using only Stateless Session Beans or Option C caching, you should only see "Evicted = 0" in the trace example shown above.

- b. Analyze your trace log. Look for the trace string *Cache limit exceeded*.
 - You might find more than one instance of this string. Examine them all to find the highest capacity value of beans in the EJB Cache. Reset your EJB Cache size to about 110% of this number. Setting the EJB Cache size is explained in a later step.
 - You might find no instances of this string. This means that you have not exceeded the capacity of the EJB Cache (which is your end goal), but not seeing it during your initial analysis could also mean that your cache is too large and using unnecessary memory. In this case, you still must tune your cache by reducing the cache size until your cache limit is not exceeded, then increasing it to the optimum value. Setting the EJB Cache size is explained in a later step.

Your ultimate goal is to set the cache limit to a value that does not waste resources, but also does not get exceeded. A good set-up gives you a trace with only the *Cache limit not reached* message, and a ratio where the capacity number can be near, but below, 100% of the EJB Cache setting.

Note: It is recommended that you do not set your cache size to anything less than the default of 2053.

5. Modify the cache settings based on your analysis. See EJB cache settings for information about how to do this.
6. Stop your server, delete all logs, and restart your server.
7. Repeat the previous steps until you are satisfied with your settings.
8. Disable the EJB Cache trace. With the cache properly tuned, you can remove the trace, remove old logs, and restart your server.

What to do next

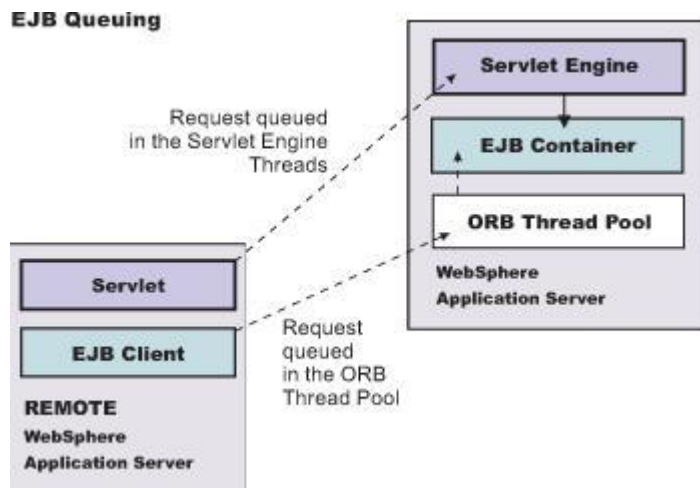
From your analysis, it is possible to set the EJB cache optimally from an EJB Container perspective, but perhaps not optimally from a WebSphere Application Server perspective. A larger cache size provides more hits and better EJB cache performance, but uses more memory. Memory used by the cache is not available to other areas of the product, potentially causing the overall performance to suffer. In a system with ample memory, this might not be an issue and properly tuning the EJB cache might increase overall performance. However, you should take into account this system performance versus EJB cache performance when configuring the cache.

EJB method Invocation Queuing

Method invocations to enterprise beans are only queued for remote clients making the method call. An example of a remote client is an Enterprise JavaBeans (EJB) client running in a separate Java virtual machine (JVM) (another address space) from the enterprise bean. In contrast, no queuing occurs if the EJB client, either a servlet or another enterprise bean, is installed in the same JVM on which the EJB method runs, and on the same thread of execution as the EJB client.

Remote enterprise beans communicate by using the Remote Method Invocation over Internet Inter-ORB Protocol (RMI-IIOP). Method invocations initiated over RMI-IIOP are processed by a server-side object request broker (ORB). The thread pool acts as a queue for incoming requests. However, if a remote method request is issued and there are no more available threads in the thread pool, a new thread is created. After the method request completes the thread is destroyed. Therefore, when the ORB is used to process remote method requests, the EJB container is an open or closed queue, due to the use of unbounded threads.

The following illustration depicts the two queuing options of enterprise beans.



Tuning applications that use the Java Persistence API

Configuring OpenJPA caching to improve performance

The OpenJPA implementation gives you the option of storing frequently used data in the memory to improve performance. OpenJPA provides concurrent data and concurrent query caches that support applications to save persistent object data and query results in memory to share among threads and for use in future queries.

About this task

OpenJPA data cache functionality

The OpenJPA data cache is a cache of persistent object data that operates at the EntityManagerFactory level. This optional-use cache is designed to increase performance while remaining in full compliance with the Java Persistence API (JPA) standard. This means that enabling the caching option can increase the performance of your application, with no changes to your code. The OpenJPA data cache is designed to provide significant performance increases over cacheless operations and ensures that behavior is identical in both cache-enabled and cacheless operations.

When enabled, the cache is examined before accessing the data store. The cache stores data when objects are committed and when persistent objects are loaded from the data store. If operating in a single

Java virtual machine (JVM) environment, the JVM maintains and shares a data cache across all EntityManager instances obtained from a particular EntityManagerFactory. The OpenJPA data cache cannot do this in a distributed environment because caches in different JVMs, created from different EntityManagerFactory objects are not synchronized.

Using the OpenJPA cache in a multi-JVM environment can be done by configuring the OpenJPA second level (L2) cache provider plug-in. See the topic, Dynamic cache provider for the JPA 2.0 second level cache, and the section "Using the dynamic cache L2 Cache Provider in a clustered environment", for more information. Configuring the DynaCache plug-in allows for the Data and Query cache content to be replicated and consistent across JVMs. Other alternatives include setting up an event notification framework or using a third-party distributed cache such as IBM WebSphere eXtreme Scale.

Enabling and configuring the OpenJPA data cache

You can enable the OpenJPA data cache for a single or a multiple JVM environment, set its default element size, including soft references, and specify **timeout** values.

To set up and configure the OpenJPA data cache, do the following:

1. To enable the cache for a single JVM, set the `openjpa.DataCache` property to `true`, and set the `openjpa.RemoteCommitProvider` property to `sjvm`:

```
<property name="openjpa.DataCache" value="true"/>
<property name="openjpa.RemoteCommitProvider" value="sjvm"/>
```

To enable the data cache in a distributed environment, the `openjpa.RemoteCommitProvider` must be configured specifically for the environment, or a third-party cache management utility can be used.

2. The maximum cache size can be adjusted by setting the `CacheSize` property:

```
<property name="openjpa.DataCache" value="true(CacheSize=5000...
```

By default, the OpenJPA data cache holds 1000 elements. Objects that are pinned into the cache are not counted when determining if the cache size exceeds its maximum size. If the cache overflows, it evicts random elements. You can preserve evicted elements longer with the `SoftReferenceSize` property. By default, soft references are unlimited. If you must, you can limit the number of soft references or set to 0 to disable soft references completely:

```
<property name="openjpa.DataCache" value="true(CacheSize=5000 SoftReferenceSize=0 ...
```

3. You can specify that a cache is cleared at certain times. The `EvictionSchedule` property of the OpenJPA cache implementation accepts a cron style eviction schedule. The cron format specifies the minute, hour of day, day of month, day of month, and day of the week beginning with 1 for Sunday; the * symbol (asterisk), indicates match all. To schedule a cache to evict at 45 minutes past 3 PM on Sunday every month you would add this property:

```
<property name="openjpa.DataCache" value="true(CacheSize=5000 SoftReferenceSize=0 EvictionSchedule='15,45 * * 1')/>
```

4. You also can specify a cache timeout value for a single class by setting the `timeout` metadata extension to the amount of time in milliseconds that the data of the class is valid; for example:

```
@Entity
@DataCache(timeout=10000)
public class Employee {
    ...
}
```

Refer to the `org.apache.openjpa.persistence.DataCache` Javadoc for more information.

After configuring your data cache, you can use it after you restart your application.

Refresh with active DataCache

Refreshing an entity may lead to different behavior with or without a `DataCache` when a separate process or part of the same application are updated or even deleted the corresponding record in the database. By

default, entities are refreshed from the database even when DataCache is active. Therefore, with the default configuration the refresh behaves identically with or without a DataCache. However, a persistence unit can be configured to refresh entities from DataCache with the property `openjpa.RefreshFromDataCache` for improved performance. Under this configuration, any out-of-band changes that occur in the database record do not appear in the refreshed state of the entity.

Note: Regardless of the `openjpa.RefreshFromDataCache` setting, the DataCache is always bypassed for refresh when locks are active, such as for a pessimistic transaction, in a persistence context. An application may activate `openjpa.RefreshFromDataCache` but can still bypass the DataCache while refreshing an entity by explicitly evicting the entity from DataCache before refresh.

Query Caching functions

OpenJPA provides a concurrent query cache that supports applications to save persistent object data and query results in memory to share among threads and for use in future queries. The query cache stores the object IDs returned by query executions. When you run a query, OpenJPA assembles a key based on the query properties and the parameters used at execution time, and checks for a cached query result. If one is found, the object IDs in the cached result are looked up, and the resultant persistence-capable objects are returned. Otherwise, the query is executed against the database, and the object IDs loaded by the query are put into the cache.

Configuring or disabling the query cache

You can configure the query cache settings in a similar way to the data cache. The interface provided to the query cache is the `org.apache.openjpa.persistence.QueryResultCache` class. You can access this class through the `OpenJPAEntityManagerFactory`.

The default query cache implementation caches 100 query executions in a *least-recently-used* cache. This can be changed by setting the cache size in the `CacheSize` plug-in property. Like the data cache, the query cache also has a backing soft reference map that can be changed using the `SoftReferenceSize` property. To keep queries in the cache at all times, you can pin them to a cache. To change the query cache properties do the following:

1. Modify the `CacheSize` property of the `openjpa.QueryCache`:

```
<property name="openjpa.QueryCache" value="true(CacheSize=1000, ...
```

2. Change the `SoftReferenceSize` property to enable and control the size of this map:

```
<property name="openjpa.QueryCache" value="true(CacheSize=1000, SoftReferenceSize=100)"/>
```

The `SoftReferenceSize` table is disabled by default. Setting the size enables it.

3. Pin or unpin queries in the cache through the `QueryResultCache` with this syntax:

```
public void pin(Query q);  
public void unpin(Query q);
```

Modifying these properties allows you to make better use of the query cache.

Extending a cache

OpenJPA provides classes that may be extended for further functionality.

- As previously mentioned, if you want to implement a distributed cache that uses an unsupported method for communications, create an implementation of `org.apache.openjpa.event.RemoteCommitProvider`.
- If you are adding new behavior, extend `org.apache.openjpa.datacache.DataCacheImpl`.
- To use your own storage mechanism, extend `org.apache.openjpa.datacache.AbstractDataCache`.
- To add query functionality, you can extend the default `org.apache.openjpa.datacache.QueryCacheImpl`.

- Implement your own storage mechanism for query results by extending `org.apache.openjpa.datacache.AbstractQueryCache`

OpenJPA query SQL cache

OpenJPA provides a cache that provides caching of SQL strings used by find operations performed on the entity manager and some queries to manage eagerly fetched relationships. When this cache is enabled, SQL queries used by these operations are generated one time per entity manager factory and can be reused. This cache is enabled by default but can also be configured through the `openjpa.jdbc.QuerySQLCache` configuration property.

Configuring or disabling the SQL query cache

The query SQL cache can be configured or disabled through the `openjpa.jdbc.QuerySQLCache` property. By default, this property is set to `true`. When the property is set to `true`, the cache is enabled and uses the `org.apache.openjpa.util.CacheMap` class for its cache store. The `CacheMap` is a managed cache, meaning that it limits the number of cache entries and has a cache eviction scheme to manage memory usage. If the cache is set to `all` the `org.apache.openjpa.lib.util.ConcurrentHashMap` class is used as a cache store. The `ConcurrentHashMap` is not a managed cache so entries remain in the cache for the lifetime of an entity manager factory. This caching mechanism can provide better performance at the expense of increased memory usage. A custom cache store class can also be specified if it implements the `java.util.Map` interface. To disable the cache, specify the value `false`. See the following examples on how to configure or disable the SQL query cache:

- To use an unmanaged cache:

```
<property name="openjpa.jdbc.QuerySQLCache" value="all"/>
```
- To specify a custom cache class:

```
<property name="openjpa.jdbc.QuerySQLCache" value="com.mycompany.MyCustomCache"/>
```
- To use an unmanaged cache:

```
<property name="openjpa.jdbc.QuerySQLCache" value="false"/>
```

What to do next

You can read more about Caching in the OpenJPA for all caching extensions in the Apache OpenJPA User Guide.

Configuring the WSJPA Object Cache to improve performance

The WebSphere Java Persistence API (WSJPA) extension to OpenJPA provides a read-only object cache that can improve performance in certain use cases.

Before you begin

This cache can improve the performance of an application which has a set of data that is used in a static, read-only method. For example, accessing basic persistent fields and persisting unidirectional relationships to a read-only type. This set of data has a number of restrictions that are documented here.

The WSJPA object cache is a non-distributed cache of read-only Entities that operates at the `EntityManagerFactory` level. These cached instances are shared by all `EntityManagers` in the Java virtual machine (JVM), but not managed by any. When enabled, the `ObjectCache` is examined before accessing the OpenJPA `DataCache` and database. When persistent objects are loaded from the database they are stored in the OpenJPA `ObjectCache`. The `ObjectCache` can be used with the OpenJPA `DataCache` and `QueryCache` for even greater performance. Types that are included in the `ObjectCache` **must not** be eligible to be cached in the OpenJPA `DataCache`. The object cache is not to be confused with a second-level cache as defined by the JPA 2.0 specification.

Types that are included in the OpenJPA ObjectCache are as follows:

- Strictly **read-only** from the application point-of-view
 - Passing a read only type into the following operations results in an UnsupportedOperationException.
 - Passing a read only Entity into EntityManager.merge(...).
 - Passing a read only Entity into EntityManager.persist(...).
 - Passing a read only Entity into EntityManager.remove(...).
 - Calling a setter method on a read-only type that was returned by the WebSphere JPA runtime results in an UnsupportedOperationException.
- Restricted to having only basic fields. An exception occurs on EntityManager creation if this rule is broken.
- Not to intersect types that are cacheable by the OpenJPA L2 cache (openjpa.DataCache). An exception occurs on EntityManager creation if this rule is broken.
- Passing a read-only Entity into EntityManager.contains(...) always returns false, even if it was just returned from a find/query operation.

About this task

You can enable the object cache for a single JVM environment, specify the types that are included in this cache, set its maximum element size and specify timeout values. To set up and configure the object cache, do the following:

Attention: The preferred property name is wsjpa.ObjectCache, but openjpa.ObjectCache is also a valid configuration.

Procedure

1. To enable the cache for a single JVM, set the wsjpa.ObjectCache property to true and specify a list of Types.

```
<property name="wsjpa.ObjectCache" value="true(Types=com.ibm.wsjpa.Foo; com.ibm.wsjpa.Bar)"/>
```

2. Adjust the maximum cache size by setting the MaxSize property:

```
<property name=" wsjpa.ObjectCache" value="true(Types=com.ibm.wsjpa.Foo; com.ibm.wsjpa.Bar, MaxSize=5000)"/>
```

By default, the object cache holds 1000 elements. If the cache overflows, it evicts random elements. Evicted elements are preserved in a soft reference map. The size of the soft reference map is unlimited and cannot be configured.

3. Specify that the ObjectCache should be cleared at certain times. The EvictionSchedule property of the object cache implementation accepts a cron-style eviction schedule. The cron format specifies the minute, hour of day, day of month, and day of the week beginning with 1 for Sunday; the * (asterisk) indicates match all. To schedule a cache to evict at 3:45 PM on Sunday every month, add this property:

```
<property name="openjpa.DataCache" value="true(Types=com.ibm.wsjpa.Foo; com.ibm.wsjpa.Bar, MaxSize=5000, EvictionSchedule='15,45 * * 1')"/>
```

In addition to the cron syntax you can specify an interval style eviction schedule. The format of this property is a plus sign (+) followed by the number of minutes between each time that the cache should be evicted. To schedule a cache to evict every 20 minutes add this property:

```
<property name="openjpa.DataCache" value="true(Types=com.ibm.wsjpa.Foo; com.ibm.wsjpa.Bar, MaxSize=5000, EvictionSchedule='+20')"/>
```

What to do next

Read about caching in the OpenJPA User Guide for information about all caching extensions.

Chapter 5. Tuning Messaging resources

This page provides a starting point for finding information about the use of asynchronous messaging resources for enterprise applications with WebSphere Application Server.

WebSphere Application Server supports asynchronous messaging based on the Java Message Service (JMS) and the Java EE Connector Architecture (JCA) specifications, which provide a common way for Java programs (clients and Java EE applications) to create, send, receive, and read asynchronous requests, as messages.

JMS support enables applications to exchange messages asynchronously with other JMS clients by using JMS destinations (queues or topics). Some messaging providers also allow WebSphere Application Server applications to use JMS support to exchange messages asynchronously with non-JMS applications; for example, WebSphere Application Server applications often need to exchange messages with WebSphere MQ applications. Applications can explicitly poll for messages from JMS destinations, or they can use message-driven beans to automatically retrieve messages from JMS destinations without explicitly polling for messages.

WebSphere Application Server supports the following messaging providers:

- The WebSphere Application Server default messaging provider (which uses service integration as the provider).
- The WebSphere MQ messaging provider (which uses your WebSphere MQ system as the provider).
- Third-party messaging providers that implement either a JCA Version 1.5 resource adapter or the ASF component of the JMS Version 1.0.2 specification.

Tuning messaging

To tune asynchronous messaging, you can, for example, configure message-driven bean (MDB) throttling for a JCA 1.5-compliant messaging provider.

About this task

To tune your asynchronous messaging, complete one or more of the following steps.

Procedure

- Configure MDB throttling for a JCA 1.5-compliant messaging provider:
 - For the default messaging provider, complete the task “Configuring MDB throttling for the default messaging provider” on page 32.
 - For the WebSphere MQ messaging provider, set the **maximum server sessions** property on the WebSphere MQ messaging provider activation specifications advanced properties page, or the **maxPoolSize** property when using the `createWMQActivationSpec` or `modifyWMQActivationSpec` wsadmin commands.
 - For a third-party JCA 1.5-compliant JMS messaging provider, refer to the generic method of message throttling described in “Throttling inbound message flow for JCA 1.5 message-driven beans” on page 10.
- Tune messaging destinations for the WebSphere MQ or V5 default messaging provider.

Tuning messaging performance with service integration technologies

To help optimize performance, you can set tuning properties that control the performance of message-driven beans and other messaging applications deployed to use service integration technologies.

About this task

To optimize the performance of messaging with service integration technologies, you can use the administrative console to set various parameters as described in the steps below. You can also set these parameters by using the wsadmin tool.

Procedure

- View the Available Message Count on a destination.

Viewing the Available Message Count on a destination enables you to determine whether your message consumers are able to cope with your current workload. If the available message count on a given destination is too high, or is increasing over time, consider some of the tuning recommendations in this topic.

 1. Enable AvailableMessageCount statistics for a queue. If you restart the administrative server, enable **AvailableMessageCount** statistics again because such runtime settings are not preserved when the server is restarted.
 - a. In the navigation pane, click **Monitoring and Tuning -> Performance Monitoring Infrastructure (PMI)**.
 - b. In the content pane, click **server_name**.
 - c. Click the Runtime tab.
 - d. In the Currently monitored statistic set, click **Custom**.
 - e. On the Custom monitoring level panel, click **SIB Service > SIB Messaging Engines > engine_name > Destinations > Queues > queue_name**.
 - f. Select the AvailableMessageCount option.
 - g. Click **Enable** at the top of the panel.
 2. View the available message count for a queue.
 - a. In the navigation pane, click **Monitoring and Tuning -> Performance Viewer -> Current activity**.
 - b. In the content pane, click **server_name**.
 - c. Click **Performance Modules > SIB Service > SIB Messaging Engines > engine_name > Destinations > Queues > queue_name**.
 - d. Click **View Module(s)** at the top of the Resource Selection panel, located on the left side. This displays the AvailableMessageCount data in the Data Monitoring panel, located on the right side. You can use the Data Monitoring panel to manage the collection of monitoring data; for example, you can use the buttons to start or stop logging, or to change the data displayed as either a table or graph.
- Monitor MDB Thread Pool Size for the Default Message Provider.

You might experience a performance bottleneck if there are insufficient threads available for the message-driven beans. There is a trade-off between providing sufficient threads to maximize the throughput of messages and configuring too many threads, which can lead to CPU starvation of the threads in the application server. If you notice that the throughput for express nonpersistent, reliable nonpersistent, or reliable persistent messaging has fallen as a result of increasing the size of the default thread pool, then decrease the size of the thread pool and reassess the message throughput.

 1. View or change the number of threads in the default thread pool for an application server. By default, message-driven beans use the default thread pool.
 - a. Click **Servers -> Server Types -> WebSphere application servers -> server_name -> [Additional Properties] Thread Pools > Default**. By default the Minimum size value is set to 5 and the Maximum size value is set to 20. The best performance is obtained by setting the Maximum size value to the expected maximum concurrency for all message-driven beans. For high throughput using a single message bean, 41 was found to be the optimal Maximum size value.
 - b. Change the Maximum size value, then click **OK**.

2. Optional: Create your own thread pool. The default thread pool is also used by other WebSphere Application Server components, so you might want to define a separate thread pool for the message-driven beans. This reduces thread contention for the default thread pool.
 - a. Click **Servers -> Server Types -> WebSphere application servers -> server_name -> [Additional Properties] Thread Pools**.
 - b. Create a new thread pool.
 - c. Create sufficient threads to support the maximum amount of concurrent work for the message-driven beans.
 - d. Change the SIB JMS Resource Adapter to use the new thread pool:
 - 1) Click **Resources -> Resource Adapters -> Resource adapters**.
 - 2) If you cannot see any SIB JMS Resource Adapter instances in the list, expand **Preferences** and enable **Show built-in resources**.
 - 3) Select the **SIB JMS Resource Adapter** with the appropriate scope depending upon the scope of the connection factories.
 - 4) Add the name of the new thread pool in the **Thread pool alias** box.
 - 5) Click **Apply** .
3. Save your changes to the master configuration.
- Tune MDB performance with the default messaging provider.
 1. Click **Resources -> JMS -> Activation specifications -> activation_specification_name**.
 2. Set the maximum batch size for this activation specification.

Delivering batches of messages to each MDB endpoint can improve performance, particularly when used with Acknowledge mode set to Duplicates-ok auto-acknowledge. However, if message ordering must be retained across failed deliveries, set this parameter to 1.
 3. Set the maximum number of concurrent endpoints for this activation specification.

The maximum concurrent endpoints parameter controls the amount of concurrent work that can be processed by a message bean. The parameter is used with message-driven beans. Increasing the number of concurrent endpoints can improve performance but can increase the number of threads in use at one time. To benefit from a change in this parameter, there should be sufficient threads available in the MDB thread pool to support the concurrent work. However, if message ordering must be retained across failed deliveries, set this parameter to 1.
 4. Save your changes to the master configuration.

For additional information about tuning the throttling of message-driven beans, including controlling the maximum number of instances of each message bean and the message batch size for serial delivery, see “Configuring MDB throttling for the default messaging provider” on page 32.

- Ensure that application servers hosting one or more messaging engines are provided with an appropriate amount of memory for the message throughput you require.

You can tune the initial and maximum Java Virtual Machine (JVM) heap sizes when adding a server to a messaging bus, that is when you create a messaging engine. You have the option to do this in any of the following cases:

 - When adding a single server to a bus
 - When adding a cluster to a bus
 - When adding a new server to an existing cluster that is itself a bus member

For an application server that is a bus member of at least one bus, or a member of a cluster that is a bus member of at least one bus, the recommended initial and maximum heap sizes are 768MB.

When you are adding a cluster to a bus, you are recommended to increase the initial and maximum JVM heap sizes for every server in the cluster to 768MB.

 - Increasing the initial heap size improves the performance for small average message sizes
 - Increasing the maximum heap size improves the performance for higher average message sizes
- Reduce the occurrence of OutOfMemoryError exceptions

If the cumulative size of the set of messages being processed within a transaction by the service integration bus is large enough to exhaust the JVM heap, `OutOfMemoryError` exceptions occur. Consider one of the following options for reducing the occurrence of `OutOfMemoryError` exceptions when processing a large set of messages within a transaction.

- Increase the JVM heap sizes for the application server.
- Reduce the cumulative size of the set of messages being processed within the transaction.
- Change the maximum connections in a connection factory for the default messaging provider.
The maximum connections parameter limits the number of local connections. The default is 10. This parameter should be set to a number equal to or greater than the number of threads (enterprise beans) concurrently sending messages.
 1. Click **Resources** -> **JMS** -> **Topic connection factories** -> *factory_name* > **[Additional Properties] Connection pool properties**.
 2. Enter the required value in the **Maximum connections** field.
 3. Click **Apply**.
 4. Save your changes to the master configuration.
- Tune reliability levels for messages.

The reliability level chosen for the messages has a significant impact on performance. In order of decreasing performance (fastest first), the reliability levels are:

Best effort nonpersistent

Express nonpersistent

Reliable nonpersistent

Reliable persistent

Assured persistent

For MDB point-to-point messaging, best effort nonpersistent throughput is more than six times greater than assured persistent. For more information about reliability levels, see *Message reliability levels - JMS delivery mode and service integration quality of service*.

Configuring MDB throttling for the default messaging provider

Use this task to configure the throttling of messages for message-driven beans that you have deployed as JCA 1.5 resources on the default messaging provider.

Before you begin

The throttling support described in this topic only applies to the default messaging provider (the service integration JMS Resource Adapter).

For the WebSphere MQ messaging provider, you configure message throttling by setting the **maximum server sessions** property on the WebSphere MQ messaging provider activation specifications panel, or the **maxPoolSize** property when using the `createWMQActivationSpec` or `modifyWMQActivationSpec` wsadmin commands.

If you have a third-party JCA 1.5-compliant JMS messaging provider, refer to the generic method of message throttling described in “Throttling inbound message flow for JCA 1.5 message-driven beans” on page 10.

About this task

Use this task if you want to throttle messages for a message-driven bean deployed as a J2EE Connector Architecture (JCA) 1.5 resource on the default messaging JMS provider.

The default messaging provider (the service integration JMS Resource Adapter) uses a special type of message throttling. You can leave the message-driven bean pools to the default size of 500.

The default messaging provider enables the throttling of message delivery to a message-driven bean through the **Maximum concurrent endpoints** configuration option on the JMS activation specification used to deploy the bean.

- The maximum number of instances of each message-driven bean is controlled by the Maximum concurrent endpoint setting in the activation specification used to deploy the message-driven bean. This maximum concurrency limit helps prevent a temporary build up of messages from starting too many MDB instances. By default, the maximum number of concurrent MDB instances is set to 10.

The Maximum concurrent endpoints field limits the number of endpoints (instances of a given message-driven bean) that process messages concurrently. If the maximum has been reached, new messages are not accepted from the messaging engine for delivery until an endpoint finishes its current processing.

If the available message count (queue depth) associated with a message-driven bean is frequently high, and if your server can handle more concurrent work, you can benefit from increasing the maximum concurrency setting.

If you set the maximum concurrency for a message-driven bean, be sure that you specify a value smaller than the maximum number of endpoint instances that can be created by the adapter that the message-driven bean is bound to. If necessary, increase the endpoint instance limit.

- An activation specification also has a **Maximum batch size** that refers to how many messages can be allocated to an endpoint in one batch for serial delivery. So, for example, if you have set the Maximum concurrent endpoints property to 10 and the Maximum batch Size property to 3, then there can be up to 10 endpoints each processing up to 3 messages giving a total of 30 messages allocated to that message-driven bean. If there are multiple message-driven beans deployed against a single activation specification then these maximum values apply to each message-driven bean individually.
- Take care to ensure that you always set the Maximum concurrent endpoints property is always less than the JCA pool size.

Note: You might want to tune the throttling of your message-driven beans, which is especially important on z/OS®. Workload arriving on the destination the message-driven bean is consuming from might use up more server resource and therefore obstruct other activities. An example of this is when restarting MDB applications you find a backlog of messages. The number of messages can be throttled so that the message-driven bean can process them in the most efficient manner.

To configure the message throttling support of the default messaging provider (the service integration bus JMS Resource Adapter), use the administrative console to complete the following steps.

Procedure

- Tune the maximum number of instances of a message-driven bean.

The maximum concurrency is set in the activation specification used to deploy the message-driven bean.

1. Click **Resources -> Resource Adapters -> J2C activation specifications -> *activation_specification_name* -> [Additional Properties] J2C activation specification custom properties.**
 2. View the maxConcurrency custom property. The default is value is 10. For high throughput primitive MDB tests, 40 was found to be an optimal value.
 3. Optional: To change the maxConcurrency setting, click the value field. This displays a panel for you to type a new value. In the Value field, type the new value then click **OK**. Save your changes to the master configuration.
- Tune the maximum batch size for a message-driven bean.

By default, only a single message is delivered to a message-driven bean instance at one time. You can improve performance by batching messages to the message-driven bean. Each message-driven bean

instance then receives a number (between 1 and the batch size) of messages at a time. A change in the maximum concurrency is likely to be beneficial if the available message count (queue depth) associated with the message-driven bean is frequently high. For more information about the available message count, see View the Available Message Count on a destination. The maximum batch size is set in the activation specification used to deploy the message-driven bean.

1. Click **Resources -> Resource Adapters -> J2C activation specifications -> *activation_specification_name* -> [Additional Properties] J2C activation specification custom properties.**
2. View the `maxBatchSize` custom property. The default value is 1. For high throughput primitive MDB tests, 5 was found to be optimal value (providing a 20 per cent gain over batch size 1).
3. Optional: To change the `maxBatchSize` setting, click the value field. This displays a panel for you to type a new value. In the Value field, type the new value then click **OK**. Save your changes to the master configuration.

Chapter 6. Tuning messaging destinations for the WebSphere MQ or V5 default messaging providers

Use this task to configure the properties of a messaging destination to optimize performance of applications that use the WebSphere MQ messaging provider or the V5 default messaging provider.

About this task

To optimize performance, configure destination properties to best fit your applications. For the V5 default messaging provider, you should also consider queue attributes of the JMS server that is associated with the queue name.

Procedure

- To optimize queue performance by configuring queue destination properties see “Performance for WebSphere MQ queues.”
- To configure the JMS queue destination properties to best fit your WebSphere Application Server Version 5 applications see “Performance for Version 5 default messaging queues” on page 36.

Performance for WebSphere MQ queues

To optimize performance, configure the queue destination properties to best fit your message-driven bean (MDB) or other applications that use the queue destinations.

For example:

- When MDB applications are configured to queues on WebSphere MQ for z/OS, the INDEX by MSGID is very important.
- Setting the Expiry property to SPECIFIED and the Specified Expiry property to 30000 milliseconds for the expiry timeout, reduces the number of messages that can be queued.

To ensure that there are enough underlying WebSphere MQ resources available for the queue, you must ensure that you configure the queue destination properties adequately for use by your message-driven beans or other applications that use the queue.

You must also consider the queue attributes associated with the queue name you created with WebSphere MQ. Inappropriate queue attributes can reduce the performance of WebSphere operations. You can use WebSphere MQ commands to change queue attributes for the queue name.

BOQNAME

The excessive backout requeue name. This attribute can be set to a local queue name that can hold the messages that were rolled back by the WebSphere applications. This queue name can be a system dead letter queue.

BOTHRESH

The backout threshold and can be set to a number when the threshold is reached, the message is moved to the queue name specified in BOQNAME.

INDXTYPE

Set this attribute to MSGID to cause an index of message identifiers to be maintained, which can improve WebSphere MQ retrieval of messages.

DEFSOPT

Set this attribute to SHARED (for shared input from the queue).

SHARE

This attribute must be specified (so that multiple applications can get messages from this queue).

For more information about using these properties, see the following sections of the WebSphere MQ information center:

- For BOQNAME and BOTHRESH, see “Handling poison messages” in the *Using Java* section
- *Script (MQSC) Command Reference*

Performance for Version 5 default messaging queues

To optimize performance, configure the JMS queue destination properties to best fit your WebSphere Application Server Version 5 applications.

For example, setting the Expiry property to SPECIFIED and the Specified Expiry property to 30000 milliseconds for the expiry timeout, reduces the number of messages that can be queued.

To ensure that there are enough underlying WebSphere MQ resources available for the queue, you must ensure that you configure the queue destination properties adequately for your application usage.

For queue destinations configured on a WebSphere Application Server Version 5 node, you should also consider queue attributes of the internal JMS server that are associated with the queue name.

Inappropriate queue attributes can reduce the performance of WebSphere operations.

BOQNAME

The excessive backout requeue name. This can be set to a local queue name that can hold the messages that were rolled back by the WebSphere applications. This queue name can be a system dead letter queue.

BOTHRESH

The backout threshold and can be set to a number once the threshold is reached, the message will be moved to the queue name specified in BOQNAME.

For more information about using these properties, see the *Handling poison messages in WebSphere MQ classes for JMS* and *Script (MQSC)* sections in the WebSphere MQ information center.

Chapter 7. Throttling inbound message flow for JCA 1.5 message-driven beans

This topic describes how to throttle message delivery for message-driven beans (MDB) which are deployed as message endpoints for Java Platform, Enterprise Edition (Java EE) Connector Architecture (JCA) Version 1.5 inbound resource adapters.

Before you begin

The throttling of messages as described in this topic does not apply to the two JCA 1.5-compliant messaging providers that are supplied with WebSphere Application Server:

- The default messaging provider.
- The WebSphere MQ messaging provider.

For the default messaging provider, you configure message throttling as described in the related tasks. For the WebSphere MQ messaging provider, you configure message throttling by setting the **maximum server sessions** property on the WebSphere MQ messaging provider activation specifications panel, or the **maxPoolSize** property when using the `createWMQActivationSpec` or `modifyWMQActivationSpec` `wsadmin` commands.

If you have a third-party JCA 1.5-compliant JMS messaging provider, check with your supplier to see whether the method of message throttling described in this topic is appropriate for their messaging provider.

About this task

For installations that use resource adapters that implement the Java EE Connector Architecture (JCA) Version 1.5 message delivery support, the WebSphere Application Server provides message throttling support to control the delivery of messages to endpoint message-driven beans (MDB). You can use this support to avoid overloading the server with a flood of inbound messages.

Message delivery is throttled on a message-driven bean basis by limiting the maximum number of endpoint instances that can be created by the adapter that the MDB is bound to. When the adapter attempts to create an endpoint instance, a proxy for the MDB instance is created and returned as allowed by the JCA 1.5 architecture. There is a one-to-one correspondence between proxies and MDB instances, and like the MDB instances, the proxies are pooled based on the minimum and maximum pool size values associated with the message-driven bean. Throttling is performed through the management of the proxy pool.

At the time the adapter attempts to create an endpoint, if the number of endpoint proxies currently created is equal to the maximum size of the pool, adapter `createEndPoint` processing returns an `UnavailableException`. When this displays, the adapter does not issue any more `createEndPoint()` requests until it has released at least one endpoint back to the server for reuse. Thus, installations can control the throttling of message delivery to a JCA 1.5 MDB based on the setting of the maximum size of the pool associated with each JCA 1.5 message-driven bean.

You can specify the pool size by using the `com.ibm.websphere.ejbcontainer.poolsize` JVM system property to define the minimum and maximum pool size of stateless, message-driven, and entity beans. For a message-driven bean that supports JCA 1.5, the maximum pool size value specified limits how many message endpoint instances can be created for that message-driven bean. For example, if the installation sets the maximum size of a JCA 1.5 MDB pool to 5, then at most 5 messages can be concurrently delivered to 5 instances of the message-driven bean. This property can be specified using the `wsadmin` scripting tool or by specifying it under the administrative console as an environmental variable.

Procedure

1. Open the administrative console.
2. Select **Servers > Server Types > WebSphere application server > *server_name***.
3. Under **Server Infrastructure**, expand **Java and Process Management > Process Definition**.
4. Under Additional Properties, select **Java Virtual Machine**.
5. Under Additional Properties, select **Custom Properties**.
6. Select **New**. A panel with three **General Properties** fields displays. This is where you set the property.
7. In the **Name** field, enter `com.ibm.websphere.ejbcontainer.poolsize`.
8. To complete the **Value** field, see the EJB container system properties for values.
9. After defining the value of the property, click **OK**. You are now prompted to save the changes you have made.
10. Click **Save**.

Chapter 8. Monitoring server session pools for listener ports

You can minimize the number of resources that server sessions use by enabling server session pool monitoring and defining the timeout value to be applied to a server session.

About this task

Each listener port uses one or more server sessions, which are held in a server session pool. Each server session is associated with a JMS session, which is taken from the JMS session pool that is associated with the JMS connection factory that the listener port is configured to use.

By default, server session pool monitoring is disabled. When a listener port uses a server session the listener port does not release the server session from the server session pool until the listener port is shut down. This means that the associated JMS session is not released into the JMS session pool until the listener port is shut down, even if the listener port is not processing any messages. Consequently the resources that the JMS session uses, for example TCP/IP connections, can be held for a long time, and this can cause problems for resource-constrained systems.

To minimize the number of resources that server sessions use, you must monitor the server session pools. When you enable server session pool monitoring each server session in each server session pool that a listener port uses is monitored to determine how much time has elapsed since the server session was last used. If the elapsed time is greater than the timeout value that you have configured, the server session is removed from the server session pool and its associated JMS session is returned to the JMS session pool. The returned JMS session can be either reused by another application or closed, depending on your JMS session pool settings. You can also configure additional pooling mechanisms, depending on your JMS provider.

Note: Server session pool monitoring cannot be used if the message listener service is operating in non-Application Server Facilities (non-ASF) mode, that is if the `NON.ASF.RECEIVE.TIMEOUT` message listener service custom property is set to a non-zero value.

Procedure

To enable server session pool monitoring, configure the following message listener service custom properties on each application server as required.

SERVER.SESSION.POOL.REAP.TIME

To enable server session pool monitoring, set this property to the time in seconds between checks on server session pools (this must be a non-negative value).

SERVER.SESSION.POOL.UNUSED.TIMEOUT

To specify the default server session pool timeout, set this property to the required number of seconds for the timeout. When this property is set to a non-negative value, it is compared with the time that has elapsed since a server session was used. If the timeout value is less than the elapsed time, the server session is removed from the server session pool and its JMS session is returned to the JMS session pool. For example, if the timeout value is one second and the time that has elapsed since a particular server session was used is two seconds, that server session is removed from the server session pool and its JMS session is returned to the JMS session pool.

SERVER.SESSION.POOL.UNUSED.TIMEOUT.*lpname*

To override the default `SERVER.SESSION.POOL.UNUSED.TIMEOUT` value for the listener port with the name *lpname*, set this property to the appropriate value:

- To override the `SERVER.SESSION.POOL.UNUSED.TIMEOUT` for the specified listener port, set this property to a non-negative value defining the required number of seconds for the server session timeout for this listener port.

- To disable server session pool monitoring for the specified listener port, set this property to a negative value.

The value that you set for this property applies to all message-driven beans that are using the specified listener port.

Example

For example, consider an application server that is configured with listener ports lp1, and lp2.

The following rules apply:

No properties set

If none of the properties are set, server session pool monitoring is disabled and JMS sessions used by server sessions are not returned to the JMS session pool until the listener port (lp1 or lp2), or its associated message-driven bean, is shut down.

SERVER.SESSION.POOL.REAP.TIME and SERVER.SESSION.POOL.UNUSED.TIMEOUT set

Consider, for example, the following settings:

```
SERVER.SESSION.POOL.REAP.TIME=60
SERVER.SESSION.POOL.UNUSED.TIMEOUT=120
```

The server session pool of both listener ports (lp1 and lp2) is checked for inactive server sessions every 60 seconds. If a server session is detected as being inactive for more than 120 seconds, it is removed from the server session pool and its JMS session is returned to the JMS session pool. Taking into account the SERVER.SESSION.POOL.REAP.TIME value, the server session pool could be removed from the session pool between two and three minutes after the server session was last used.

SERVER.SESSION.POOL.REAP.TIME and SERVER.SESSION.POOL.UNUSED.TIMEOUT set, and overrides set for SERVER.SESSION.POOL.UNUSED.TIMEOUT.*lpname*

Consider, for example, the following settings:

```
SERVER.SESSION.POOL.REAP.TIME=60
SERVER.SESSION.POOL.UNUSED.TIMEOUT=120
SERVER.SESSION.POOL.UNUSED.TIMEOUT.lp2=-1
SERVER.SESSION.POOL.UNUSED.TIMEOUT.lp1=60
```

The server session pool for listener port lp2 is not checked because it has a negative timeout value. In the server session pool for listener port lp1, any server sessions that are inactive for more than 60 seconds are removed from the server session pool.

Chapter 9. Tuning Object Request Broker (ORB)

This page provides a starting point for finding information about the Object Request Broker (ORB). The product uses an ORB to manage communication between client applications and server applications as well as among product components. These Java Platform, Enterprise Edition (Java EE) standard services are relevant to the ORB: Remote Method Invocation/Internet Inter-ORB Protocol (RMI/IIOP) and Java Interface Definition Language (Java IDL).

The ORB provides a framework for clients to locate objects in the network and call operations on those objects as though the remote objects were located in the same running process as the client, providing location transparency.

Tuning Object Request Brokers

Object Request Broker tuning guidelines

Use the guidelines in this document any time the Object Request Broker (ORB) is used in a workload.

The ORB is used whenever enterprise beans are accessed through a remote interface. If you experience particularly high or low CPU consumption, you might have a problem with the value of one of the following parameters. Examine these core tuning parameters for every application deployment.

Thread pool adjustments Size

Tune the size of the ORB thread pool according to your workload. Avoid suspending threads because they have no work ready to process. If threads do not have work ready to process, CPU time is consumed by calling the `Object.wait` method, performing a context switch. Tune the thread pool size such that the length of time that the threads wait is short enough to prevent them from being destroyed because they are idle too long.

The thread pool size is dependent on your workload and system. In typical configurations, applications need 10 or fewer threads per processor.

However, if your application is performing a very slow backend request, like a request to a database system, a server thread blocks waiting for the backend request to complete. With backend requests, CPU use is fairly low. In this case, increasing the load does not increase CPU use or throughput. Your thread dumps indicate that nearly all the threads are in a call out to the backend resource. In this case, consider increasing the number of threads per processor until throughput improves and thread dumps show that the threads are in other areas of the run time besides the backend call. You should adjust the number of threads only if your backend resource is tuned correctly.

The **Allow thread allocation beyond maximum thread size** parameter also affects thread pool size, but do not use this parameter unless your back end stops for long periods of time, causing the blocking of all the run-time threads waiting for the backend system instead of processing other work that does not involve the backend system.

You can adjust the thread pool size settings in the administrative console. Click **Servers > Server Types > Application servers > *server_name* > Container services > ORB service > Thread pool**. You can adjust the minimum and maximum number of threads.

Thread pool timeout

Each inbound and outbound request through the ORB requires a thread from the ORB thread pool. In heavy load scenarios or scenarios where ORB requests nest deeply, it is possible for a Java virtual machine (JVM) to have all threads from the ORB thread pool attempting to send requests. Meanwhile, the

remote JVM ORB that process these requests has all threads from its ORB thread pool attempting to send requests. As a result, progress is never made, threads are not released back to the ORB thread pool, and the ORB is unable to process requests. As a result, there is a potential deadlock. Using the administrative console, you can adjust this behavior through the ORB `com.ibm.websphere.orb.threadPoolTimeout` custom property. For more information, see the documentation about the Object Request Broker custom properties.

Fragment size

The ORB separates messages into fragments to send over the ORB connection. You can configure this fragment size through the `com.ibm.CORBA.FragmentSize` parameter.

To determine and change the size of the messages that transfer over the ORB and the number of required fragments, perform the following steps:

1. In the administrative console, enable ORB tracing in the ORB Properties page.
2. Enable ORBRas tracing from the logging and tracing page.
3. Increase the trace file sizes because tracing can generate a lot of data.
4. Restart the server and run at least one iteration (preferably several) of the case that you are measuring.
5. Look at the traceable file and do a search for `Fragment to follow: Yes`.

This message indicates that the ORB transmitted a fragment, but it still has at least one remaining fragment to send before the entire message arrives. A `Fragment to follow: No value` indicates that the particular fragment is the last in the entire message. This fragment can also be the first, if the message fit entirely into one fragment.

If you go to the spot where `Fragment to follow: Yes` is located, you find a block that looks similar to the following example:

```
Fragment to follow: Yes
Message size:      4988 (0x137C)
--
Request ID:       1411
```

This example indicates that the amount of data in the fragment is 4988 bytes and the Request ID is 1411. If you search for all occurrences of `Request ID: 1411`, you can see the number of fragments that are used to send that particular message. If you add all the associated message sizes, you have the total size of the message that is being sent through the ORB.

6. You can configure the fragment size by setting the `com.ibm.CORBA.FragmentSize` ORB custom property.

Interceptors

Interceptors are ORB extensions that can set up the context before the ORB runs a request. For example, the context might include transactions or activity sessions to import. If the client creates a transaction, and then flows the transaction context to the server, then the server imports the transaction context onto the server request through the interceptors.

Most clients do not start transactions or activity sessions, so most systems can benefit from removing the interceptors that are not required.

To remove the interceptors, manually edit the `server.xml` file and remove the interceptor lines that are not needed from the ORB section.

Connection Cache Adjustments

Depending on an application server's workload, and throughput or response-time requirements, you might need to adjust the size of the ORB's connection cache. Each entry in the connection cache is an object that represents a distinct TCP/IP socket endpoint, identified by the hostname or TCP/IP address, and the port number used by the ORB to send a GIOP request or a GIOP reply to the remote target endpoint. The

purpose of the connection cache is to minimize the time required to establish a connection by reusing ORB connection objects for subsequent requests or replies. (The same TCP/IP socket is used for the request and corresponding reply.)

For each application server, the number of entries in the connection cache relates directly to the number of concurrent ORB connections. These connections consist of both the inbound requests made from remote clients and outbound requests made by the application server. When the server-side ORB receives a connection request, it uses an existing connection from an entry in the cache, or establishes a new connection and adds an entry for that connection to the cache.

The ORB Connection cache maximum and Connection cache minimum properties are used to control the maximum and minimum number of entries in the connection cache at a given time. When the number of entries reaches the value specified for the Connection cache maximum property, and a new connection is needed, the ORB creates the requested connection, adds an entry to the cache and searches for and attempts to remove up to five inactive connection entries from the cache. Because the new connection is added before inactive entries are removed, it is possible for the number of cache entries to temporarily exceed the value specified for the Connection cache maximum property.

An ORB connection is considered inactive if the TCP/IP socket stream is not in use and there are no GIOP replies pending for any requests made on that connection. As the application workload diminishes, the ORB closes the connections and removes the entries for these connections from the cache. The ORB continues to remove entries from the cache until the number of remaining entries is at or below the value specified for the Connection cache maximum property. The number of cache entries is never less than the value specified for the Connection cache minimum property, which must be at least five connections less than the value specified for the Connection cache maximum property.

Adjustments to the connection cache in the client-side ORB are usually not necessary because only a small number of connections are made on that side.

JNI Reader Threads

By default, the ORB uses a Java thread for processing each inbound connection request it receives. As the number of concurrent requests increases, the storage consumed by a large number of reader threads increases and can become a bottleneck in resource-constrained environments. Eventually, the number of Java threads created can cause out-of-memory exceptions if the number of concurrent requests exceeds the system's available resources.

To help address this potential problem, you can configure the ORB to use JNI reader threads where a finite number of reader threads, implemented using native OS threads instead of Java threads, are created during ORB initialization. JNI reader threads rely on the native OS TCP/IP asynchronous mechanism that enables a single native OS thread to handle I/O events from multiple sockets at the same time. The ORB manages the use of the JNI reader threads and assigns one of the available threads to handle the connection request, using a round-robin algorithm. Ordinarily, JNI reader threads should only be configured when using Java threads is too memory-intensive for your application environment.

The number of JNI reader threads you should allocate for an ORB depends on many factors and varies significantly from one environment to another, depending on available system resources and workload requirements. The following potential benefits might be achieved if you use JNI threads:

- Because a fixed number of threads is allocated, memory usage is reduced. This reduction provides significant benefit in environments with unusually large and sustained client-request workloads.
- The time needed to dynamically create and destroy Java threads is eliminated because a fixed number of JNI threads is created and allocated during ORB initialization.
- Each JNI thread can handle up to 1024 socket connections and interacts directly with the asynchronous I/O native OS mechanism, which might provide enhanced performance of network I/O processing.

Chapter 10. Tuning Service integration

This page provides a starting point for finding information about service integration.

Service integration provides asynchronous messaging services. In asynchronous messaging, producing applications do not send messages directly to consuming applications. Instead, they send messages to destinations. Consuming applications receive messages from these destinations. A producing application can send a message and then continue processing without waiting until a consuming application receives the message. If necessary, the destination stores the message until the consuming application is ready to receive it.

Tuning messaging engines

Use this task to set tuning properties for the service integration environment.

About this task

The service integration environment includes properties that you can set to improve the performance of a messaging engine or the component of the messaging engine that manages the data store. These properties are known collectively as tuning properties. You can set these properties either with the WebSphere Application Server administrative console or by editing the `sib.properties` file.

Tip: Properties set with the administrative console take precedence over properties set in the `sib.properties` file.

Procedure

- Set tuning properties by using the administrative console:
 - Set the tuning properties of a messaging engine.
 - Control the memory buffers used by a messaging engine.
- Use the administrative console to tune the data source.
- Set tuning properties for any of the components mentioned above by editing the `sib.properties` file.

Setting tuning properties of a messaging engine

You can set the tuning properties for a messaging engine to improve its performance.

About this task

You can set the following tuning property for a messaging engine:

`sib.trm.retry`

The messaging engine to messaging engine connection retry interval, in seconds. The retry interval is the time delay left between attempts to contact neighboring messaging engines with which communications exist. The default retry interval is 30 seconds.

To set the tuning properties for a messaging engine, use the administrative console to complete the following steps.

Procedure

1. In the navigation pane, click **Service integration -> Buses -> *bus_name* -> [Topology] Messaging engines -> *engine_name* -> [Additional Properties] Custom properties**.
2. Type the name of the property that you want to set.
3. Type the value that you want to set for that property.
4. Click **OK**.

5. Save your changes to the master configuration.
6. Restart the messaging engine for the changes to take effect.

Controlling the memory buffers used by a messaging engine

Every messaging engine manages two memory buffers that contain messages and message-related data. You can improve the interaction of a messaging engine with its data store by tuning the properties that set the sizes of the two buffers.

About this task

You can set the following properties to improve the interaction of a messaging engine with its data store:

sib.msgstore.discardableDataBufferSize

The size in bytes of the data buffer that the messaging engine uses to contain data for which the quality of service attribute is best effort nonpersistent. The default value is 320000, which is approximately 320 kilobytes.

The discardable data buffer contains all data for which the quality of service attribute is best effort nonpersistent. That data comprises both data that is involved in active transactions, and any other best effort nonpersistent data that the messaging engine has neither discarded nor consumed. The messaging engine holds this data entirely within this memory buffer and never writes the data to the data store. When the messaging engine adds data to the discardable data buffer, for example when the messaging engine receives a best effort nonpersistent message from a client, the messaging engine might discard data already in the buffer to make space. The messaging engine can discard only data that is not involved in active transactions. This behavior enables the messaging engine to discard best effort nonpersistent messages.

Increasing the size of the discardable data buffer allows more best effort nonpersistent data to be handled before the messaging engine begins to discard messages.

If the messaging engine attempts to add data to the discardable data buffer when insufficient space remains after discarding all the data that is not involved in active transactions, the messaging engine throws a `com.ibm.ws.sib.msgstore.OutOfCacheSpace` exception. Client applications can catch this exception, wrapped inside API-specific exceptions such as `javax.jms.JMSEException`.

sib.msgstore.cachedDataBufferSize

The size in bytes of the data buffer that the messaging engine uses to contain data for which the quality of service attribute is *better than* best effort nonpersistent and that is held in the data store. The default value is 320000, which is approximately 320 kilobytes.

The purpose of the cached data buffer is to optimize the performance of the messaging engine by caching in memory the data that the messaging engine might otherwise have to read from the data store. As it writes data to the data store and reads from the data store, the messaging engine attempts to add that data to the cached data buffer. The messaging engine might discard data already in the buffer to make space.

sib.msgstore.transactionSendLimit

The maximum number of operations that the messaging engine includes in each transaction. For example, each JMS send or receive is an operation that counts towards the transaction send limit. The default value is 100.

Attention: The messaging engine uses approximate calculations to manage the data it holds in the memory buffers. Neither of the **DataBufferSize** properties gives an accurate indication of the amount of memory that the messaging engine consumes in the JVM heap. The messaging engine can consume considerably more heap storage than the **DataBufferSize** properties indicate.

To set the properties of a messaging engine to improve its interaction with its data store, use the administrative console to complete the following steps:

Procedure

1. In the navigation pane, click **Service integration -> Buses -> bus_name -> [Topology] Messaging engines -> engine_name -> [Additional Properties] Custom properties**.
2. Type the name of the property that you want to set.
3. Type the value that you want to set for that property.
4. Click **OK**.
5. Save your changes to the master configuration.

What to do next

Remember: When you change any of these properties, the new values do not take effect until you restart the messaging engine.

Tuning the JDBC data source of a messaging engine

The messaging engine needs to have the correct configuration for JDBC data source to achieve messaging performance on a service integration bus.

Before you begin

Consider whether you must configure the connection pool for the JDBC data source to achieve your requirements for messaging performance.

About this task

The messaging engine uses the connection pool to obtain its connections to the database. With a heavy workload, a messaging engine might require a large number of concurrent connections to avoid delays waiting for connections to become available in the pool. For example, a very heavily loaded messaging engine might need 50 or more connections. Complete the following steps to configure the connection pool to meet your performance requirements:

Procedure

1. Ensure that the configuration of your relational database management system (RDBMS) permits the number of connections that you require. Refer to the documentation for your RDBMS for more information.
2. Use the administrative console to set the connection pool parameters for your data source. Navigate to **Resources -> JDBC -> Data sources -> data_source_name -> [Additional Properties] Connection pool properties**.

- a. Set **Maximum connections** to the number of connections you require, for example, at least 50. The default number of connections is 10.

Tip: If your messaging engine times out when requesting a database connection, check the error log. If the error log contains error message CWSIS1522E, increase the number of connections and ensure that the configuration of your RDBMS permits that number of connections.

- b. Set **Purge policy** to `EntirePool`. This policy enables the connection pool to release all connections when the messaging engine stops.

Important: You must set this value if the messaging engine can failover in a cluster.

Setting tuning properties by editing the sib.properties file

Use this task to set tuning properties for the service integration environment by editing the `sib.properties` file

About this task

You can set the following tuning properties to improve the performance of components in the service integration environment.

Properties for a messaging engine

sib.trm.retry

The messaging engine to messaging engine connection retry interval, in seconds. The retry interval is the time delay left between attempts to contact neighboring messaging engines with which communications exist. The default retry interval is 30 seconds.

Properties for the component of a messaging engine that manages the data store

sib.msgstore.discardableDataBufferSize

The size in bytes of the data buffer that the messaging engine uses to contain data for which the quality of service attribute is best effort nonpersistent. The default value is 320000, which is approximately 320 kilobytes.

The discardable data buffer contains all data for which the quality of service attribute is best effort nonpersistent. That data comprises both data that is involved in active transactions, and any other best effort nonpersistent data that the messaging engine has neither discarded nor consumed. The messaging engine holds this data entirely within this memory buffer and never writes the data to the data store. When the messaging engine adds data to the discardable data buffer, for example when the messaging engine receives a best effort nonpersistent message from a client, the messaging engine might discard data already in the buffer to make space. The messaging engine can discard only data that is not involved in active transactions. This behavior enables the messaging engine to discard best effort nonpersistent messages.

Increasing the size of the discardable data buffer allows more best effort nonpersistent data to be handled before the messaging engine begins to discard messages.

If the messaging engine attempts to add data to the discardable data buffer when insufficient space remains after discarding all the data that is not involved in active transactions, the messaging engine throws a `com.ibm.ws.sib.msgstore.OutOfCacheSpace` exception. Client applications can catch this exception, wrapped inside API-specific exceptions such as `javax.jms.JMSEException`.

sib.msgstore.cachedDataBufferSize

The size in bytes of the data buffer that the messaging engine uses to contain data for which the quality of service attribute is *better than* best effort nonpersistent and that is held in the data store. The default value is 320000, which is approximately 320 kilobytes.

The purpose of the cached data buffer is to optimize the performance of the messaging engine by caching in memory the data that the messaging engine might otherwise have to read from the data store. As it writes data to the data store and reads from the data store, the messaging engine attempts to add that data to the cached data buffer. The messaging engine might discard data already in the buffer to make space.

sib.msgstore.transactionSendLimit

The maximum number of operations that the messaging engine includes in each transaction. For example, each JMS send or receive is an operation that counts towards the transaction send limit. The default value is 100.

To set these properties by editing the `sib.properties` file, complete the following steps:

Procedure

1. Navigate to the `profile_root/properties` directory, where `profile_root` is the directory in which profile-specific information is stored.

2. If the directory does not contain a `sib.properties` file, then copy the template `sib.properties` files from the `app_server_root/properties` directory, where `app_server_root` is the root directory for the installation of WebSphere Application Server.
3. Using a text editor, open the `sib.properties` file and add the name and value of the property that you want to set. The format is `name=value`. For example `sib.trm.retry=60`

Tuning messaging performance with service integration technologies

To help optimize performance, you can set tuning properties that control the performance of message-driven beans and other messaging applications deployed to use service integration technologies.

About this task

To optimize the performance of messaging with service integration technologies, you can use the administrative console to set various parameters as described in the steps below. You can also set these parameters by using the `wsadmin` tool.

Procedure

- View the Available Message Count on a destination.

Viewing the Available Message Count on a destination enables you to determine whether your message consumers are able to cope with your current workload. If the available message count on a given destination is too high, or is increasing over time, consider some of the tuning recommendations in this topic.

1. Enable `AvailableMessageCount` statistics for a queue. If you restart the administrative server, enable **AvailableMessageCount** statistics again because such runtime settings are not preserved when the server is restarted.
 - a. In the navigation pane, click **Monitoring and Tuning -> Performance Monitoring Infrastructure (PMI)**.
 - b. In the content pane, click **server_name**.
 - c. Click the Runtime tab.
 - d. In the Currently monitored statistic set, click **Custom**.
 - e. On the Custom monitoring level panel, click **SIB Service > SIB Messaging Engines > engine_name > Destinations > Queues > queue_name**.
 - f. Select the `AvailableMessageCount` option.
 - g. Click **Enable** at the top of the panel.
 2. View the available message count for a queue.
 - a. In the navigation pane, click **Monitoring and Tuning -> Performance Viewer -> Current activity**.
 - b. In the content pane, click **server_name**.
 - c. Click **Performance Modules > SIB Service > SIB Messaging Engines > engine_name > Destinations > Queues > queue_name**.
 - d. Click **View Module(s)** at the top of the Resource Selection panel, located on the left side. This displays the `AvailableMessageCount` data in the Data Monitoring panel, located on the right side. You can use the Data Monitoring panel to manage the collection of monitoring data; for example, you can use the buttons to start or stop logging, or to change the data displayed as either a table or graph.
- Monitor MDB Thread Pool Size for the Default Message Provider.

You might experience a performance bottleneck if there are insufficient threads available for the message-driven beans. There is a trade-off between providing sufficient threads to maximize the throughput of messages and configuring too many threads, which can lead to CPU starvation of the threads in the application server. If you notice that the throughput for express nonpersistent, reliable

nonpersistent, or reliable persistent messaging has fallen as a result of increasing the size of the default thread pool, then decrease the size of the thread pool and reassess the message throughput.

1. View or change the number of threads in the default thread pool for an application server. By default, message-driven beans use the default thread pool.
 - a. Click **Servers -> Server Types -> WebSphere application servers -> *server_name* -> [Additional Properties] Thread Pools > Default**. By default the Minimum size value is set to 5 and the Maximum size value is set to 20. The best performance is obtained by setting the Maximum size value to the expected maximum concurrency for all message-driven beans. For high throughput using a single message bean, 41 was found to be the optimal Maximum size value.
 - b. Change the Maximum size value, then click **OK**.
 2. Optional: Create your own thread pool. The default thread pool is also used by other WebSphere Application Server components, so you might want to define a separate thread pool for the message-driven beans. This reduces thread contention for the default thread pool.
 - a. Click **Servers -> Server Types -> WebSphere application servers -> *server_name* -> [Additional Properties] Thread Pools**.
 - b. Create a new thread pool.
 - c. Create sufficient threads to support the maximum amount of concurrent work for the message-driven beans.
 - d. Change the SIB JMS Resource Adapter to use the new thread pool:
 - 1) Click **Resources -> Resource Adapters -> Resource adapters**.
 - 2) If you cannot see any SIB JMS Resource Adapter instances in the list, expand **Preferences** and enable **Show built-in resources**.
 - 3) Select the **SIB JMS Resource Adapter** with the appropriate scope depending upon the scope of the connection factories.
 - 4) Add the name of the new thread pool in the **Thread pool alias** box.
 - 5) Click **Apply** .
 3. Save your changes to the master configuration.
- Tune MDB performance with the default messaging provider.
 1. Click **Resources -> JMS -> Activation specifications -> *activation_specification_name***.
 2. Set the maximum batch size for this activation specification.

Delivering batches of messages to each MDB endpoint can improve performance, particularly when used with Acknowledge mode set to `Duplicates-ok auto-acknowledge`. However, if message ordering must be retained across failed deliveries, set this parameter to 1.
 3. Set the maximum number of concurrent endpoints for this activation specification.

The maximum concurrent endpoints parameter controls the amount of concurrent work that can be processed by a message bean. The parameter is used with message-driven beans. Increasing the number of concurrent endpoints can improve performance but can increase the number of threads in use at one time. To benefit from a change in this parameter, there should be sufficient threads available in the MDB thread pool to support the concurrent work. However, if message ordering must be retained across failed deliveries, set this parameter to 1.
 4. Save your changes to the master configuration.

For additional information about tuning the throttling of message-driven beans, including controlling the maximum number of instances of each message bean and the message batch size for serial delivery, see “Configuring MDB throttling for the default messaging provider” on page 32.

- Ensure that application servers hosting one or more messaging engines are provided with an appropriate amount of memory for the message throughput you require.

You can tune the initial and maximum Java Virtual Machine (JVM) heap sizes when adding a server to a messaging bus, that is when you create a messaging engine. You have the option to do this in any of the following cases:

- When adding a single server to a bus
- When adding a cluster to a bus
- When adding a new server to an existing cluster that is itself a bus member

For an application server that is a bus member of at least one bus, or a member of a cluster that is a bus member of at least one bus, the recommended initial and maximum heap sizes are 768MB.

When you are adding a cluster to a bus, you are recommended to increase the initial and maximum JVM heap sizes for every server in the cluster to 768MB.

- Increasing the initial heap size improves the performance for small average message sizes
- Increasing the maximum heap size improves the performance for higher average message sizes
- Reduce the occurrence of `OutOfMemoryError` exceptions

If the cumulative size of the set of messages being processed within a transaction by the service integration bus is large enough to exhaust the JVM heap, `OutOfMemoryError` exceptions occur. Consider one of the following options for reducing the occurrence of `OutOfMemoryError` exceptions when processing a large set of messages within a transaction.

 - Increase the JVM heap sizes for the application server.
 - Reduce the cumulative size of the set of messages being processed within the transaction.
- Change the maximum connections in a connection factory for the default messaging provider.

The maximum connections parameter limits the number of local connections. The default is 10. This parameter should be set to a number equal to or greater than the number of threads (enterprise beans) concurrently sending messages.

 1. Click **Resources** -> **JMS** -> **Topic connection factories** -> *factory_name* > **[Additional Properties] Connection pool properties**.
 2. Enter the required value in the **Maximum connections** field.
 3. Click **Apply**.
 4. Save your changes to the master configuration.
- Tune reliability levels for messages.

The reliability level chosen for the messages has a significant impact on performance. In order of decreasing performance (fastest first), the reliability levels are:

Best effort nonpersistent

Express nonpersistent

Reliable nonpersistent

Reliable persistent

Assured persistent

For MDB point-to-point messaging, best effort nonpersistent throughput is more than six times greater than assured persistent. For more information about reliability levels, see *Message reliability levels - JMS delivery mode and service integration quality of service*.

Configuring MDB throttling for the default messaging provider

Use this task to configure the throttling of messages for message-driven beans that you have deployed as JCA 1.5 resources on the default messaging provider.

Before you begin

The throttling support described in this topic only applies to the default messaging provider (the service integration JMS Resource Adapter).

For the WebSphere MQ messaging provider, you configure message throttling by setting the **maximum server sessions** property on the WebSphere MQ messaging provider activation specifications panel, or the **maxPoolSize** property when using the `createWMQActivationSpec` or `modifyWMQActivationSpec` wsadmin commands.

If you have a third-party JCA 1.5-compliant JMS messaging provider, refer to the generic method of message throttling described in “Throttling inbound message flow for JCA 1.5 message-driven beans” on page 10.

About this task

Use this task if you want to throttle messages for a message-driven bean deployed as a J2EE Connector Architecture (JCA) 1.5 resource on the default messaging JMS provider.

The default messaging provider (the service integration JMS Resource Adapter) uses a special type of message throttling. You can leave the message-driven bean pools to the default size of 500.

The default messaging provider enables the throttling of message delivery to a message-driven bean through the **Maximum concurrent endpoints** configuration option on the JMS activation specification used to deploy the bean.

- The maximum number of instances of each message-driven bean is controlled by the Maximum concurrent endpoint setting in the activation specification used to deploy the message-driven bean. This maximum concurrency limit helps prevent a temporary build up of messages from starting too many MDB instances. By default, the maximum number of concurrent MDB instances is set to 10.

The Maximum concurrent endpoints field limits the number of endpoints (instances of a given message-driven bean) that process messages concurrently. If the maximum has been reached, new messages are not accepted from the messaging engine for delivery until an endpoint finishes its current processing.

If the available message count (queue depth) associated with a message-driven bean is frequently high, and if your server can handle more concurrent work, you can benefit from increasing the maximum concurrency setting.

If you set the maximum concurrency for a message-driven bean, be sure that you specify a value smaller than the maximum number of endpoint instances that can be created by the adapter that the message-driven bean is bound to. If necessary, increase the endpoint instance limit.

- An activation specification also has a **Maximum batch size** that refers to how many messages can be allocated to an endpoint in one batch for serial delivery. So, for example, if you have set the Maximum concurrent endpoints property to 10 and the Maximum batch Size property to 3, then there can be up to 10 endpoints each processing up to 3 messages giving a total of 30 messages allocated to that message-driven bean. If there are multiple message-driven beans deployed against a single activation specification then these maximum values apply to each message-driven bean individually.
- Take care to ensure that you always set the Maximum concurrent endpoints property is always less than the JCA pool size.

Note: You might want to tune the throttling of your message-driven beans, which is especially important on z/OS. Workload arriving on the destination the message-driven bean is consuming from might use up more server resource and therefore obstruct other activities. An example of this is when restarting MDB applications you find a backlog of messages. The number of messages can be throttled so that the message-driven bean can process them in the most efficient manner.

To configure the message throttling support of the default messaging provider (the service integration bus JMS Resource Adapter), use the administrative console to complete the following steps.

Procedure

- Tune the maximum number of instances of a message-driven bean.

The maximum concurrency is set in the activation specification used to deploy the message-driven bean.

1. Click **Resources -> Resource Adapters -> J2C activation specifications -> *activation_specification_name* -> [Additional Properties] J2C activation specification custom properties.**
 2. View the `maxConcurrency` custom property. The default value is 10. For high throughput primitive MDB tests, 40 was found to be an optimal value.
 3. Optional: To change the `maxConcurrency` setting, click the value field. This displays a panel for you to type a new value. In the Value field, type the new value then click **OK**. Save your changes to the master configuration.
- Tune the maximum batch size for a message-driven bean.

By default, only a single message is delivered to a message-driven bean instance at one time. You can improve performance by batching messages to the message-driven bean. Each message-driven bean instance then receives a number (between 1 and the batch size) of messages at a time. A change in the maximum concurrency is likely to be beneficial if the available message count (queue depth) associated with the message-driven bean is frequently high. For more information about the available message count, see [View the Available Message Count on a destination](#). The maximum batch size is set in the activation specification used to deploy the message-driven bean.

1. Click **Resources -> Resource Adapters -> J2C activation specifications -> *activation_specification_name* -> [Additional Properties] J2C activation specification custom properties.**
2. View the `maxBatchSize` custom property. The default value is 1. For high throughput primitive MDB tests, 5 was found to be optimal value (providing a 20 per cent gain over batch size 1).
3. Optional: To change the `maxBatchSize` setting, click the value field. This displays a panel for you to type a new value. In the Value field, type the new value then click **OK**. Save your changes to the master configuration.

Chapter 11. Tuning messaging engine data stores

Obtain an overview of improving the performance of messaging engine data stores.

About this task

- “Tuning the JDBC data source of a messaging engine” on page 47
- “Controlling the memory buffers used by a messaging engine” on page 46
- Sharing connections to benefit from one-phase commit optimization

Tuning the JDBC data source of a messaging engine

The messaging engine needs to have the correct configuration for JDBC data source to achieve messaging performance on a service integration bus.

Before you begin

Consider whether you must configure the connection pool for the JDBC data source to achieve your requirements for messaging performance.

About this task

The messaging engine uses the connection pool to obtain its connections to the database. With a heavy workload, a messaging engine might require a large number of concurrent connections to avoid delays waiting for connections to become available in the pool. For example, a very heavily loaded messaging engine might need 50 or more connections. Complete the following steps to configure the connection pool to meet your performance requirements:

Procedure

1. Ensure that the configuration of your relational database management system (RDBMS) permits the number of connections that you require. Refer to the documentation for your RDBMS for more information.
2. Use the administrative console to set the connection pool parameters for your data source. Navigate to **Resources -> JDBC -> Data sources -> *data_source_name* -> [Additional Properties] Connection pool properties**.
 - a. Set **Maximum connections** to the number of connections you require, for example, at least 50. The default number of connections is 10.

Tip: If your messaging engine times out when requesting a database connection, check the error log. If the error log contains error message CWSIS1522E, increase the number of connections and ensure that the configuration of your RDBMS permits that number of connections.

- b. Set **Purge policy** to `EntirePool`. This policy enables the connection pool to release all connections when the messaging engine stops.

Important: You must set this value if the messaging engine can failover in a cluster.

Controlling the memory buffers used by a messaging engine

Every messaging engine manages two memory buffers that contain messages and message-related data. You can improve the interaction of a messaging engine with its data store by tuning the properties that set the sizes of the two buffers.

About this task

You can set the following properties to improve the interaction of a messaging engine with its data store:

sib.msgstore.discardableDataBufferSize

The size in bytes of the data buffer that the messaging engine uses to contain data for which the quality of service attribute is best effort nonpersistent. The default value is 320000, which is approximately 320 kilobytes.

The discardable data buffer contains all data for which the quality of service attribute is best effort nonpersistent. That data comprises both data that is involved in active transactions, and any other best effort nonpersistent data that the messaging engine has neither discarded nor consumed. The messaging engine holds this data entirely within this memory buffer and never writes the data to the data store. When the messaging engine adds data to the discardable data buffer, for example when the messaging engine receives a best effort nonpersistent message from a client, the messaging engine might discard data already in the buffer to make space. The messaging engine can discard only data that is not involved in active transactions. This behavior enables the messaging engine to discard best effort nonpersistent messages.

Increasing the size of the discardable data buffer allows more best effort nonpersistent data to be handled before the messaging engine begins to discard messages.

If the messaging engine attempts to add data to the discardable data buffer when insufficient space remains after discarding all the data that is not involved in active transactions, the messaging engine throws a `com.ibm.ws.sib.msgstore.OutOfCacheSpace` exception. Client applications can catch this exception, wrapped inside API-specific exceptions such as `javax.jms.JMSEException`.

sib.msgstore.cachedDataBufferSize

The size in bytes of the data buffer that the messaging engine uses to contain data for which the quality of service attribute is *better than* best effort nonpersistent and that is held in the data store. The default value is 320000, which is approximately 320 kilobytes.

The purpose of the cached data buffer is to optimize the performance of the messaging engine by caching in memory the data that the messaging engine might otherwise have to read from the data store. As it writes data to the data store and reads from the data store, the messaging engine attempts to add that data to the cached data buffer. The messaging engine might discard data already in the buffer to make space.

sib.msgstore.transactionSendLimit

The maximum number of operations that the messaging engine includes in each transaction. For example, each JMS send or receive is an operation that counts towards the transaction send limit. The default value is 100.

Attention: The messaging engine uses approximate calculations to manage the data it holds in the memory buffers. Neither of the **DataBufferSize** properties gives an accurate indication of the amount of memory that the messaging engine consumes in the JVM heap. The messaging engine can consume considerably more heap storage than the **DataBufferSize** properties indicate.

To set the properties of a messaging engine to improve its interaction with its data store, use the administrative console to complete the following steps:

Procedure

1. In the navigation pane, click **Service integration -> Buses -> bus_name -> [Topology] Messaging engines -> engine_name -> [Additional Properties] Custom properties**.
2. Type the name of the property that you want to set.
3. Type the value that you want to set for that property.
4. Click **OK**.
5. Save your changes to the master configuration.

What to do next

Remember: When you change any of these properties, the new values do not take effect until you restart the messaging engine.

Increasing the number of data store tables to relieve concurrency bottleneck

Service integration technologies enables users to spread the data store for a messaging engine across several tables. In typical use this is unlikely to have a significant influence. However, if statistics suggest a concurrency bottleneck on the *SIBnnn* tables for a data store, you might try to solve the problem by increasing the number of tables.

About this task

For more information about the set of tables in a data store see [Data store tables](#).

SIB000	contains information about the structure of the data in the other two tables - the "stream table"
SIB001	contains persistent objects - the "permanent item table"
SIB002	contains nonpersistent objects that have been saved to the data store to reduce the messaging engine memory requirement - the "temporary item table"

Having multiple tables means you can relieve any performance bottleneck you might have in your system. You can modify *SIBnnn* tables of the data store of a messaging engine. You can increase the number of permanent and temporary tables (*SIB001* and *SIB002*), although there is no way to increase the number of stream tables (*SIB000*).

Example

This example illustrates what the *SIBnnn* tables for a data store might look like after modification:

SIB000	contains information about the structure of the data in the other two tables - the "stream table"
SIB001	contains persistent objects - the "permanent item table"
SIB002	contains persistent objects - the "permanent item table"
SIB003	contains persistent objects - the "permanent item table"
SIB004	contains nonpersistent objects that have been saved to the data store to reduce the messaging engine memory requirement - the "temporary item table"
SIB005	contains nonpersistent objects that have been saved to the data store to reduce the messaging engine memory requirement - the "temporary item table"

Increasing the number of item tables for a messaging engine when tables are not automatically created

If a concurrency bottleneck occurs on the item tables, increasing the number of item tables will increase the throughput of the messaging engine.

Before you begin

Before performing this task you must ensure that the messaging engine is using a data store, and that its **Create tables** option is set to `False`.

Procedure

1. Relevant performance monitoring tools show that the throughput of a messaging engine is inefficient.
2. Use your database performance monitoring tools to examine lock statistics for the item tables for evidence of a bottleneck. Consult database documentation on how to interpret the locking statistics.
3. Create tables and increase data store attributes.
 - Create tables for the data store schema. For more information see [Creating data store tables](#)
 - Increase number of permanent tables or temporary tables, or both, for the data store. For more information see [Configuring a messaging engine data store to use a data source](#)

You can only increase the number of permanent tables or temporary tables, not decrease them.

4. Stop and restart the WebSphere Application Server so that configuration changes take effect. The extra tables are used when the messaging engine restarts.
5. Observe the effect on throughput and lock statistics by checking performance monitoring tools. Consider whether any improvement is sufficient and modifying the data store attributes further would be beneficial

Increasing the number of item tables for a messaging engine when tables are automatically created

If a concurrency bottleneck occurs on the item tables, increasing the number of items tables will increase the throughput of the messaging engine.

Before you begin

Before performing this task you must ensure that the messaging engine is using a data store, and that its **Create tables** option is set to `True`.

Procedure

1. Relevant performance monitoring tools show that the throughput of a messaging engine is insufficient.
2. Use your database performance monitoring tools to examine lock statistics for the item tables for evidence of a bottleneck. Consult your database documentation on how to interpret the locking statistics.
3. Increase the attributes for the messaging engine data store: the number of permanent tables or temporary tables, or both. For more information see [Configuring a messaging engine data store to use a data source](#). You can only increase the number of permanent tables or temporary tables, not decrease them.
4. Stop and restart the messaging engine so that configuration changes take effect. The extra tables are created when the messaging engine starts again.
5. Observe the effect on throughput and lock statistics by checking performance monitoring tools. Consider whether any improvement is sufficient and modifying the data store attributes further would be beneficial

Tuning one-phase commit optimization

If you have configured your messaging engine to use a data store, you can achieve better performance by configuring both the messaging engine and container-managed persistent (CMP) beans to share the same data source.

About this task

You must configure both the CMP bean and the messaging engine resource authorization so that they share the same data source.

Procedure

1. Open the administrative console.
2. Click **Applications -> Application Types -> WebSphere enterprise applications -> *application_name* -> [Enterprise Java Bean Properties] Map data sources for all 2.x CMP beans.**
3. On the content pane, select the check boxes next to all the CMP beans.
4. Select *Per application* in the **Resource authorization** drop-down list.
5. Modify the messaging engine resource authorization to *Per application* by modifying the property file `sib.properties` and adding the custom property `sib.msgstore.jdbcResAuthForConnections=Application`.

Tuning the detection of database connection loss

If a messaging engine is configured to use a data store and cannot connect to its data store, for example because the database that contains the data store is not running, the messaging engine does not start. You can tune your system to increase the chance of a successful start of the messaging engine.

About this task

In a single-server environment, when you start the application server the messaging engine attempts to start. If the database is unavailable for more than 15 minutes, the messaging engine might enter the stopped state and need to be started manually.

In a high availability environment, a messaging engine starts either as part of the server or cluster startup, or as part of the failover process. During messaging engine startup, the messaging engine attempts to connect to the data store, for up to 15 minutes by default. If one of the following statements remains true during that time, the messaging engine cannot start on the server, and the server is disabled for high availability:

- The database is unavailable or not running.
- In a failover situation, the database does not detect the loss of the network connection to the original application server, and therefore does not release the locks on the data store.

This disabled state can propagate to all members of the cluster. You must manually re-enable the servers to maintain your high availability environment.

You can increase the chance of the messaging engine starting successfully by configuring various parameters, such as the 15 minute default timeout, on the database server or application server.

Procedure

1. On the database server, configure the operating system to minimize the amount of time taken to detect the loss of a network connection to an application server. Refer to the documentation for the operating system for details. For example, the following table lists the relevant parameters for Windows and AIX® operating systems:

Table 3. TCP/IP parameters. The first column of the table provides the list of TCP/IP parameters for the Windows operating systems. The second column of the table provides the list of TCP/IP parameters for the AIX operating systems. The third column provides the description of the parameters.

Parameter name on Windows operating systems	Parameter name on AIX operating systems	Description
KeepAliveTime	tcp_keepidle	The amount of time (in milliseconds on Windows operating systems and in 0.5 seconds on AIX operating systems) to wait before sending a keepalive request for an inactive connection.
KeepAliveInterval	tcp_keepintvl	The amount of time (in milliseconds on Windows operating systems and in 0.5 seconds on AIX operating systems) to wait for a response.
TCPMaxDataRetransmissions	tcp_keepcnt	The number of requests to send before ending the connection.

You can calculate the total amount of time taken for the database server to detect the failure of the connection to the application server, by using the following formula:

time to detect connection failure = keep alive time + (keep alive interval x number of requests)

For example, for a Windows system with the parameters set according to the following table, the total amount of time taken for the database server to detect the failure of the connection to the application server is 350 seconds.

Table 4. Example parameter values. The first column provides the parameter names. The second column provides a sample value for the parameters.

Parameter	Value
KeepAlive	300000 milliseconds
KeepAliveInterval	10000 milliseconds
TCPMaxDataRetransmissions	5

Your database product might also have relevant parameters that you can configure, for example, the IDLE THREAD TIMEOUT parameter in DB2 for z/OS.

When the database server detects the loss of the connection to the application server, the database releases the locks on the data store. The messaging engine can now access the data store and can therefore start successfully.

2. On the application server, tune the messaging engine to wait for an appropriate amount of time for the data store to become available. By default, the messaging engine will attempt to connect to the data store every 2 seconds for 15 minutes. Complete the rest of this step if you want to adjust these timings.
 - a. Click **Service integration -> Buses -> bus_name -> [Topology] Messaging engines -> engine_name -> [Additional Properties] Custom properties** to navigate to the custom properties panel for the messaging engine.
 - b. Click **New**.
 - c. Type `sib.msgstore.jdbcInitialDataSourceWaitTimeout` in the Name field and an appropriate value in the Value field. This property is the time, in milliseconds, to wait for the data store to become available. The default value is 900000 (15 minutes). This time includes the time required to establish a connection to the database and to obtain the required table locks.
Ensure that the value of this property is greater than the total time taken for the database server to detect the loss of a network connection, as configured in step 1.
 - d. Click **OK**.
 - e. Click **New**.
 - f. Type `sib.msgstore.jdbcStaleConnectionRetryDelay` in the Name field and an appropriate value in the Value field. This property is the time, in milliseconds, to wait between attempts to connect to the

data store. The default value is 2000 (2 seconds). For example, if you set the `sib.msgstore.jdbcInitialDatasourceWaitTimeout` property to 600000, and the `sib.msgstore.jdbcStaleConnectionRetryDelay` property to 3000, the messaging engine will attempt to connect every 3 seconds until 10 minutes has passed.

- g. Click **OK**.
- h. Save your changes to the master configuration.
- i. Restart the application server.
- j. If you have a cluster, repeat the previous steps to add these properties for every messaging engine in the cluster.

Results

By configuring these parameters and custom properties, you minimize the amount of time taken for the database server to detect the loss of a network connection, and ensure that the messaging engine waits for a reasonable amount of time for the database connection to recover before attempting to start.

What to do next

You might want to configure the messaging engine and server to restart in the event of a database connection failure. This behavior reduces the risk of the messaging engine being in an inconsistent state when the database connection is restored.

Chapter 12. Setting tuning properties for a mediation

Use this task to tune a mediation for performance by using the administrative console.

Before you begin

Review the guidance on when it is appropriate to tune a mediation for performance in the topic Performance tuning for mediations.

About this task

You can set the following tuning property in the administrative console to improve the performance of a mediation:

sib:SkipWellFormedCheck

Whether you want to omit the well formed check that is performed on messages after they have been processed by the mediation. Either true or false.

Note: This property is overridden for messages that have the delivery option assured persistent, and a well formed check is always performed.

To set, or unset, one or more tuning properties for a mediation, use the administrative console to complete the following steps:

Procedure

1. Display the mediation context information:
 - a. Click **Service integration -> Buses -> bus_name -> [Destination resources] Mediations**.
 - b. In the content pane, select the name of the mediation for which you want to configure tuning information.
 - c. Click **[Additional Properties] Context information**.
2. In the content pane, click **New**.
3. Type the name of the property in the **Name** field.
4. Select the type **Boolean** from the drop-down list.
5. Type **true** in the **Context Value** field to set the property, or type **false** to unset the property.
6. Click **OK**.
7. Save your changes to the master configuration.

Chapter 13. Enabling CMP entity beans and messaging engine data stores to share database connections

Use this task to enable container-managed persistence (CMP) entity beans to share the database connections used by the data store of a messaging engine. Performing this task has been estimated to provide a potential performance improvement of 15% for overall message throughput, but can only be used for entity beans connected to the application server that contains the messaging engine.

About this task

To enable CMP entity beans to share the database connections used by the data store of a messaging engine, complete the following steps.

Procedure

1. Configure the data store to use a data source that is not XA-capable. For more information about configuring a data store, see *Configuring a JDBC data source for a messaging engine*.
2. Select the **Share data source with CMP** option. This option is provided on the JMS connection factory or JMS activation specification used to connect to the service integration bus that hosts the bus destination that is used to store and process messages for the CMP bean.

For example, to select the option on a unified JMS connection factory, complete the following steps:

- a. Display the default messaging provider. In the navigation pane, click **Resources -> JMS -> JMS providers**.
- b. Select the default provider for which you want to configure a unified connection factory.
- c. Optional: Change the **Scope** check box to set the level at which the connection factory is to be visible, according to your needs.
- d. In the content pane, under Additional Properties, click **Connection factories**.
- e. Optional: To create a new unified JMS connection factory, click **New**.

Specify the following properties for the connection factory:

Name Type the name by which the connection factory is known for administrative purposes.

JNDI name

Type the JNDI name that is used to bind the connection factory into the namespace.

Bus name

Type the name of the service integration bus that the connection factory is to create connections to. This service integration bus hosts the destinations that the JMS queues and topics are assigned to.

- f. Optional: To change the properties of an existing connection factory, select its name from one of the connection factories displayed. The properties for the connection factory are displayed in the content pane.
- g. Select the check box for the **Share data source with CMP** field.
- h. Click **OK**.
- i. Save your changes to the master configuration.

The JMS connection factory can only be used to connect to a “local” messaging engine that is in the application server on which the CMP beans are deployed.

3. Deploy the CMP beans onto the application server that contains the messaging engine, and specify the same data source as that used by the messaging engine. You can use the administrative console to complete the following steps:
 - a. Optional: To determine the data source used by the messaging engine, click **Servers -> Server Types -> WebSphere application servers -> server_name -> [Server messaging] Messaging engines -> engine_name -> [Additional Properties] Message store**.

The **Data source name** field displays the name of the data source, which is by default:

`jdbc/com.ibm.ws.sib/engine_name`

- b. Click **Applications -> New Application -> New Enterprise Application**.
- c. On the first Preparing for the application installation page, specify the full path name of the source application file (.ear file, otherwise known as an EAR file), then click **Next**.
- d. On the second Preparing for the application installation page, complete the following steps:
 - 1) Select the check box for **Generate Default Bindings**. Data source bindings (for EJB 1.1 JAR files) are generated based on the JNDI name, data source, user name, and password options. This results in default data source settings for each EJB JAR file. No bean-level data source bindings are generated.
 - 2) Under Connection Factory Bindings, select the check box for **Default connection factory bindings**;, then type the JNDI name for the data source and optionally select a **Resource authorization** value.
 - 3) Click **Next** to display the Install New Application pages. The contents of the application that you are installing determines which pages are available.
4. If your application uses EJB modules that contain CMP beans that are based on the EJB 1.x specification, for **Map default data sources for modules containing 1.x entity beans**, specify a JNDI name for the default data source for the EJB modules. The default data source for the EJB modules is optional if data sources are specified for individual CMP beans.
5. If your application has CMP beans that are based on the EJB 1.x specification, for **Map data sources for all 1.x CMP**, specify a JNDI name for data sources to be used for each of the 1.x CMP beans. The data source attribute is optional for individual CMP beans if a default data source is specified for the EJB module that contains CMP beans.
6. Click **Finish**. If neither a default data source for the EJB module nor a data source for individual CMP beans are specified, a validation error displays and the installation is cancelled.
7. Complete other pages as needed.
8. On the Summary page, verify the cell, node, and server onto which the application modules will install.
 - a. Beside **Cell/Node/Server**, click **Click here**.
 - b. Verify the settings on the Map modules to servers page that is displayed. Ensure that the application server that is specified contains the messaging engine and its data store.
 - c. Specify the web servers as targets that will serve as routers for requests to this application. This information is used to generate the plug-in configuration file (plugin-cfg.xml) for each web server.
 - d. Return to the Summary page.
 - e. Click **Finish**.

Results

For more information about installing applications, see Installing enterprise application files with the console.

Chapter 14. Tuning security

This page provides a starting point for finding information about how to maintain, improve and harden your security configurations.

Tuning, hardening, and maintaining security configurations

After installing WebSphere Application Server, there are several considerations for tuning, strengthening, and maintaining your security configuration.

About this task

The following topics are covered in this section:

Procedure

- **Tuning security configurations** You can tune your security configuration to balance performance with function. You can achieve this balance following considerations for tuning general security, Common Secure Interoperability version 2 (CSIv2), Lightweight Directory Access Protocol (LDAP) authentication, web authentication, and authorization. For more information on tuning security, see “Tuning security configurations.”
- **Hardening security configurations** Several methods exist that you can use to protect your infrastructure and applications from different forms of attack. For more information on hardening your security, see “Hardening security configurations” on page 73.
- **Securing passwords in files** Password encryption and encoding can add protection to passwords existing in files. For more information on encoding and encrypting passwords, see “Securing passwords in files” on page 76.

What to do next

For additional information about hardening security configurations, see the WebSphere Application Server security web page.

Tuning security configurations

You can tune security to balance performance with function. You can achieve this balance following considerations for tuning general security, Common Secure Interoperability version 2 (CSIv2), Lightweight Directory Access Protocol (LDAP) authentication, web authentication, and authorization.

About this task

Performance issues typically involve trade-offs between function and speed. Usually, the more function and the more processing that are involved, the slower the performance. Consider what type of security is necessary and what you can disable in your environment. For example, if your application servers are running in a Virtual Private Network (VPN), consider whether you can disable Secure Sockets Layer (SSL). If you have a lot of users, can they be mapped to groups and then associated to your Java Platform, Enterprise Edition (Java EE) roles? These questions are things to consider when designing your security infrastructure.

Procedure

- Consider the following recommendations for tuning general security.
 - Consider disabling Java 2 security manager if you know exactly what code is put onto your server and you do not need to protect process resources. Remember that in doing so, you put your local resources at some risk.

- Consider propagating new security settings to all nodes before restarting the deployment manager and node agents, to change the new security policy.

If your security configurations are not consistent across all servers, you get access denied errors. Therefore, you must propagate new security settings when enabling or disabling administrative security.

Configuration changes are generally propagated using configuration synchronization. If auto-synchronization is enabled, you can wait for the automatic synchronization interval to pass, or you can force synchronization before the synchronization interval expires. If you are using manual synchronization, you must synchronize all the nodes.

If the cell is in a configuration state and the security policy is mixed with nodes that have security enabled and disabled, you can use the syncNode utility to synchronize the nodes where the new settings are not propagated.

For more detailed information about enabling security in a distributed environment, see Enabling security for the realm.

- Consider increasing the cache and token timeout if you feel your environment is secure enough. By increasing these values, you have to re-authenticate less often. This action supports subsequent requests to reuse the credentials that already are created. The downside of increasing the token timeout is the exposure of having a token hacked and providing the hacker more time to hack into the system before the token expires. You can use security cache properties to determine the initial size of the primary and secondary hashtable caches, which affect the frequency of rehashing and the distribution of the hash algorithms.

See the article Authentication cache settings for a list of these properties.

- Consider changing your administrative connector from Simple Object Access Protocol (SOAP) to Remote Method Invocation (RMI) because RMI uses stateful connections while SOAP is completely stateless. Run a benchmark to determine if the performance is improved in your environment.
- Use the wsadmin script to complete the access IDs for all the users and groups to speed up the application startup. Complete this action if applications contain many users or groups, or if applications are stopped and started frequently. WebSphere Application Server maps user and group names to unique access IDs in the authorization table. The exact format of the access ID depends on the repository. The access ID can only be determined during and after application deployment. Authorization tables created during assembly time do not have the proper access IDs. See the Commands for the AdminApp article for more information about how to update access IDs.
- Consider tuning the Object Request Broker (ORB) because it is a factor in enterprise bean performance with or without security enabled. Refer to the information about ORB tuning guidelines.
- If using SSL, enable the SSL session tracking mechanism option as described in the information about session management settings.
- In some cases, using the unrestricted Java Cryptography Extension (JCE) policy file can improve performance. Refer to the information about tuning Web Services Security.
- Distributing the workload to multiple Java virtual machines (JVMs) instead of a single JVM on a single machine can improve the security performance because there is less contention for authorization decisions.
- Consider the following steps to tune Common Secure Interoperability version 2 (CSlv2).
 - Consider using Secure Sockets Layer (SSL) client certificates instead of a user ID and password to authenticate Java clients. Because you are already making the SSL connection, using mutual authentication adds little overhead while it removes the service context that contains the user ID and password completely.
 - If you send a large amount of data that is not very security sensitive, reduce the strength of your ciphers. The more data you have to bulk encrypt and the stronger the cipher, the longer this action takes. If the data is not sensitive, do not waste your processing with 128-bit ciphers.
 - Consider putting only an asterisk (*) in the trusted server ID list (meaning trust all servers) when you use identity assertion for downstream delegation. Use SSL mutual authentication between servers to provide this trust. Adding this extra step in the SSL handshake performs better than having to fully

authenticate the upstream server and check the trusted list. When an asterisk (*) is used, the identity token is trusted. The SSL connection trusts the server through client certificate authentication.

- Ensure that stateful sessions are enabled for CSiv2. This is the default, but requires authentication only on the first request and on any subsequent token expirations.
- Consider changing the values for the CSiv2 session cache. Changing these values can avoid resource shortages. Refer to the Common Secure Interoperability Version 2 outbound communications topic for more information.
- If you are communicating only with WebSphere Application Server Version 5 or higher servers, make the Active Authentication Protocol CSI, instead of CSI and SAS. This action removes an interceptor invocation for every request on both the client and server sides.

Important: SAS is supported only between Version 6.0.x and previous version servers that have been federated in a Version 6.1 cell.

- Consider the following steps to tune Lightweight Directory Access Protocol (LDAP) authentication.
 1. In the administration console, click **Security > Global security**.
 2. Under User account repository, click the **Available realm definitions** drop-down list, select **Standalone LDAP registry** and click **Configure**.
 3. Select the **Ignore case for authorization** option in the stand-alone LDAP registry configuration, when case-sensitivity is not important.
 4. Select the **Reuse connection** option.
 5. Use the cache features that your LDAP server supports.
 6. Choose either the IBM Tivoli Directory Server or SecureWay™ directory type, if you are using an IBM Tivoli Directory Server. The IBM Tivoli Directory Server yields improved performance because it is programmed to use the new group membership attributes to improve group membership searches. However, authorization must be case insensitive to use IBM Tivoli Directory Server.
 7. Choose either iPlanet Directory Server (also known as Sun ONE) or Netscape as the directory if you are an iPlanet Directory user. Using the iPlanet Directory Server directory can increase performance in group membership lookup. However, use **Role** only for group mechanisms.
- Consider the following steps to tune web authentication.
 - Increase the cache and token timeout values if you feel your environment is secure enough. The web authentication information is stored in these caches and as long as the authentication information is in the cache, the login module is not invoked to authenticate the user. This supports subsequent requests to reuse the credentials that are already created. A disadvantage of increasing the token timeout is the exposure of having a token stolen and providing the thief more time to hack into the system before the token expires.
 - Enable single sign-on (SSO). To configure SSO, click **Security > Global security**. Under Web security, click **Single sign-on (SSO)**.

SSO is only available when you configure **LTPA** as the authentication mechanism in the Authentication mechanisms and expiration panel. Although you can select Simple WebSphere Authentication Mechanism (SWAM) as the authentication mechanism on the Authentication mechanisms and expiration panel, SWAM is deprecated in Version 8.0 and does not support SSO. When you select SSO, a single authentication to one application server is enough to make requests to multiple application servers in the same SSO domain. Some situations exist where SSO is not a desirable and you do not want to use it in those situations.
 - Disable or enable the **Web Inbound Security Attribute Propagation** option on the Single sign-on (SSO) panel if the function is not required. In some cases, having the function enabled can improve performance. This improvement is most likely for higher volume cases where a considerable number of user registry calls reduces performance. In other cases, having the feature disabled can improve performance. This improvement is most likely when the user registry calls do not take considerable resources.
 - The following two custom properties might help to improve performance when security attribute propagation is enabled:

- **com.ibm.CSI.propagateFirstCallerOnly**
The default value of this property is true. When this custom property is set to true the first caller in the propagation token that stays on the thread is logged when security attribute propagation is enabled. When this property is set to false, all of the caller switches are logged, which can affect performance.
- **com.ibm.CSI.disablePropagationCallerList**
When this custom property is set to true the ability to add a caller or host list in the propagation token is completely disabled. This function is beneficial when the caller or host list in the propagation token is not needed in the environment.
- Consider the following steps to tune authorization.
 - Map your users to groups in the user registry. Associate the groups with your Java Platform, Enterprise Edition (Java EE) roles. This association greatly improves performance when the number of users increases.
 - Judiciously assign method-permissions for enterprise beans. For example, you can use an asterisk (*) to indicate all the methods in the method-name element. When all the methods in enterprise beans require the same permission, use an asterisk (*) for the method-name to indicate all methods. This indication reduces the size of deployment descriptors and reduces the memory that is required to load the deployment descriptor. It also reduces the search time during method-permission match for the enterprise beans method.
 - Judiciously assign security-constraints for servlets. For example, you can use the *.jsp URL pattern to apply the same authentication data constraints to indicate all JavaServer Pages (JSP) files. For a given URL, the exact match in the deployment descriptor takes precedence over the longest path match. Use the *.jsp, *.do, *.html extension match if no exact matches exist and longest path matches exist for a given URL in the security constraints.
- Use new tuning parameters when using Java 2 security. The new tuning parameters can improve performance significantly, and introduce a new concept called *Read-only Subject*, which enables a new cache for J2C Auth Subjects when using container-managed auth data aliases. If the J2C auth subject does not need to be modified after it is created, the following new tuning parameters can be used to improve Java 2 Security performance:
 - com.ibm.websphere.security.auth.j2c.cacheReadOnlyAuthDataSubjects=true
 - com.ibm.websphere.security.auth.j2c.readOnlyAuthDataSubjectCacheSize=50 (This is the maximum number of subjects in the hashtable of the cache. Once the cache reaches this size, some of the entries are purged. For better performance, this size should be equal to the number of unique subjects (cache based on uniqueness of user principal + auth data alias + managed connection factory instance) when role-based security and Java 2 security are used together).
- Use new tuning parameters to improve the performance of Security Attribute Propagation. The new tuning parameters can be set through custom properties in the administrative console to reduce the extra overhead of Security Attribute Propagation:
 - com.ibm.CSI.disablePropagationCallerList=true
 - com.ibm.CSI.propagateFirstCallerOnly=true (use if you want to track the first caller only).

Results

You always have a trade off between performance, feature, and security. Security typically adds more processing time to your requests, but for a good reason. Not all security features are required in your environment. When you decide to tune security, create a benchmark before making any change to ensure that the change is improving performance.

What to do next

In a large scale deployment, performance is very important. Running benchmark measurements with different combinations of features can help you to determine the best performance versus the benefit of

configuration for your environment. Continue to run benchmarks if anything changes in your environment, to help determine the impact of these changes.

Secure Sockets Layer performance tips

Use this page to learn about Secure Sockets Layer (SSL) performance tips. Be sure to consider that performance issues typically involve trade-offs between function and speed. Usually, the more function and the more processing that are involved, the slower the performance.

The following are two types of Secure Sockets Layer (SSL) performance:

- Handshake
- Bulk encryption and decryption

When an SSL connection is established, an SSL handshake occurs. After a connection is made, SSL performs bulk encryption and decryption for each read-write. The performance cost of an SSL handshake is much larger than that of bulk encryption and decryption.

To enhance SSL performance, decrease the number of individual SSL connections and handshakes.

Decreasing the number of connections increases performance for secure communication through SSL connections, as well as non-secure communication through simple Transmission Control Protocol/Internet Protocol (TCP/IP) connections. One way to decrease individual SSL connections is to use a browser that supports HTTP 1.1. Decreasing individual SSL connections can be impossible if you cannot upgrade to HTTP 1.1.

Another common approach is to decrease the number of connections (both TCP/IP and SSL) between two WebSphere Application Server components. The following guidelines help to verify the HTTP transport of the application server is configured so that the Web server plug-in does not repeatedly reopen new connections to the application server:

- Verify that the maximum number of keep alives are, at minimum, as large as the maximum number of requests per thread of the web server (or maximum number of processes for IBM HTTP Server on UNIX). Make sure that the web server plug-in is capable of obtaining a keep alive connection for every possible concurrent connection to the application server. Otherwise, the application server closes the connection after a single request is processed. Also, the maximum number of threads in the web container thread pool should be larger than the maximum number of keep alives, to prevent the keep alive connections from consuming the web container threads.

Note: HTTP Transports have been deprecated. For instructions on how to set a maximum keep alive value for channel based configurations, see HTTP transport channel settings.

- Increase the maximum number of requests per keep alive connection. The default value is 100, which means the application server closes the connection from the plug-in after 100 requests. The plug-in then has to open a new connection. The purpose of this parameter is to prevent denial of service attacks when connecting to the application server and preventing continuous send requests to tie up threads in the application server.
- Use a hardware accelerator if the system performs several SSL handshakes.

Hardware accelerators currently supported by WebSphere Application Server only increase the SSL handshake performance, not the bulk encryption and decryption. An accelerator typically only benefits the web server because Web server connections are short-lived. All other SSL connections in WebSphere Application Server are long-lived.

The IBM Cryptographic Coprocessor is not supported for use with WebSphere Application Server. However, you can use the IBM Cryptographic Coprocessor to improve SSL performance for other products, such as IBM HTTP Server for iSeries®, which is powered by Apache.


- Use an alternative cipher suite with better performance.

The performance of a cipher suite is different with software and hardware. Just because a cipher suite performs better in software does not mean a cipher suite will perform better with hardware. Some

algorithms are typically inefficient in hardware, for example, Data Encryption Standard (DES) and triple-strength DES (3DES); however, specialized hardware can provide efficient implementations of these same algorithms.

The performance of bulk encryption and decryption is affected by the cipher suite used for an individual SSL connection. The following chart displays the performance of each cipher suite. The test software calculating the data was Java Secure Socket Extension (JSSE) for both the client and server software, which used no cryptographic hardware support. The test did not include the time to establish a connection, but only the time to transmit data through an established connection. Therefore, the data reveals the relative SSL performance of various cipher suites for long running connections.

Before establishing a connection, the client enables a single cipher suite for each test case. After the connection is established, the client times how long it takes to write an integer to the server and for the server to write the specified number of bytes back to the client. Varying the amount of data had negligible effects on the relative performance of the cipher suites.



An analysis of the above data reveals the following:

- Bulk encryption performance is only affected by what follows the WITH in the cipher suite name. This is expected since the portion before the WITH identifies the algorithm used only during the SSL handshake.
- MD5 and Secure Hash Algorithm (SHA) are the two hash algorithms used to provide data integrity. MD5 is generally faster than SHA, however, SHA is more secure than MD5.
- DES and RC2 are slower than RC4. Triple DES is the most secure, but the performance cost is high when using only software.
- The cipher suite providing the best performance while still providing privacy is SSL_RSA_WITH_RC4_128_MD5. Even though SSL_RSA_EXPORT_WITH_RC4_40_MD5 is cryptographically weaker than RSA_WITH_RC4_128_MD5, the performance for bulk encryption is the same. Therefore, as long as the SSL connection is a long-running connection, the difference in the performance of high and medium security levels is negligible. It is recommended that a security level of high be used, instead of medium, for all components participating in communication only among WebSphere Application Server products. Make sure that the connections are long running connections.

Tuning security performance

Use the following procedures to tune the performance, without compromising your security settings.

About this task

Enabling security decreases performance. The following tuning parameters provide ways to minimize this performance impact.

Procedure

- Disable security on any application servers that do not need security. You can disable security in the administrative console by clicking **Security > Global security** and deselecting the **Enable administrative security** option.
- Fine-tune the **Authentication cache timeout** value on the Authentication mechanisms and expiration panel in the administrative console. For more information, see the Global security settings topic.
- Configure the security cache properties. For more information, see the Authentication cache settings topic.
- Enable the **Enable SSL ID tracking** option on the Session management panel in the administrative console.
- Improve the performance of Web Services Security by downloading a Java Cryptography Extension (JCE) unlimited jurisdiction policy file that does not have restrictions on cryptography strength. See the information about tuning Web Services Security for Version 8.0 applications for details.
- Read the Secure Sockets Layer performance tips and “Tuning security configurations” on page 67 topics for more information.

Hardening security configurations

There are several methods that you can use to protect the WebSphere Application Server infrastructure and applications from different forms of attack. Several different techniques can help with multiple forms of attack. Sometimes a single attack can leverage multiple forms of intrusion to achieve the end goal.

About this task

For example, in the simplest case, network sniffing can be used to obtain passwords and those passwords can then be used to mount an application-level attack. The following issues are discussed in IBM WebSphere Developer Technical Journal: WebSphere Application Server V5 advanced security and system hardening:

Procedure

- Take preventative measures to protect the infrastructure.
- Make applications less vulnerable to attack.
- At a minimum, ensure administrative security is enabled in all WebSphere processes. This protects access to the administrative ConfigService interface and managed beans (MBeans) that enables control over the WebSphere process if it is compromised.
- Ensure Secure Sockets Layer (SSL) is used whenever possible, and mutual SSL whenever possible. However, mutual SSL requires all clients to supply a trusted personal certificate in order to connect.
- Remove any unnecessary certificate authority (CA) signer certificates from your trust stores.
- Change default keystore passwords during or after profile creation using the AdminTask `changeMultipleKeyStorePasswords` command.
- Change your Lightweight Third-Party Authentication (LTPA) keys periodically. You can configure the automatic regeneration of LTPA keys if necessary.
- Common Secure Interoperability version 2 (CSIv2) inbound Basic authentication is supported in this release of WebSphere Application Server. The authentication default is 'required'.

What to do next

Note: In this release of WebSphere Application Server, more security hardening features of the server are enabled by default. However, if the features are not enabled after migration you can enable them yourself. See the Security hardening features enablement and migration article for more information.

For additional information about hardening security configurations, see the WebSphere Application Server security web page.

Enablement and migration considerations of Security hardening features

In this release of WebSphere Application Server, more security hardening features of the server are enabled out-of-the-box by default. When migrating, the settings that were enabled prior to migration are retained. However, if the features are not enabled after migration you can enable them yourself.

Note: To ensure that WebSphere Application Server configuration is set to be secure by default, the following defaults have been changed as part of the new security hardening features in WebSphere Application Server Version 8.0:

- Enablement of Secure Sockets Layer (SSL)-required on Common Secure Interoperability version 2 (CSlv2) transport by default

The following settings for the CSlv2 transport layer exist: **TCP/IP** for a TCP/IP connection, **SSL-supported** for a TCP/IP or an SSL connection, and **SSL-required** for an SSL connection only. SSL-required is the new default in this release of WebSphere Application Server. Switching to **SSL-required** as the default setting ensures that all CSlv2 connections into and out of the server are using the secure SSL connection.

- Enablement of the HttpOnly attribute on LTPA cookies by default

When the `com.ibm.ws.security.addHttpOnlyAttributeToCookies` custom property is set to `true`, the HttpOnly attribute is added to those security cookies (LTPA and WASReqURL cookies) that are created by the server. The HttpOnly attribute is a browser attribute created to prevent client side applications (such as Java scripts) from accessing cookies to prevent some cross-site scripting vulnerabilities. This attribute is now configurable in the administrative console. Prior to WebSphere Application Server Version 8.0, the `com.ibm.ws.security.addHttpOnlyAttributeToCookies` custom property default was `false`. For WebSphere Application Server Version 8.0, the default is now `true` for both the LTPA cookie and the Session Cookie.

For more information see the custom property `com.ibm.ws.security.addHttpOnlyAttributeToCookies` in the Security custom properties article.

- Enablement of session security integration by default

Only authenticated users can access sessions created in secure pages. The session management facility uses the security infrastructure to determine the authenticated identity associated with a client HTTP request, and either retrieves or creates a session. For more information on session security, read the Session security support article.

Along with enabling session security integration, credential persistence is enabled as well. This allows login information to be available to unprotected web clients to enable additional access to user information. For more information on credential persistence, see the "Use available authentication data when an unprotected URI is accessed" feature in the web authentication settings article.

Enabling the new security hardening features after migration

If the new security features are not enabled after migration, you can enable them yourself using the administrative console or by scripting.

Enablement of SSL by default on CSiv2

To enable SSL by default for inbound and outbound transports on CSiv2:

If you are using the administrative console, select **Security > Global security > RMI/IOP > CSiv2 inbound communications**. In the Transport box, select **SSL- required** from the pull-down list and then click **Apply**.

Repeat the same steps for CSiv2 outbound communications and click **Security > Global security > RMI/IOP > CSiv2 outbound communications**. In the Transport box, select **SSL- required** from the menu list and then click **Apply**.

If you want to enable SSL by default for inbound and outbound transports on CSiv2 using scripting, use the `configureCSIInbound` and `configureCSIOutbound` commands. See the [Configuring Common Secure Interoperability authentication using scripting](#) topic for more information.

For the client side, edit the `com.ibm.client.props` file. Change `com.ibm.CSI.performTransportAssocSSLTLSRequired` to `true` and change `com.ibm.CSI.performTransportAssocSSLTLSSupported` to `false`.

Enablement of the HttpOnly cookie attribute

To enable the HttpOnly attribute on cookies attribute by default:

If you are using the administrative console, click **Security > Global security > Custom properties**. Click **New** and enter `com.ibm.ws.security.addHttpOnlyAttributeToCookies` for the Name and `true` for the Value.

You can also enable the HttpOnly attribute using the administrative console by clicking **Security > Global security > Single sign-on (SSO)**. Click **Set security cookies to HTTPOnly to help prevent cross-site scripting attacks**, and then click **Apply**.

To enable the HttpOnly attribute on cookies attribute by default using scripting, use the `setAdminActiveSecuritySettings` command.

Enablement of session security integration

To enable session security integration for each server by using the administrative console, select **Servers > Server types > WebSphere application servers > <server1> > Session management**. Select the **security integration** check box.

To enable persisting credentials from the administrative console, click **Security > Global security > Web and SIP security > General settings**. Select the **Use available authentication data when an unprotected URI is accessed** check box.

Security hardening features enablement troubleshooting

When the new security hardening features are enabled you might see some differences in system behavior depending upon which environment you might have used in the past.

For example, if you are coming from an environment where CSiv2 transport was set to the previous default of SSL-supported, you do not experience any differences, as SSL-supported communicates with both TCP/IP and SSL connections. If a problem is encountered, however, certificates might not have been exchanged correctly to enable the client and server to communicate. Read about the [Secure communications using Secure Sockets Layer \(SSL\)](#) topic for more information.

If you worked in an environment where TCP/IP is used for the connection to CSiv2, you might experience connection problem to the SSL-enabled CSiv2 connection. The server configuration can be modified to **SSL-supported** or to **TCP/IP** if SSL is not required.

For the HttpOnly attribute, when the attribute is added to the security cookies, the browser prevents client side scripts from accessing these cookies. In most case this should be the default behavior to minimize

cross-site scripting vulnerabilities. If there is an absolute need to allow client-side scripts to access WebSphere security cookies, and you are aware of the possible consequences, then the setting of the HttpOnly attribute can be disabled.

However, the HttpOnly attribute can possibly uncover client-side scripts that are used to access WebSphere cookies, and can then use them even though it was not intended to do so. If this happens, the web application that enables the scripts to access the WebSphere cookies must be evaluated.

For session security integration enablement, when session integrated security is enabled you might receive an `UnauthorizedSessionRequestException` exception on servlets if they access a session that belongs to authenticated identities other than to the identity that currently owns the session. If you do not want this checking to occur, you can disable session security from the server that is experiencing the problem.

Securing passwords in files

Password encoding and encryption deters the casual observation of passwords in server configuration and property files.

About this task

The following topics can be used to add protection for passwords located in files:

Procedure

- Password encoding and encryption Passwords are automatically encoded with a simple masking algorithm in various WebSphere Application Server ASCII configuration files. Additionally, you can manually encode passwords in properties files that are used by Java clients and by administrative commands for WebSphere Application Server. For more information on password encoding and encryption, see “Password encoding and encryption.”
- Encoding passwords in files WebSphere Application Server contains some encoded passwords that are not encrypted. The **PropFilePasswordEncoder** utility is included to encode these passwords. For more information on encoding passwords in a file, see “Encoding passwords in files” on page 79.
- Enabling custom password encryption You need to protect passwords that are contained in your WebSphere Application Server configuration. You can added protection by creating a custom class for encrypting the passwords. For more information on custom password encryption, see “Enabling custom password encryption” on page 86.

Password encoding and encryption

Password encoding deters the casual observation of passwords in server configuration and property files.

By default, passwords are automatically encoded with a simple masking algorithm in various WebSphere Application Server ASCII configuration files. Additionally, you can manually encode passwords in properties files that are used by Java clients and by administrative commands for WebSphere Application Server.

The default encoding algorithm is referred to as XOR. An alternate OS400 encoding algorithm can be used with WebSphere Application Server for IBM i that exploits native validation list (*VLDL) objects only. With the OS400 algorithm, passwords are stored in an encrypted form within a validation list. The configuration files contain indexes to the stored passwords instead of the masked passwords, as is done with the XOR algorithm.

Encoded passwords use the following syntax:

```
{algorithm}encoded_password
```

where {algorithm} is a tag that specifies the algorithm that is used to encode the password, which is either XOR or OS400. The *encoded_password* variable is the encoded value of the password. When a server or client needs to decode a password, it uses the tag to determine what algorithm to use and then uses that algorithm to decode the encoded password.

Java clients use passwords from the `sas.client.props` file, which is in the *profile_root/properties* directory.

To use password encoding with Java clients, the passwords must be manually encoded in the `sas.client.props` file using the **PropFilePasswordEncoder** tool.

The administrative commands for WebSphere Application Server use passwords from the `soap.client.props` file, which is also located in the *profile_root/properties* directory, for SOAP connections. Some administrative commands optionally use passwords from the `sas.client.props` file in the *profile_root/properties* for Remote Method Invocation (RMI) connections. To use password encoding with administrative commands, you must manually encode the passwords in the `soap.client.props` and `sas.client.props` files using the **PropFilePasswordEncoder** tool.

Attention: Whether you select to use the OS400 encoding algorithm or the default encoding algorithm, encoding is not sufficient to fully protect passwords. Native security is the primary mechanism for protecting passwords that are used in the configuration and property files for WebSphere Application Server.

Issues to consider when you use the OS400 password encoding algorithm

The following issues are important for you to consider before deciding to use the OS400 password encoding algorithm:

- You must set the QRETSVRSEC operating system value to 1 to use on the system that hosts the Java client application or WebSphere Application Server. With this setting, WebSphere Application Server can retrieve the encrypted passwords from the validation list.
 - Attention:** The QRETSVRSEC system value affects access to the encrypted data in all of the validation lists on your operating system. Do not use the OS400 password encoding algorithm if this setting is not consistent with your security policy for your operating system.
- You can use the OS400 algorithm with server instances only when all of the server instances within the administrative domain for WebSphere Application Server reside on the same IBM i system. Consider the following related issues:
 - Administrative domains for WebSphere Application Server can extend across multiple IBM i systems. You can use the OS400 password algorithm only when all of the servers within an administrative domain reside on the same IBM i system.
 - Server configuration XML files contain encoded passwords. If the passwords that are contained in the XML files are encoded using the OS400 encoding algorithm, those encodings are valid only for the Application Server profiles on the same IBM i system on which the passwords were originally encoded. Copies of configuration files that contain passwords that are encoded using the OS400 encoding algorithm cannot be used to configure servers on other IBM i systems.
 - All server instances within an administrative domain must be configured to use the same native validation list (*VLDL) object.
- For Java clients, you can use the OS400 password algorithm on any IBM i system. However, option 1 must be installed on the system that hosts the Java client.
- If an error occurs while a password is encoded using the OS400 encoding algorithm, the XOR encoding algorithm is used to encode the password. An error might occur if an administrator manually creates the validation list object and grants insufficient authority to the validation list object for the IBM i QEJB user profile.

Object and file security:

This topic discusses the various objects and files that contain sensitive information and need to be protected.

Secure integrated file system files

In addition to enterprise beans and servlets, WebSphere Application Server accesses integrated file system stream files. The following files might contain sensitive information. It is recommended that you give these files close consideration to ensure that unauthorized access is not granted.

- In the `/properties` subdirectory of your profile, the following files can contain user IDs and passwords:
 - `sas.client.props`
 - `soap.client.props`
 - `sas.stdclient.properties`
 - `sas.tools.properties`
 - `wssserver.key`

By default, the `/properties` subdirectory is located in the `profile_root` directory. Each of the previous files is shipped with `*PUBLIC` authority set to `*EXCLUDE`. The QEJBSVR user profile is granted `*RW` authority to these files. Additional protection is available through password encoding. For more information, see “Password encoding and encryption” on page 76.

- In the `/etc` subdirectory of your profile, protect all of the key (KDB) files and trust (JKS) files that you create for your WebSphere Application Server profile.

For the JKS files, the QEJBSVR user profiles should have `*R` authority and `*PUBLIC` should have `*EXCLUDE` authority.

For the KDB files, the user profile that the web server is running under should have `*RX` authority and `*PUBLIC` should have `*EXCLUDE` authority.

Secure database resources for WebSphere Application Server

WebSphere Application Server uses tables to persist data for user applications such as enterprise beans persistence and servlet session data. You have several options for controlling which user profiles are allowed access to this user data. For more information, see Database access security.

Secure WebSphere Application Server files

When you enable WebSphere Application Server security, the server user profile and password are placed into server configuration files, which should be maintained in a secure way using operating system security. Additionally, you can password protect some WebSphere Application Server resources. These passwords are also placed in server configuration files. The server automatically encodes passwords to deter casual observation, but password encoding alone is not sufficient protection.

The following files are located in the `/config` subdirectory of your profile and they can contain user identifiers and passwords:

- `cells/cell_name/security.xml`
- `cells/cell_name/nodes/node_name/resources.xml`
- `cells/cell_name/nodes/node_name/servers/server_name/server.xml`

For example, for the default profile, the `server_name` is `server1`.

The server user profile and password are used for authenticating the server when it initializes. This authentication is required for the following reasons:

- The user ID and password are used as the system identity for the server when an enterprise bean security is deployed to use `SYSTEM_IDENTITY` for method delegation. In this case, the user ID and password are used when method calls are made from one enterprise bean to another.

- The user ID and password are used to authenticate servers for inter-server communication. Because security for these files can be compromised, use a non-default user profile for the server identity and password. The default user profile is QEJBSVR. If you use the local OS user registry, you might choose to create and use a user profile that has no special authorities. For more information, see Running application servers under specific user profiles.

Secure user profiles for WebSphere Application Server

When WebSphere Application Server is first installed, by default, it uses the following user profiles:

QEJB This profile provides access to some administrative data, including passwords.

QEJBSVR

This profile provides the context in which your WebSphere Application Server runs. For security or administrative purposes, you might want to create other user profiles under which to run various parts of WebSphere Application Server. For more information, see Running application servers under specific user profiles.

Encoding passwords in files

The purpose of password encoding is to deter casual observation of passwords in server configuration and property files. Use the **PropFilePasswordEncoder** utility to encode passwords stored in properties files. WebSphere Application Server does not provide a utility for decoding the passwords. Encoding is not sufficient to fully protect passwords. Native security is the primary mechanism for protecting passwords used in WebSphere Application Server configuration and property files.

About this task

WebSphere Application Server contains several encoded passwords in files that are not encrypted. WebSphere Application Server provides the **PropFilePasswordEncoder** utility, which you can use to encode passwords. The purpose of password encoding is to deter casual observation of passwords in server configuration and property files. The **PropFilePasswordEncoder** utility does not encode passwords that are contained within XML or XMI files.

Table 5. XML and XMI files that contain encoded passwords. Instead, WebSphere Application Server automatically encodes the passwords in these files. XML and XMI files that contain encoded passwords include the following:

File name	Additional information
<code>profile_root/config/cells/cell_name/security.xml</code>	The following fields contain encoded passwords: <ul style="list-style-type: none"> • LTPA password • JAAS authentication data • User registry server password • LDAP user registry bind password • Keystore password • Truststore password • Cryptographic token device password
<code>war/WEB-INF/ibm_web_bnd.xml</code>	Specifies the passwords for the default basic authentication for the resource-ref bindings within all the descriptors, except in the Java cryptography architecture
<code>ejb_jar/META-INF/ibm_ejbjar_bnd.xml</code>	Specifies the passwords for the default basic authentication for the resource-ref bindings within all the descriptors, except in the Java cryptography architecture
<code>client_jar/META-INF/ibm_appclient_bnd.xml</code>	Specifies the passwords for the default basic authentication for the resource-ref bindings within all the descriptors, except in the Java cryptography architecture
<code>ear/META-INF/ibm_application_bnd.xml</code>	Specifies the passwords for the default basic authentication for the run as bindings within all the descriptors
<code>profile_root/config/cells/cell_name/nodes/node_name/servers/security.xml</code>	The following fields contain encoded passwords: <ul style="list-style-type: none"> • Keystore password • Truststore password • Cryptographic token device password • Session persistence password • DRS client data replication password

Table 5. XML and XMI files that contain encoded passwords (continued). Instead, WebSphere Application Server automatically encodes the passwords in these files. XML and XMI files that contain encoded passwords include the following:

File name	Additional information
<i>profile_root</i> /config/cells/ <i>cell_name</i> /nodes/ <i>node_name</i> /servers/server1/resources.xml	The following fields contain encoded passwords: <ul style="list-style-type: none"> • WAS40Datasource password • mailTransport password • mailStore password • MQQueue queue mgr password
<i>profile_root</i> /config/cells/ <i>cell_name</i> /ws-security.xml	
ibm-webservices-bnd.xmi	
ibm-webservicesclient-bnd.xmi	

Table 6. The PropFilePasswordEncoder utility - Partial File List. You use the PropFilePasswordEncoder utility to encode the passwords in properties files. These files include:

File name	Additional information
<i>profile_root</i> /properties/sas.client.props	Specifies the passwords for the following files: <ul style="list-style-type: none"> • com.ibm.ssl.keyStorePassword • com.ibm.ssl.trustStorePassword • com.ibm.CORBA.loginPassword
<i>profile_root</i> /properties/sas.tools.properties	Specifies passwords for: <ul style="list-style-type: none"> • com.ibm.ssl.keyStorePassword • com.ibm.ssl.trustStorePassword • com.ibm.CORBA.loginPassword
<i>profile_root</i> /properties/sas.stdclient.properties	Specifies passwords for: <ul style="list-style-type: none"> • com.ibm.ssl.keyStorePassword • com.ibm.ssl.trustStorePassword • com.ibm.CORBA.loginPassword
<i>profile_root</i> /properties/wsserver.key	
<i>profile_root</i> /profiles/AppSrvXX/properties/sib.client.ssl.properties	Specifies passwords for: <ul style="list-style-type: none"> • com.ibm.ssl.keyStorePassword • com.ibm.ssl.trustStorePassword
<i>profile_root</i> /UDDIReg/scripts/UDDIUtilityTools.properties	Specifies passwords for: <ul style="list-style-type: none"> • trustStore.password

To encode a password again in one of the previous files, complete the following steps:

Procedure

1. Access the file using a text editor and type over the encoded password. The new password is shown is no longer encoded and must be re-encoded.
2. Use the **PropFilePasswordEncode** script in the *profile_root/bin/* directory to encode the password again.

If you are encoding files that are not SAS properties files, type PropFilePasswordEncoder "*file_name*" *password_properties_list*

"*file_name*" is the name of the SAS properties file and *password_properties_list* is the name of the properties to encode within the file.

Note: Only the password should be encoded in this file by using the **PropFilePasswordEncoder** tool.

Use the **PropFilePasswordEncoder** tool to encode WebSphere Application Server password files only. The utility cannot encode passwords that are contained in XML files or other files that contain open and close tags. To change passwords in these files, use the administrative console or an assembly tool such as the Rational Application Developer.

Results

If you reopen the affected files, the passwords are encoded. WebSphere Application Server does not provide a utility for decoding the passwords.

Example

The following example shows how to use the **PropFilePasswordEncoder** tool:

```
PropFilePasswordEncoder C:\WASV8\WebSphere\AppServer\profiles\AppSrv\properties
\sas.client.props com.ibm.ssl.keyStorePassword,com.ibm.ssl.trustStorePassword
```

where:

PropFilePasswordEncoder is the name of the utility that you are running from the profile_root/profiles/profile_name/bin directory.

C:\WASV6\WebSphere\AppServer\profiles\AppSrv\properties\sas.client.props is the name of the file that contains the passwords to encode.

com.ibm.ssl.keyStorePassword is a password to encode in the file.

com.ibm.ssl.trustStorePassword is a second password to encode in the file.

Manually encoding passwords in properties files:

To use password encoding with WebSphere Application Server administrative commands and Java clients, passwords must be manually encoded in the soap.client.props and sas.client.props files using the **PropFilePasswordEncoder** tool.

Before you begin

To run the script, your user profile must have *ALLOBJ authority.

About this task

Use the **PropFilePasswordEncoder** utility to encode the passwords in properties files. The **PropFilePasswordEncoder** utility is a Qshell script. Complete the following steps to manually encode the passwords:

Procedure

1. Sign on the server with a user profile that has all object (*ALLOBJ) special authority.
2. Run the Start Qshell (STRQSH) command on a command line to start the Qshell environment.
3. Use the **PropFilePasswordEncoder** utility to encode the passwords.

For example, to encode the passwords for properties in the sas.client.props file for the default WebSphere Application Server profile (in a default installation), enter the following command:

```
profile_root/bin/PropFilePasswordEncoder
-profileName server1
profile_root/properties/sas.client.props -SAS
```

For example, to encode the passwords for properties in the soap.client.props file for the default stand-alone application server profile, enter the following command:

```
profile_root/bin/PropFilePasswordEncoder
-profileName server1
profile_root/properties/soap.client.props
com.ibm.SOAP.loginPassword,com.ibm.ssl.keyStorePassword,
com.ibm.ssl.trustStorePassword
```

For more information on the sas.client.props utility, see the “PropFilePasswordEncoder command reference” on page 82.

Results

The passwords are encoded in the `soap.client.props` and `sas.client.props` files.

What to do next

See “Restoring or replacing damaged validation list objects” on page 85 for information on how to restore or replace a damaged validation list object.

PropFilePasswordEncoder command reference:

The **PropFilePasswordEncoder** command encodes passwords that are located in plain text property files. This command encodes both Secure Authentication Server (SAS) property files and non-SAS property files. After you encode the passwords, a decoding command does not exist.

To encode passwords, you must run this command from the directory:

- `app_server_root/bin`

Note: If you need to custom encode passwords in property files, manually edit the `PropFilePasswordEncoder.sh` or `PropFilePasswordEncoder.bat` file before issuing this command. See the topic [Implementing custom password encryption](#) for a description of the lines that need to be added to this file.

To run this script, your user profile must have `*ALLOBJ` authority.

Syntax

The command syntax is as follows:

```
PropFilePasswordEncoder "fileName" { passwordPropertiesList  
  [-SAS ] [ -profileName profile ] [ -help | -? ]
```

Important: You must specify either the `passwordPropertiesList` parameter or the `-SAS` parameter.

Parameters

The following option is available for the **PropFilePasswordEncoder** command:

fileName

This required parameter specifies the name of the file in which passwords are encoded.

passwordPropertiesList

This parameter is required if you are encoding passwords in property files other than the `sas.client.props` file. Specify one or more password properties that you want to encode. The password properties list should be delimited by commas.

-SAS

This parameter is required if you are encoding passwords in the `sas.client.props` file.

-profileName

This parameter is optional. The profile value specifies an application server profile name. The script uses the password encoding algorithm that it retrieves from the specified profile. If you do not specify this parameter, the script uses the default profile.

-noBackup

This parameter is optional and the default. The script does not create a backup file. The default value can be altered by adding following Java System Property:

```
"-Dcom.ibm.websphere.security.util.createBackup=true".
```

-Backup

This parameter is optional. The script creates a backup file, <file_name>.bak, which contains passwords in clear text.

-help or -?

If you specify this parameter, the script ignores all other parameters and displays usage text.

The following command encodes the passwords in the `sas.client.props` file for the default stand-alone application server profile:

```
app_server_root/bin/PropFilePasswordEncoder  
profile_root/default/properties/sas.client.props -SAS
```

The following command encodes the passwords in the `soap.client.props` file for the default stand-alone application server profile:

```
app_server_root/bin/PropFilePasswordEncoder  
profile_root/default/properties/soap.client.props  
com.ibm.SOAP.loginPassword,com.ibm.ssl.keyStorePassword,com.ibm.ssl.trustStorePassword
```

Attention: These commands are displayed on multiple lines for illustrative purposes only.

Enabling the non-default OS/400 password encoding algorithm:

The purpose of password encoding is to deter casual observation of passwords in server configuration and property files.

Before you begin

Make sure all server profiles within the administration console reside on the same IBM i system.

About this task

By default, passwords are automatically encoded with a simple masking algorithm in various ASCII configuration files for WebSphere Application Server. You can manually encode passwords in properties files that are used by Java clients and by Application Server administrative commands.

For a description of the OS400 encoding algorithm, see “Password encoding and encryption” on page 76. To enable the OS400 password encoding algorithm for a WebSphere Application Server profile, complete these steps:

Procedure

1. Set the `os400.security.password` properties to turn on the OS400 password encoding algorithm and to specify which the validation list object to use.

Use the same validation list object for all WebSphere Application Server profiles. However, it is not recommended if you do not back up the objects and data for all profiles simultaneously. Consider your backup and restore policy when you decide what validation list object to use for each WebSphere Application Server profile.

To set the properties, complete one of these steps:

- Use the **-os400passwords** and **-validationlist** options for the `manageprofiles -create` utility, which is located in the `app_server_root/bin` directory, to set the properties when creating the profile. To create a WebSphere Application Server profile named `prod`, and to enable that profile for the OS400 encoding algorithm using the `/QSYS.LIB/QUSRSYS.LIB/WAS.VLDL` validation list object, you can complete the following steps:
 - a. Run the Start Qshell (STRQSH) command on the IBM i command line.
 - b. In Qshell, run the following command:

```

app_server_root/bin/manageprofiles
-create -profileName prod -startingPort 10150
-templatePath default -os400passwords
-validationlist /QSYS.LIB/QUSRSYS.LIB/WAS.VLDL

```

The previous command is on multiple lines for illustration purposes only.

- Set the Java system properties in the **setupCmdLine** Qshell script of the WebSphere Application Server profile. To enable the OS400 password encoding algorithm, edit the *profile_root/bin/setupCmdLine* script using the following steps:
 - a. Set the `os400.security.password.encoding.algorithm` property to `OS400`. The default setting is `XOR`.
 - b. Set the `os400.security.password.validation.list.object` property to the absolute name of the validation list that you need to use. The default setting is `/QSYS.LIB/QUSRSYS.LIB/EJSADMIN.VLDL`.
 - c. Save the file.
- 2. Grant the QEJB user profile run authority (*X) to the library that contains the validation list. If QEJB already has the minimum required authority (*X) to access the library, then proceed to the next step.
 - a. Use the Display Authority (DSPAUT) to check for the minimum required authority if the validation list is created in the `/QSYS.LIB/WSADMIN.LIB` file.

For example:

```
DSPAUT OBJ('/QSYS.LIB/WSADMIN.LIB')
```

- b. Use the Change Authority (CHGAUT) command to grant run authority to the QEJB profile only if the QEJB profile does not already have this authority.

For example:

```
CHGAUT OBJ('/QSYS.LIB/WSADMIN.LIB') USER(QEJB) DTAAUT(*X)
```

3. Create a native validation list object (*VLDL). This step is optional for server profiles. The validation list object is created when the server is started. For remote profiles, create the validation list if the validation list does not already exist on the system that hosts the remote profile. Also, consider your backup and restore policy when you decide what validation list object to use with each remote profile.

Attention: When you use the OS400 password encoding algorithm, the Java client is not required to reside on the same IBM i system as the WebSphere Application Server profile that the client accesses. To create a validation list object, perform the following steps with an IBM i user profile that has *ALLOBJ special authority:

- a. Sign on the server with a user profile that has the *ALLOBJ special authority.
- b. Use the Create Validation List (**CRTVLDL**) command to create the validation list object.

For example, to create the WSVLIST validation list object in the WSADMIN.LIB library, use the following command:

```
CRTVLDL VLDL(WSADMIN/WSVLIST)
```

- c. Grant the QEJB user profile *RWX authority to the validation list object. For example, to grant *RWX authority to the WSVLIST validation list object in the WSADMIN library, use the following command:

```
CHGAUT OBJ('/QSYS.LIB/WSADMIN.LIB/WSVLIST.VLDL') USER(QEJB) DTAAUT(*RWX)
```

4. Use the Change System Value (**CHGSYSVAL**) command to set the QRETSVRSEC system value to 1. For example:

```
CHGSYSVAL SYSVAL(QRETSVRSEC) VALUE('1')
```

5. For server profile, start or restart the server and wait until the server is ready for service before attempting to manually encode passwords in properties files that belong to the profile.

Results

You have enabled the OS400 password encoding algorithm.

What to do next

After completing the previous steps and restarting the server, you can manually encode passwords in properties files. See “Manually encoding passwords in properties files” on page 81 for more information.

Changing encoding algorithm from OS400 to XOR:

Use these steps to change your encoding algorithm from OS400 to XOR.

Before you begin

WebSphere Application Server supports both the OS400 and the XOR encoding algorithms. The default encoding algorithm is XOR. For more conceptual information on these algorithms, see “Password encoding and encryption” on page 76.

About this task

There might be instances where you need to use the OS400 encoding algorithm. See “Enabling the non-default OS/400 password encoding algorithm” on page 83.

If you use the OS400 encoding algorithm, but change the `os400.security.password.encoding.algorithm` property value to XOR as described in the first step, then all of the passwords remain encoded with the OS400 algorithm. However, after you restart the server, when you make changes to the passwords through the administrative console, the passwords are encoded using the XOR algorithm.

The following steps describe how to change your encoding algorithm from OS400 to XOR:

Procedure

1. Set the `os400.security.password.encoding.algorithm` property to XOR for each WebSphere Application Server profile that uses the Validation list object.
2. Stop all of the servers.
3. Edit the configuration files and change all of the encoded passwords to their unencoded values.
4. Edit the properties files and change all of the encoded passwords to their clear text values.

Results

You have changed your encoding algorithm from OS400 to XOR.

What to do next

After you completing these steps, you can encode the passwords in the properties files. See “Manually encoding passwords in properties files” on page 81 for more information.

Restoring or replacing damaged validation list objects:

Periodically, you should save your validation list objects with the other configuration data objects that are used by WebSphere Application Server. Use this task if you need to restore or replace a damaged validation list object.

About this task

You can share validation lists between multiple WebSphere Application Server profiles. For example, if you have two profiles of WebSphere Application Server, default and prod, both profiles can use the `/QSYS.LIB/QUSRSYS.LIB/EJSADMIN.VLDL` validation list.

To restore or replace a damaged validation list object, complete the following steps:

Procedure

1. Replace the encoded passwords with the unencoded value of the password for all of the WebSphere Application Server profiles that use the validation list object. To replace the password values, complete the following steps:
 - a. Stop each of the servers.
 - b. Set the `os400.security.password.validation.list.object` property for all of the servers to the absolute name of the new validation list that you want to use. You can use an existing validation list object or specify a new object. For new validation list objects, create them manually or use the objects that are created automatically when the server is restarted. For more information on manually creating validation list objects, see “Manually encoding passwords in properties files” on page 81.
 - c. Edit the configuration files and set each encoded password to the appropriate clear text value.
2. Edit the `sas.client.props` and `soap.client.props` files and set each encoded password to the appropriate unencoded value before manually encoding the passwords.
3. Restart the servers for all of the WebSphere Application Server profiles whose validation list objects that are replaced.

Results

After restarting the server, you have successfully replaced a damaged validation list object.

What to do next

For additional information on backing up your data objects, see “Backing up security configuration files” on page 88.

Enabling custom password encryption

You need to protect passwords that are contained in your WebSphere Application Server configuration. After creating your server profile, you can add protection by creating a custom class for encrypting the passwords.

Before you begin

Create your custom class for encrypting passwords. For more information, see Plug point for custom password encryption.

About this task

Complete the following steps to enable custom password encryption.

Procedure

1. Add the following system properties for every server and client process. For server processes, update the `server.xml` file for each process. Add these properties as a `genericJvmArgument` argument preceded by a **-D** prefix.

```
com.ibm.wsspi.security.crypto.customPasswordEncryptionClass=  
com.acme.myPasswordEncryptionClass  
com.ibm.wsspi.security.crypto.customPasswordEncryptionEnabled=true
```

Tip: If the custom encryption class name is `com.ibm.wsspi.security.crypto.CustomPasswordEncryptionImpl`, it is automatically enabled when this class is present in the classpath. Do not define the system properties that are listed previously when the custom implementation has this package and class name. To disable encryption for this class, you must specify `com.ibm.wsspi.security.crypto.customPasswordEncryptionEnabled=false` as a system property.

2. Choose one of the following methods to configure the WebSphere Application Server runtime to load the custom encryption implementation class:
 - Place the custom encryption class in a Java archive (JAR) file that resides in the `${WAS_INSTALL_ROOT}/classes` directory, which you have created.

gotcha: WebSphere Application Server does not create the `${WAS_INSTALL_ROOT}/classes` directory. For more information on the classes directory, see the topic, "Creating a classes subdirectory in your profile for custom classes".
 - Place the custom encryption class in a Java archive (JAR) file that resides in the `${WAS_HOME}/lib/ext` directory.
3. Restart all server processes.
4. Edit each configuration document that contains a password and save the configuration. All password fields are then run through the **WSEncoderDecoder** utility, which calls the plug point when it is enabled. The `{custom:alias}` tags are displayed in the configuration documents. The passwords, even though they are encrypted, are still Base64-encoded. They seem similar to encoded passwords, except for the tags difference.
5. Encrypt any passwords that are in client-side property files using the **PropsFilePasswordEncoder** (.bat or .sh) utility. This utility requires that the properties listed previously are defined as system properties in the script to encrypt new passwords instead of encoding them.
6. To decrypt passwords from client Java virtual machines (JVMs), add the properties listed previously as system properties for each client utility.
7. Ensure that all nodes have the custom encryption classes in their class paths prior to enabling this function. The order in which enablement occurs is important. When adding a new node to a cell that contains password encryption, the new node must contain the custom encryption classes prior to using the **addNode** command. Consider the following WebSphere Application Server, Network Deployment enablement scenarios:
 - a. The StandAloneProfile profile is encrypting passwords with a different key prior to federation to a deployment manager cell. For this scenario, you must uninstall custom password encryption to ensure that the configuration has `{xor}` tags preceding the passwords prior to running the **addNode** command. The same implementation of the plug point must be in the `/classes` directory prior to running the **addNode** command, and the proper configuration properties are set so that the new node can recognize the encrypted password format of the `security.xml` file after federation completes.
 - b. The StandAloneProfile profile does not have password encryption configured prior to federation to a deployment manager cell. The same implementation of the plug point must be in the `/classes` directory prior to running the **addNode** command, and the proper configuration properties are set so that the new node can recognize the encrypted password format of the `security.xml` file after federation completes.
 - c. If enabling custom password encryption in a cell with multiple nodes present, update the correct configuration properties and have the custom password encryption implementation class located on all nodes. Stop all processes in the cell, and then start the deployment manager. Use the administrative console to edit the security configuration and then save it. Verify that the passwords are encrypted by looking at the `security.xml` file to see if the passwords are preceded by `{custom:alias}` tags.
 - d. Run the **syncNode** command on each node, and start each one individually. If any nodes fail to start, make sure that they have custom password encryption enabled properly in each `security.xml` file and that the implementation class is in the appropriate `/classes` directory for the platform.

Results

Custom password encryption is enabled.

What to do next

If custom password encryption fails or is no longer required, see “Disabling custom password encryption.”

Disabling custom password encryption:

If custom password encryption fails or is no longer required, perform this task to disable custom password encryption.

Before you begin

Enable custom password encryption.

About this task

Complete the following steps to disable custom password encryption.

Procedure

1. Change the `com.ibm.wsspi.security.crypto.customPasswordEncryptionEnabled` property to be `false` in the `security.xml` file, but leave the `com.ibm.wsspi.security.crypto.customPasswordEncryptionClass` property configured. Any passwords in the model that still have the `{custom:alias}` tag are decrypted by using the customer password encryption class.
2. If an encryption key is lost, any passwords that are encrypted with that key cannot be retrieved. To recover a password, retype the password in the password field in plaintext and save the document. The new password must be written out using encoding with the `{xor}` tag with scripting or from the administrative console.

```
com.ibm.wsspi.security.crypto.customPasswordEncryptionClass=  
    com.acme.myPasswordEncryptionClass  
com.ibm.wsspi.security.crypto.customPasswordEncryptionEnabled=false
```

3. Restart all processes to make the changes effective.
4. Edit each configuration document that contains an encrypted password and save the configuration. All password fields are then run through the **WSEncoderDecoder** utility, which calls the plug point in the presence of the `{custom:alias}` tag. The `{xor}` tags display in the configuration documents again after the documents are saved.
5. Decrypt and encode any passwords that are in client-side property files using the **PropsFilePasswordEncoder** (.bat or .sh) utility. If the encryption class is specified, but custom encryption is disabled, running this utility converts the encryption to encoding and causes the `{xor}` tags to display again.
6. Disable custom password encryption from the client Java virtual machines (JVMs) by adding the system properties listed previously to all client scripts. This action enables the code to decrypt passwords, but this action is not used to encrypt them again. The `{xor}` algorithm becomes the default for encoding. Leave the custom password encryption class defined for a time in case any encrypted passwords still exist in the configuration.

Results

Custom password encryption is disabled.

Backing up security configuration files

Back up your security configuration files to prevent the loss of information due to a potential system failure.

About this task

Consider backing up the following security information:

Procedure

- Back up your user profiles.

When you use local OS security, back up your user profiles, using the normal save procedures for user profiles. For more information, see the Backup and Recovery Guide by searching for "Saving group and user profiles" in the iSeries information center.

For information about the Directory Services Product (LDAP server), see the iSeries information center.

For information about Lotus® Domino®, see the Lotus Domino reference library.

- Back up your security property files.

Security settings are saved in several properties files. By default, these properties are located in the *profile_root/properties* directory. The default stand-alone profile name is default. If you define additional WebSphere Application Server profiles, there are additional properties files located in the directories for those profiles.

The following command saves all of the properties in the */SAS* subdirectory:

```
SAV DEV('/QSYS.lib/wsa1ib.lib/wsasavf.file')
OBJ(('profile_root/properties/sas*'))
```

This previous command is on two lines for illustrative purposes only. Enter it as one continuous line

You can save security property files while WebSphere Application Server is running.

- Back up your HTTP configuration.

The following information applies to IBM HTTP Server. If you are using Lotus Domino HTTP Server, see the Notes®.net Documentation Library.

Changes to the HTTP configuration are often made to enable WebSphere Application Server to serve servlets and JavaServer Pages (JSP) file requests and to enable WebSphere Application Server security. Consider saving your HTTP configuration as a part of your WebSphere Application Server backup and recovery. The IBM HTTP Server configurations are stored as members of the QATMHTTTPC file in the QUSRSYS library. HTTP server instances are members of the QATMHINSTC file in the QUSRSYS library. The following example commands back up these files:

```
SAVOBJ OBJ(QUSRSYS/QATMHTTTPC)
```

```
SAVOBJ OBJ(QUSRSYS/QATMHINSTC)
```

- Back up your key files.

The key files contain certificates that are used by the security infrastructure for WebSphere Application Server. These certificates are also used for HTTPS transport between servers. Save all of the files in the *profile_root/etc* directory. Key files are contained in the *profile_root/etc* directory, but administrators might create and store these files in other directories.

- Back up your validation lists.

Passwords are stored as encrypted data in validation list objects when you use the OS/400® password encoding algorithm. The default validation list is */QSYS.LIB/QUSRSYS.LIB/EJSADMIN.VLDL*, but you can change it in the administrative console by specifying it as a system property for the application server. See the information about administering application servers.

What to do next

For more information on backup and recovery strategies, see the following references:

- iSeries Information Center
- iSeries Backup and Recovery Version 5
- Security: Resources for Learning

Chapter 15. Tuning Session Initiation Protocol (SIP) applications

This page provides a starting point for finding information about SIP applications, which are Java programs that use at least one Session Initiation Protocol (SIP) servlet written to the JSR 116 specification.

SIP is used to establish, modify, and terminate multimedia IP sessions including IP telephony, presence, and instant messaging.

Tuning your environment to run SIP servlets

You might want to implement one or more of the following tuning suggestions if you are encountering performance issues while running SIP servlets.

About this task

The following recommendations are provided as is without any implied warranty.

Procedure

- Make the following changes to your Java virtual machine (JVM) settings.
To make these changes, in the administrative console, click **Servers > Server Types > WebSphere application servers > *server_name* > Java and Process Management > Process Definition > Java Virtual Machine**.
 1. Select **Enable Verbose garbage collection**.
 2. Set **Initial heap size** to 1284.
 3. Set **Maximum heap size** to 1284.
 4. Specify `-Xmn150m -Xgcpolicy:gencon -Xtgc:parallel -Xgc:scvNoAdaptiveTenure,scvTenureAge=1,stdGlobalCompactToSatisfyAllocate -Xdump:heap:events=user,request=exclusive+prepwalk+compact -Xcompactexplicitgc -XX:MaxDirectMemorySize=256000000 -Xdump:heap:events=user,request=exclusive+prepwalk+compact` in the **Generic JVM arguments** field.
 5. Change the following JVM log settings for the application server.
To make these changes, in the administrative console, click **Servers > Server Types > WebSphere application servers > *server_name* > Java and Process Management > Process Definition > Logging and Tracing > JVM Logs**.
 - Under System.out, specify 20 for the File size maximum Size property and 2 for the Maximum number of historical log files property.
 - Under System.err, specify 20 for the File size maximum Size property and 2 for the Maximum number of historical log files property.
 6. Change the following node agent detection monitoring policy settings
These setting changes enable the JVM to detect if the operating system has hung.
To make these changes, in the administrative console, click **Servers > Server Types > WebSphere application servers > *server_name* > Java and Process Management > MonitoringPolicy**.
 - Set the Maximum startup attempts property to 2.
 - Set the Ping interval property to 30.
 - Set the Ping timeout property to 60.
- Create a thread pool, and specify SipContainerPool for the Name property, and 45 for both the Minimum size, and Maximum size properties.
- Make the following changes to your SIP container settings.

To make these changes, in the administrative console, click **Servers > Server Types > WebSphere application servers > server_name > SIP Container Settings > sip_container**.

1. Set the Maximum application sessions property to 36000.
2. Set the Maximum messages per averaging period property to 26640.
3. Set the Maximum dispatch queue size property to 5220.
4. Set the Averaging period property to 30000.
5. Set the Statistics update rate property to 3000.
6. Select SipContainerPool from the list of available thread pools.
7. Click **Custom Properties > New**, and add the following SIP container custom properties to your SIP container settings.

Table 7. SIP container custom property settings. This table indicates settings that you might want to specify for the listed SIP container custom properties if you need to improve the performance of your SIP servlets.

Name	Value
javax.sip.max.object.pool.size	1000
javax.sip.stat.report.interval	0
max.tu.pool.size	1000
weight.overload.watermark	1
end.of.service.replication	true
immediate.replication	false
thread.message.queue.max.size	5000
bootstrap.batch.size	1000
bootstrap.blocking.delay	500

8. From the SIP container settings page, click **SIP container transport chains > SIPInboundDefaultUDP > UDP Inbound channel (UDP 1) > Custom Properties > New**, and add the following two custom properties.

Table 8. UDP inbound channel (UDP 1) custom property settings. This table indicates settings that you might want to specify for the listed UDP inbound channel (UDP 1) custom properties if you need to improve the performance of your SIP servlets.

Name	Value
sendBufferSizeSocket	3000000
receiveBufferSizeSocket	3000000

- Make the following changes to your proxy server settings for each of your proxy servers.
To make these changes, in the administrative console, click **Servers > Server Types > Proxy servers > proxy_server_name > Proxy server settings**.
 1. Create two new ports.
To add these ports, click **Ports** on the **Proxy servers > proxy_server_name** page.
Specify the following settings for the first port:
 - Select **User-defined Port**.
 - Specify SIP_LB_ADDRESS in the **Port name** field.
 - Specify the name of the Load Balancer (cluster) address in the **Host** field.
 - Specify 5060 in the **Port** field.
 Specify the following settings for the second port:
 - Select **User-defined Port**.
 - Specify HTTP_LB_ADDRESS in the **Port name** field.
 - Specify the name of the Load Balancer (cluster) address in the **Host** field.

- Specify 80 in the **Port** field.
- 2. On the **Ports** page, change the value of the **Host** field for the DCS_UNICAST_ADDRESS port to the host name of the management network for the proxy server. This network handles WebSphere Application Server inter-process communication, such as Distribution and Consistency Services (DCS) communication.
- 3. On the **Ports** page, change the value of the **Host** field for the PROXY_HTTPS_ADDRESS, PROXY_HTTP_ADDRESS, PROXY_SIPS_ADDRESS and PROXY_SIP_ADDRESS ports to the host name of the network that is handling SIP traffic for the proxy server.
- 4. Add the following custom properties.
To add these properties, from the **Proxy server settings** page, click **SIP proxy settings > Custom Properties**.

Table 9. SIP proxy server custom property settings. This table indicates settings that you might want to specify for the listed SIP proxy server custom properties if you need to improve the performance of your SIP servlets.

Name	Value
isSipComplianceEnabled	false
serverUDPInterface	The host name of the network that is handling SIP traffic for the proxy server. This value becomes the Container Facing Interface.
serverUDPPort	5080 This port becomes the Container Facing Port.
LBIPAddr	The Load Balancer IP address. This address is not the cluster address.
SIPAdvisorMethodName	Options
maxThroughputFactor	90
perSecondBurstFactor	200
burstResetFactor	120
sipProxyStartupDelay	120000
proxyTransitionPeriod	360
minDeflatorRatio	6
maxDeflatorRatio	10
keepAliveInterval	2000
keepAliveFailures	3
clusterRouteConfigUpdateDelay	60000
localOutboundTCPAddress	The IP address of the local outbound TCP connection
localOutboundTCPPort	The port number of the local outbound TCP connection
receiveBufferSizeSocket	3000000
sendBufferSizeSocket	3000000

- 5. Delete the UDP transport chain UDP_SIP_PROXY_CHAIN
To delete the UDP transport chain UDP_SIP_PROXY_CHAIN, click **SIP proxy server settings > SIP proxy server transports** on the **Proxy server settings** page, select the UDP transport chain UDP_SIP_PROXY_CHAIN, and then click **Delete**.
- 6. Change the settings of the SIP proxy server transport.
To change the settings of the SIP proxy server transport, click **SIP proxy server settings > SIP proxy server transports > New** on the **Proxy servers > proxy_server_name** page, and then specify the following settings for the new transport:
 - Specify SIP_LB_CHAIN in the **Transport chain name** field.
 - Specify proxy in the **Transport chain template** field.

- On the next administrative console page, select **Use existing port** and **SIP_LB_ADDRESS** from the lists of available options.
7. Change the settings of the HTTP proxy server transport.

To change the settings of the SIP proxy server transport, click **SIP proxy server settings > HTTP proxy server transports > New** on the **Proxy servers > proxy_server_name** page, and then specify the following settings for the new transport:

 - Specify HTTP_LB_CHAIN in the **Transport chain name** field.
 - Specify proxy in the **Transport chain template** field.
 - On the next administrative console page, select **Use existing port** and **HTTP_LB_ADDRESS** from the lists of available options.
 8. Add the following custom properties to the UDP channel settings.

To add new custom properties to the UDP channel settings, click **SIP proxy transport chains > UDPInboundChannel > Custom Properties > New** on the **SIP proxy server** settings page.

Table 10. UDPInboundChannel custom property settings for SIP proxy transport chains. This table indicates settings that you might want to specify for the listed UDPInboundChannel custom properties for SIP proxy transport chains if you need to improve the performance of your SIP servlets.

Name	Value
sendBufferSizeSocket	3000000
receiveBufferSizeSocket	3000000

9. Add the following custom properties to the HTTP proxy server settings.

To add this property, click **Proxy server settings > Proxy settings > Custom Properties > New** on the **Proxy servers > proxy_server_name** page.

Table 11. HTTP proxy server custom property settings. This table indicates settings that you might want to specify for the listed HTTP proxy server custom properties if you need to improve the performance of your SIP servlets.

Name	Description	Value
localOutboundTCPAddress	Local outbound interface address	The proxy address of the network that is handling SIP traffic.
localOutboundTCPPort	1080	Local outbound port

10. Change the following JVM settings for the proxy server.

To make these changes, in the administrative console, click **Servers > Server Types > WebSphere proxy servers > proxy_server_name > Java and Process Management > Process Definition > Java virtual machine**.

 - Set **Initial heap size** to 300.
 - Set **Maximum heap size** to 450.
 - Add `-Xm60m;-Xgcpolicy:gencon;-Xgc:noAdaptiveTenure,tenureAge=8,stdGlobalCompactToSatisfyAllocate;-Xtgc:parallel-Xdump:heap:events=user,request=exclusive+prepwalk+compact;-Xl0a;-Xl0a:minimum0.03` to the **Generic JVM arguments** field.

If you are running over 50 proxy servers in your environment, you might also want to specify `-Xm600m` to `-Xm200m`
11. Change the following JVM log settings for the proxy server.

To make these changes, in the administrative console, click **Servers > Server Types > WebSphere proxy servers > proxy_server_name > Java and Process Management > Process Definition > Logging and Tracing > JVM Logs**.

 - Under System.out, specify 20 for the File size maximum Size property and 2 for the Maximum number of historical log files property.

- Under System.err, specify 20 for the File size maximum Size property and 2 for the Maximum number of historical log files property.
- 1. Add the IBM_CLUSTER_RUNRULES_TIMER_TIME custom property to the cluster run rules. To add this property, click **Administration > Custom Properties > New** on the **Proxy servers > proxy_server_name** page. Then specify IBM_CLUSTER_RUNRULES_TIMER_TIME in the **Name** field, and 1000 in the **Value** field.
- Set the size of the application server transport buffer and the proxy server transport buffer to 100 megabytes.

For the application server, in the administrative console, click **Servers > Server Types > WebSphere application servers > server_name**, and enter 100 for the Transport buffer size property.

For the proxy server, in the administrative console, click **Servers > Server Types > WebSphere proxy servers > server_name**, and enter 100 for the Transport buffer size property.
- Set the size of the deployment manager transport buffer to 100 megabytes.

In the administrative console, click **System administration > Deployment manager > Core group service**, and enter 100 for the Transport buffer size property.
- For each node agent, set the size of the node agent transport buffer to 100 megabytes.

In the administrative console, click **System administration > Node agents > node_agent Core group service**, and enter 100 for the Transport buffer size property.
- Make the following changes to your core group settings.
 1. Change the settings for the IBM_CS_FD_PERIOD_SECS, IBM_CS_FD_CONSECUTIVE_MISSED, and IBM_CS_DATASTACK_MEG core group custom properties.

The following values are recommended for these custom properties:

Table 12. Core group custom property settings. This table indicates settings that you might want to specify for the listed core group custom properties if you need to improve the performance of your SIP servlets.

Name	Description	Value
IBM_CS_FD_PERIOD_SECS	This custom property specifies, in seconds, how frequently the Failure Detection Protocol checks the core group network connections that the discovery protocol establishes. The Failure Detection Protocol notifies the Discovery Protocol if a connection failure occurs.	2
IBM_CS_FD_CONSECUTIVE_MISSED	This custom property specifies the consecutive number of heartbeats that must be missed before the core group member is considered failed. Do not configure this custom property unless you are running in a mixed cell environment that includes at least one core groups that contains a mixture of Version 7.0 and Version 6.x processes.	3

Table 12. Core group custom property settings (continued). This table indicates settings that you might want to specify for the listed core group custom properties if you need to improve the performance of your SIP servlets.

Name	Description	Value
IBM_CS_DATASTACK_MEG	<p>This custom property specifies, in megabytes, the amount of memory that is available for synchronization.</p> <p>Do not configure this custom property unless you are running in a mixed cell environment that includes at least one core groups that contains a mixture of Version 7.0 and Version 6.x processes.</p>	100

Chapter 16. Tuning Transactions

This page provides a starting point for finding information about Java Transaction API (JTA) support. Applications running on the server can use transactions to coordinate multiple updates to resources as one unit of work, such that all or none of the updates are made permanent.

The product provides advanced transactional capabilities to help application developers avoid custom coding. It provides support for the many challenges related to integrating existing software assets with a Java EE environment.

Chapter 17. Tuning web applications

This page provides a starting point for finding information about web applications, which are comprised of one or more related files that you can manage as a unit, including:

- HTML files
- Servlets can support dynamic web page content, provide database access, serve multiple clients at one time, and filter data.
- Java ServerPages (JSP) files enable the separation of the HTML code from the business logic in web pages.

IBM extensions to the JSP specification make it easy for HTML authors to add the power of Java technology to web pages, without being experts in Java programming. More introduction...

Tuning URL cache

Tuning URL invocation cache

The URL invocation cache holds information for mapping request URLs to servlet resources. This cache is Web container-based, and shared for all Web container threads. A cache of the requested size is created for each Web container thread that is available to process a request. The default size of the invocation cache is 50. If more than 50 unique URLs are actively being used (each JavaServer Page is a unique URL), you should increase the size of the invocation cache.

Before you begin

A larger cache uses more of the Java heap, so you might also need to increase the maximum Java heap size. For example, if each cache entry requires 2KB, maximum thread size is set to 25, and the URL invocation cache size is 100; then 5MB of Java heap are required.

About this task

To change the size of the invocation cache:

Procedure

1. In the administrative console, click **Servers > Server Types > WebSphere application servers** and select the application server that you are tuning.
2. Click **Java and Process Management**.
3. Click **Process Definition** under Additional Properties.
4. Click **Java Virtual Machine** under Additional Properties.
5. Click **Custom Properties** under Additional Properties.
6. Specify **invocationCacheSize** in the Name field and the size of the cache in the Value field. The default size for the invocation cache is 500 entries. Since the invocation cache is no longer thread-based, the invocation cache size specified by the user is multiplied by ten to provide similar function from previous releases. For example, if you specify an invocation cache size of 50, the web container will create a cache size of 500.
7. Click **Apply** and then **Save** to save your changes.
8. Stop and restart the application server.

Results

The new cache size is used for the URL invocation cache.

Tuning sessions

Session management tuning

WebSphere Application Server session support has features for tuning session performance and operating characteristics, particularly when sessions are configured in a distributed environment. These options support the administrator flexibility in determining the performance and failover characteristics for their environment.

Table 13. Summary of tuning features. The following table summarizes the tuning features, including whether they apply to sessions tracked in memory, in a database, with memory-to-memory replication, or all. Some features are easily manipulated using administrative settings; others require code or database changes.

Feature or option	Goal	Applies to sessions in memory, database, or memory-to-memory
Write frequency	Minimize database write operations.	Database and Memory-to-Memory
Session affinity	Access the session in the same application server instance.	All
Multirow schema	Fully utilize database capacities.	Database
Base in-memory session pool size	Fully utilize system capacity without overburdening system.	All
Write contents	Allow flexibility in determining what session data to write	Database and Memory-to-Memory
Scheduled invalidation	Minimize contention between session requests and invalidation of sessions by the Session Management facility. Minimize write operations to database for updates to last access time only.	Database and Memory-to-Memory
Tablespace and row size	Increase efficiency of write operations to database.	Database (DB2 only)

Tuning parameter settings

Use this page to set tuning parameters for distributed sessions.

To view this administrative console page, click **Servers > Server Types > WebSphere application servers > *server_name* > Session management > Distributed environment settings > Custom tuning parameters**.

Tuning level

Specifies that the session management facility provides certain predefined settings that affect performance.

Select one of these predefined settings or customize a setting. To customize a setting, select one of the predefined settings that comes closest to the setting desired, click **Custom settings**, make your changes, and then click **OK**.

Very high (optimize for performance)

Write frequency	Time based
Write interval	300 seconds
Write contents	Only updated attributes
Schedule sessions cleanup	true
First time of day default	0
Second time of day default	2

High

Write frequency	Time based
Write interval	300 seconds
Write contents	All session attributes
Schedule sessions cleanup	false

Medium

Write frequency	End of servlet service
Write contents	Only updated attributes
Schedule sessions cleanup	false

Low (optimize for failover)

Write frequency	End of servlet service
Write contents	All session attributes
Schedule sessions cleanup	false

Custom settings

Write frequency default	Time based
Write interval default	10 seconds
Write contents default	All session attributes
Schedule sessions cleanup default	false

Tuning parameter custom settings

Use this page to customize tuning parameters for distributed sessions.

To view this administrative console page, click **Servers > Server Types > WebSphere application servers > *server_name* > Session management > Distributed environment settings > Custom tuning parameters > Custom settings**.

Write frequency

Specifies when the session is written to the persistent store.

End of servlet service	A session writes to a database or another WebSphere Application Server instance after the servlet completes execution.
Manual update	A programmatic sync on the IBMSession object is required to write the session data to the database or another WebSphere Application Server instance.
Time based	Session data writes to the database or another WebSphere Application Server instance based on the specified Write interval value. Default: 10 seconds

Write contents

Specifies whether updated attributes are only written to the external location or all of the session attributes are written to the external location, regardless of whether or not they changed. The external location can be either a database or another application server instance.

Only updated attributes
All session attribute

Only updated attributes are written to the persistent store.
All attributes are written to the persistent store.

Schedule sessions cleanup

Specifies when to clean the invalid sessions from a database or another application server instance.

Specify distributed sessions cleanup schedule

Enables the scheduled invalidation process for cleaning up the invalidated HTTP sessions from the external location. Enable this option to reduce the number of updates to a database or another application server instance required to keep the HTTP sessions alive. When this option is not enabled, the invalidator process runs every few minutes to remove invalidated HTTP sessions.

When this option is enabled, specify the two hours of a day for the process to clean up the invalidated sessions in the external location. Specify the times when there is the least activity in the application servers. An external location can be either a database or another application server instance.

First Time of Day (0 - 23)

Indicates the first hour during which the invalidated sessions are cleared from the external location. Specify this value as a positive integer between 0 and 23. This value is valid only when schedule invalidation is enabled.

Second Time of Day (0 - 23)

Indicates the second hour during which the invalidated sessions are cleared from the external location. Specify this value as a positive integer between 0 and 23. This value is valid only when schedule invalidation is enabled.

Best practices for using HTTP sessions

This topic presents best practices for the implementation of HTTP sessions.

best-practices: Browse the following recommendations for implementing HTTP sessions.

- **Enable Security integration for securing HTTP sessions**

HTTP sessions are identified by session IDs. A session ID is a pseudo-random number generated at the runtime. Session hijacking is a known attack HTTP sessions and can be prevented if all the requests going over the network are enforced to be over a secure connection (meaning, HTTPS). But not every configuration in a customer environment enforces this constraint because of the performance impact of SSL connections. Due to this relaxed mode, HTTP session is vulnerable to hijacking and because of this vulnerability, WebSphere Application Server has the option to tightly integrate HTTP sessions and WebSphere Application Server security. Enable security in WebSphere Application Server so that the sessions are protected in a manner that only users who created the sessions are allowed to access them.

- **Release HttpSession objects using `javax.servlet.http.HttpSession.invalidate()` when finished.**

HttpSession objects live inside the web container until:

- The application explicitly and programmatically releases it using the `javax.servlet.http.HttpSession.invalidate` method; quite often, programmatic invalidation is part of an application logout function.
- WebSphere Application Server destroys the allocated HttpSession when it expires (default = 1800 seconds or 30 minutes). The WebSphere Application Server can only maintain a certain number of HTTP sessions in memory based on session management settings. In case of distributed sessions, when maximum cache limit is reached in memory, the session management facility removes the least recently used (LRU) one from cache to make room for a session.

- **Avoid trying to save and reuse the HttpSession object outside of each servlet or JSP file.**

The HttpSession object is a function of the HttpServletRequest (you can get it only through the req.getSession method), and a copy of it is valid only for the life of the service method of the servlet or JSP file. You *cannot* cache the HttpSession object and refer to it outside the scope of a servlet or JSP file.

- **Implement the java.io.Serializable interface when developing new objects to be stored in the HTTP session.**

Serializability of a class is enabled by the class implementing the java.io.Serializable interface. Implementing the java.io.Serializable interface allows the object to properly serialize when using distributed sessions. Classes that do not implement this interface will not have their states serialized or deserialized. Therefore, if a class does not implement the Serializable interface, the JVM cannot persist its state into a database or into another JVM. All subtypes of a serializable class are serializable. An example of this follows:

```
public class MyObject implements java.io.Serializable {...}
```

Make sure all instance variable objects that are not marked transient are serializable. You cannot cache a non-serializable object.

In compliance with the Java Servlet specification, the distributed servlet container must create an IllegalArgumentException for objects when the container cannot support the mechanism necessary for migration of the session storing them. An exception is created only when you have selected distributable.

- **The HttpSession API does not dictate transactional behavior for sessions.**

Distributed HttpSession support does not guarantee transactional integrity of an attribute in a failover scenario or when session affinity is broken. Use transactional aware resources like enterprise Java beans to guarantee the transaction integrity required by your application.

- **Ensure the Java objects you add to a session are in the correct class path.**

If you add Java objects to a session, place the class files for those objects in the correct class path (the application class path if utilizing sharing across web modules in an enterprise application, or the web module class path if using the Servlet 2.2-complaint session sharing) or in the directory containing other servlets used in WebSphere Application Server. In the case of session clustering, this action applies to every node in the cluster.

Because the HttpSession object is shared among servlets that the user might access, consider adopting a site-wide naming convention to avoid conflicts.

- **Avoid storing large object graphs in the HttpSession object.**

In most applications each servlet only requires a fraction of the total session data. However, by storing the data in the HttpSession object as one large object, an application forces WebSphere Application Server to process all of it each time.

- **Utilize Session Affinity to help achieve higher cache hits in the WebSphere Application Server.**

WebSphere Application Server has functionality in the HTTP Server plug-in to help with session affinity. The plug-in reads the cookie data (or encoded URL) from the browser and helps direct the request to the appropriate application or clone based on the assigned session key. This functionality increases use of the in-memory cache and reduces hits to the database or another WebSphere Application Server instance

- **Maximize use of session affinity and avoid breaking affinity.**

Using session affinity properly can enhance the performance of the WebSphere Application Server. Session affinity in the WebSphere Application Server environment is a way to maximize the in-memory cache of session objects and reduce the amount of reads to the database or another WebSphere Application Server instance. Session affinity works by caching the session objects in the server instance of the application with which a user is interacting. If the application is deployed in multiple servers of a server group, the application can direct the user to any one of the servers. If the users starts on server1 and then comes in on server2 a little later, the server must write all of the session information to the external location so that the server instance in which server2 is running can read the database. You can avoid this database read using session affinity. With session affinity, the user starts on server1 for the first request; then for every successive request, the user is directed back to server1. Server1 has to look only at the cache to get the session information; server1 never has to make a call to the session database to get the information.

You can improve performance by not breaking session affinity. Some suggestions to help avoid breaking session affinity are:

- Combine all web applications into a single application server instance, if possible, and use modeling or cloning to provide failover support.
- Create the session for the frame page, but do not create sessions for the pages within the frame when using multi-frame JSP files. (See discussion later in this topic.)
- **When using multi-framed pages, follow these guidelines:**
 - Create a session in only one frame or before accessing any frame sets. For example, assuming there is no session already associated with the browser and a user accesses a multi-framed JSP file, the browser issues concurrent requests for the JSP files. Because the requests are not part of any session, the JSP files end up creating multiple sessions and all of the cookies are sent back to the browser. The browser honors only the last cookie that arrives. Therefore, only the client can retrieve the session associated with the last cookie. Creating a session before accessing multi-framed pages that utilize JSP files is recommended.
 - By default, JSP files get a `HTTPSession` using `request.getSession(true)` method. So by default JSP files create a new session if none exists for the client. Each JSP page in the browser is requesting a new session, but only one session is used per browser instance. A developer can use `<% @ page session="false" %>` to turn off the automatic session creation from the JSP files that do not access the session. Then if the page needs access to the session information, the developer can use `<% HttpSession session = javax.servlet.http.HttpServletRequest.getSession(false); %>` to get the already existing session that was created by the original session creating JSP file. This action helps prevent breaking session affinity on the initial loading of the frame pages.
 - Update session data using only one frame. When using framesets, requests come into the HTTP server concurrently. Modifying session data within only one frame so that session changes are not overwritten by session changes in concurrent frameset is recommended.
 - Avoid using multi-framed JSP files where the frames point to different web applications. This action results in losing the session created by another web application because the `JSESSIONID` cookie from the first web application gets overwritten by the `JSESSIONID` created by the second web application.
- **Secure all of the pages (not just some) when applying security to servlets or JSP files that use sessions with security integration enabled, .**

When it comes to security and sessions, it is all or nothing. It does not make sense to protect access to session state only part of the time. When security integration is enabled in the session management facility, all resources from which a session is created or accessed must be either secured or unsecured. You cannot mix secured and unsecured resources.

The problem with securing only a couple of pages is that sessions created in secured pages are created under the identity of the authenticated user. Only the same user can access sessions in other secured pages. To protect these sessions from use by unauthorized users, you cannot access these sessions from an unsecured page. When a request from an unsecured page occurs, access is denied and an `UnauthorizedSessionRequestException` error is created. (`UnauthorizedSessionRequestException` is a runtime exception; it is logged for you.)

- **Use manual update and either the `sync()` method or time-based write in applications that read session data, and update infrequently.**

With `END_OF_SERVICE` as write frequency, when an application uses sessions and anytime data is read from or written to that session, the `LastAccess` time field updates. If database sessions are used, a new write to the database is produced. This activity is a performance hit that you can avoid using the Manual Update option and having the record written back to the database only when data values update, not on every read or write of the record.

To use manual update, turn it on in the session management service. (See the tables above for location information.) Additionally, the application code must use the `com.ibm.websphere.servlet.session.IBMSession` class instead of the generic `HttpSession`. Within the `IBMSession` object there is a `sync` method. This method tells the WebSphere Application Server to write the data in the session object to the database. This activity helps the developer to improve overall performance by having the session information persist only when necessary.

Note: An alternative to using the manual updates is to utilize the timed updates to persist data at different time intervals. This action provides similar results as the manual update scheme.

- Implement the following suggestions to achieve high performance:
 - If your applications do not change the session data frequently, use Manual Update and the sync function (or timed interval update) to efficiently persist session information.
 - Keep the amount of data stored in the session as small as possible. With the ease of using sessions to hold data, sometimes too much data is stored in the session objects. Determine a proper balance of data storage and performance to effectively use sessions.
 - If using database sessions, use a dedicated database for the session database. Avoid using the application database. This helps to avoid contention for JDBC connections and allows for better database performance.
 - If using memory-to-memory sessions, employ partitioning (either group or single replica) as your clusters grow in size and scaling decreases.
 - Verify that you have the latest fix packs for the WebSphere Application Server.
- Utilize the following tools to help monitor session performance.
 - Run the `com.ibm.servlet.personalization.sessiontracking.IBMTrackerDebug` servlet. - To run this servlet, you must have the servlet invoker running in the web application you want to run this from. Or, you can explicitly configure this servlet in the application you want to run.
 - Use the WebSphere Application Server Resource Analyzer which comes with WebSphere Application Server to monitor active sessions and statistics for the WebSphere Application Server environment.
 - Use database tracking tools such as "Monitoring" in DB2. (See the respective documentation for the database system used.)

Chapter 18. Tuning web services

This page provides a starting point for finding information about web services.

Web services are self-contained, modular applications that can be described, published, located, and invoked over a network. They implement a services oriented architecture (SOA), which supports the connecting or sharing of resources and data in a very flexible and standardized manner. Services are described and organized to support their dynamic, automated discovery and reuse.

Tuning Web Services Security

When using Web Services Security for message-level protection of SOAP message in WebSphere Application Server, the choice of configuration options can affect the performance of the application.

Tuning Web Services Security for Version 8.0 applications

The Java Cryptography Extension (JCE) is integrated into the software development kit (SDK) Version 1.4.x and later. This is no longer an optional package. However, the default JCE jurisdiction policy file shipped with the SDK enables you to use cryptography to enforce this default policy. In addition, you can modify the web services security configuration options to achieve the best performance for web services security protected applications.

About this task

Using the unrestricted JCE policy files

Due to export and import regulations, the default JCE jurisdiction policy file shipped with the SDK enables you to use strong, but limited, cryptography only. To enforce this default policy, WebSphere Application Server uses a JCE jurisdiction policy file that might introduce a performance impact. The default JCE jurisdiction policy might have a performance impact on the cryptographic functions that are supported by Web Services Security. If you have web services applications that use transport level security for XML encryption or digital signatures, you might encounter performance degradation over previous releases of WebSphere Application Server. However, IBM and Sun Microsystems provide versions of these jurisdiction policy files that do not have restrictions on cryptographic strengths. If you are permitted by your governmental import and export regulations, download one of these jurisdiction policy files. After downloading one of these files, the performance of JCE and Web Services Security might improve.

Attention: Fix packs that include updates to the Software Development Kit (SDK) might overwrite unrestricted policy files. Back up unrestricted policy files before you apply a fix pack and reapply these files after the fix pack is applied.

Important: Your country of origin might have restrictions on the import, possession, use, or re-export to another country, of encryption software. Before downloading or using the unrestricted policy files, you must check the laws of your country, its regulations, and its policies concerning the import, possession, use, and re-export of encryption software, to determine if it is permitted.

For WebSphere Application Server platforms using IBM Developer Kit, Java Technology Edition Version 6, you can obtain unlimited jurisdiction policy files by completing the following steps:

1. Go to the following website: <http://www.ibm.com/developerworks/java/jdk/security/index.html>
2. Click **Java SE 6**
3. Scroll down and click **IBM SDK Policy files**.

The Unrestricted JCE Policy files for the SDK website is displayed.

4. Click **Sign in** and provide your IBM intranet ID and password or register with IBM to download the files.

5. Select the appropriate Unrestricted JCE Policy files and then click **Continue**.
6. View the license agreement and then click **I Agree**.
7. Click **Download Now**.

For IBM i (both 5.4 and 6.1) and the IBM Software Development Kit, Java Technology Edition Version 6, the restricted JCE jurisdiction policy files are configured, by default. You can download the unrestricted JCE jurisdiction policy files from the following website: <http://www.ibm.com/developerworks/java/jdk/security/60>:

Note: If Java Platform, Standard Edition 6 (Java SE 6) 32-bit for IBM i is the enabled Java virtual machine (JVM) for your profile, substitute /QOpenSys/QIBM/ProdData/JavaVM/jdk50/32bit/jre for /QIBM/ProdData/Java400/jdk6 as the path name in the following steps.

To configure the unrestricted jurisdiction policy files for IBM i and the IBM Software Development Kit, complete the following steps:

Procedure

1. Make backup copies of these files:

```
/QIBM/ProdData/Java400/jdk6/lib/security/local_policy.jar
/QIBM/ProdData/Java400/jdk6/lib/security/US_export_policy.jar
```

2. Download the unrestricted policy files from <http://www.ibm.com/developerworks/java/jdk/security/index.html> to the /QIBM/ProdData/Java400/jdk6/lib/security directory.
3. Go to the following website: <http://www.ibm.com/developerworks/java/jdk/security/index.html>
 - a. Click **Java SE 6**
 - b. Scroll down and click **IBM SDK Policy files**. The Unrestricted JCE Policy files for the SDK web site is displayed.
 - c. Click **Sign in** and provide your IBM intranet ID and password.
 - d. Select the appropriate Unrestricted JCE Policy files and then click **Continue**.
 - e. View the license agreement and then click **I Agree**.
 - f. Click **Download Now**.
4. Use the DSPAUT command to ensure that *PUBLIC is granted *RX data authority but that object authority is not provided to either of the local_policy.jar and US_export_policy.jar files, which are located in the /QIBM/ProdData/Java400/jdk6/lib/security directory. For example:

```
DSPAUT OBJ('/qibm/proddata/java400/jdk6/lib/security/local_policy.jar')
```

5. Use the CHGAUT command to change authorization, if needed. For example:

```
CHGAUT OBJ('/qibm/proddata/java400/jdk6/lib/security/local_policy.jar')
USER(*PUBLIC) DTAAUT(*RX) OBJAUT(*NONE)
```

Example

Using configuration options to tune WebSphere Application Server

When using WS-Security for message-level protection of SOAP message in WebSphere Application Server, the choice of configuration options can affect the performance of the application. The following guidelines will help you achieve the best performance for your WS-Security protected applications.

1. Use WS-SecureConversation when appropriate for JAX-WS applications. The use of symmetric keys with a Secure Conversation typically performs better than asymmetric keys used with X.509.

Note: The use of WS-SecureConversation is supported for JAX-WS applications only, not JAX-RPC applications.

2. Use the standard token types provided by WebSphere Application Server. Use of custom tokens is supported, but higher performance is achieved with the use of the provided token types.

3. For signatures, use only the exclusive canonicalization transform algorithm. See the W3 Recommendation web page (<http://www.w3.org/2001/10/xml-exc-c14n#>) for more information.
4. Whenever possible, avoid the use of the XPath expression to select which SOAP message parts to protect. The WS-Security policies shipped with WebSphere Application Server for JAX-WS applications use XPath expressions to specify the protection of some elements in the security header, such as Timestamp, SignatureConfirmation, and UsernameToken. The use of these XPath expressions is optimized, but other uses are not.
5. Although there are Websphere Application Server extensions to WS-Security that can be used to insert nonce and timestamp elements into SOAP message parts before signing or encrypting the message parts, you should avoid the use of these extensions for improved performance.
6. There is an option to send the base-64 encoded CipherValue of WS-Security encrypted elements as MTOM attachments. For small encrypted elements, the best performance is achieved by avoiding this option. For larger encrypted elements, the best performance is achieved by using this option.
7. When signing and encrypting elements in the SOAP message, specify the order as sign first, then encrypt.
8. When adding a timestamp element to a message, the timestamp should be added to the security header before the signature element. This is accomplished by using the Strict or LaxTimestampFirst security header layout option in the WS-Security policy configuration.
9. For JAX-WS applications, use the policy-based configuration rather than WSS API-based configuration.

What to do next

In IBM WebSphere Application Server Version 6.1 and later, Web Services Security supports the use of cryptographic hardware devices. There are two ways in which to use hardware cryptographic devices with Web Services Security. See [Hardware cryptographic device support for Web Services Security](#) for more information.

Tuning Web Services Security for Version 5.x applications

The Java Cryptography Extension (JCE) policy is integrated into the IBM Software Development Kit (SDK) Version 1.4.x and is no longer an optional package. However, due to export and import regulations, the default JCE jurisdiction policy file shipped with the SDK enables you to use strong, but limited, cryptography only.

About this task

To enforce this default policy, WebSphere Application Server uses a JCE jurisdiction policy file that might introduce a performance impact. The default JCE jurisdiction policy might have a performance impact on the cryptographic functions that are supported by Web Services Security. If you have web services applications that use transport level security for XML encryption or digital signatures, you might encounter performance degradation over previous releases of WebSphere Application Server. However, IBM and Sun Microsystems provide versions of these jurisdiction policy files that do not have restrictions on cryptographic strengths. If you are permitted by your governmental import and export regulations, download one of these jurisdiction policy files. After downloading one of these files, the performance of JCE and Web Services Security might improve.

Procedure

For IBM i and IBM Software Development Kit Version 1.4, the tuning of Web Services Security is not required. The unrestricted jurisdiction policy files for IBM Software Development Kit Version 1.4 are automatically configured when the prerequisite software is installed.

1. For IBM i (formerly known as IBM i V5R3), the unrestricted jurisdiction policy files for IBM Software Development Kit Version 1.4 are automatically configured by installing product 5722AC3, Crypto Access Provider 128-bit.

2. For IBM i 5.4, the unrestricted jurisdiction policy files for IBM Software Development Kit Version 1.4 are automatically configured by installing product 5722SS1 Option 3, Extended Base Directory Support.

Results

After following either of these sets of steps, two Java Archive (JAR) files are placed in the JVM directory.

Tuning web services reliable messaging applications

Modifying certain settings, such as heap size, can help to improve the performance of your system.

Procedure

1. Tune the heap size of the JVM. Increasing the heap size, for example to 1536, can avoid out-of-memory errors, especially in the following situations:
 - When you are composing WS-ReliableMessaging with WS-Notification.
 - When you are using a Reliable Asynchronous Message Profile (WS-I RSP) policy set.
 - When you are using WS-ReliableMessaging in situations where a lot of recovery is required.

For more information about tuning the JVM, refer to Tuning the IBM virtual machine for Java.

2. If your application uses a managed quality of service, tune the `sib.wsrn.tokenLockTimeout` property of the messaging engine that is specified in the policy binding for the application. See the service integration custom properties information to learn more about this property, and why you might want to change its value.

Tuning bus-enabled web services

You can use the administrative console or a Jacl script to tune performance settings for service integration bus-enabled web services.

About this task

Bus-enabled web services dynamically use a fast-path route through the bus where possible. This fast-path route is used if the following criteria are met:

- The inbound port and outbound port for the service are on the same server.
- There are no mediations on the path from the inbound port to the outbound port.

If you migrate web services from the WebSphere Application Server Version 5.1 Web services gateway, then your messages use this fast-path route through the bus.

Further optimizations can be made, if your configuration also meets the following criteria:

- The inbound template WSDL URI is the same location as the Outbound Target Service WSDL location URI.
- The inbound service template WSDL service name matches the outbound WSDL service name.
- The inbound service template port name matches the outbound WSDL port name.
- The mapping of the namespaces is disabled (that is, you have set the inbound service property `com.ibm.websphere.wsgw.mapSoapBodyNamespace` to `false`).
- Operation-level security is not enabled on the outbound service.

If your web services use the fast-path route, you need not tune mediations or the service integration bus. However it is good practise to do so, because a typical environment will have at least one non-fast-path (for example, mediated) service.

To improve the performance of bus-enabled web services you can tune the following parameters:

- The Java virtual machine heap size. This helps ensure there is enough memory available to process large messages, or messages with large attachments.
- The maximum number of instances of a message-driven bean that are permitted by the activation specification for the service integration technologies resource adapter. This throttles the number of concurrent clients serviced.
- The maximum batch size for batches of messages to be delivered to a client. By default, only a single message is delivered to a message-driven bean instance at one time; you can improve performance by allowing messages to be sent in batches to a message-driven bean.
- The number of threads available to service requests for each client. That is, the number of threads available in the default thread pool, the web container thread pool and the mediation thread pool for a given application server.
- The number of threads available in the mediation thread pool. This assumes that your mediations use concurrent support where appropriate, as explained in *Concurrent mediations*.

If you have mediations that act on SOAP headers, you can improve performance by inserting the associated header schemas (.xsd files) into the SDO repository.

To tune bus-enabled web services, complete one of the following two steps:

- Use the administrative console to tune bus-enabled web services, or
- Use a Jacl script to tune bus-enabled web services.

If you have mediations that act on SOAP headers, also complete the following step:

- Insert the header schemas into the SDO repository.

Procedure

- Optional: To use the administrative console to tune bus-enabled web services, complete the following steps:
 1. Use the topic *Tuning the IBM virtual machine for Java* to set the JVM heap size to a larger value than the default value (256 megabytes). The value should generally be as large as possible without incurring paging.
 2. Use the topic *Tuning service integration messaging* to tune the maximum number of instances of a message-driven bean, the maximum batch size for batches of messages for a bean, and the number of threads available to service requests for a bean.
 3. Use the topic *Tuning the application serving environment* to tune the general application serving environment, in particular the size of the web container thread pool. In a server that is exclusively serving requests to bus-enabled web services, the default thread pool and the web container thread pool should be the same size.
 4. Use the topic *Configuring the mediation thread pool* to configure the number of threads available to concurrent mediations.
- Optional: If you have mediations that act on SOAP headers, insert the associated schemas (.xsd files) into the SDO repository as described in “Including SOAP header schemas in the SDO repository.”

Including SOAP header schemas in the SDO repository

Use this task to improve mediation performance by inserting the SOAP header schema into the SDO repository.

About this task

Mediations accessing SOAP headers should ensure that the SOAP header schema is made available to the SDO repository. This simplifies access to the header fields (see *Web Services code example*) and can provide a significant performance benefit. Usually the schema (.xsd file) for a SOAP header is already available to the application developer.

Here is an example of a header (used for routing) that is passed in the SOAP message:

```
<soapenv:Header>
<hns0:myClientToken xmlns:hns0="http://www.ibm.com/wbc">
  <UseRoutingId>true</ UseRoutingId >
  <RoutingID>5</ RoutingID >
</hns0: myClientToken >
</soapenv:Header>
```

Here is an example of an associated header schema:

```
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema"
  targetNamespace="http://www.ibm.com/wbc"
  elementFormDefault="unqualified">
<xs:element name=" myClientToken">
  <xs:complexType>
    <xs:sequence>
      <xs:element name="UseRoutingId" type="xs:string"/>
      <xs:element name="RoutingID" type="xs:string"/>
    </xs:sequence>
  </xs:complexType>
</xs:element>
</xs:schema>
```

To insert the schema into the SDO repository, complete the following steps:

Procedure

1. Create a script that contains the following code:

- For Jython, create a script called `sdoXSDImport.py`:

```
#
xsdFile=sys.argv[0]
xsdKey=sys.argv[1]
sdoRep=AdminControl.queryNames("*,type=SdoRepository,node=%s" % AdminControl.
getNode)
print AdminControl.invoke(sdoRep , importResource([xsdKey , xsdFile]))
```

- For Jacl, create a script called `sdoXSDImport.jacl`:

```
#
set xsdFile [lindex $argv 0]
set xsdKey [lindex $argv 1]
set sdoRep [$AdminControl queryNames *,type=SdoRepository,node=[$AdminControl
getNode]]
puts [$AdminControl invoke $sdoRep importResource [list $xsdKey $xsdFile]]
```

Note: To create an equivalent script for removing a resource from the SDO repository, take a copy of this script and modify the final line as follows:

- Using Jython:

```
AdminControl.invoke(sdoRep , "removeResource" , [[xsdKey , "false"]])
```

- Using Jacl:

```
$AdminControl invoke $sdoRep removeResource [list $xsdKey false]
```

2. Use the `wsadmin` scripting client to insert the schema into the SDO repository by entering the following command.

Note: Run the command from within QShell.

- To use the Jython script:

```
wsadmin -lang jython -f sdoXSDImport.py your_header.xsd your_header_namespace
```

- To use the Jacl script:

```
wsadmin -f sdoXSDImport.jacl your_header.xsd your_header_namespace
```

where

- *your_header.xsd* is the name of the file that contains your header schema.
- *your_header_namespace* is the target namespace for the header. For example `http://yourCompany.com/yourNamespace`.

Chapter 19. Tuning Work area

This page provides a starting point for finding information about work areas, a WebSphere extension for improving developer productivity.

Work areas provide a capability much like that of global variables. They enable efficient sharing of information across a distributed application.

For example, you might want to add profile information as each customer enters your application. By placing this information in a work area, it is available throughout your application, eliminating the need to hand-code a solution or to read and write information to a database.

Work area service performance considerations

Work area service performance considerations

The work area service is designed to address complex data passing patterns that can quickly grow beyond convenient maintenance. A *work area* is a note pad that is accessible to any client that is capable of looking up Java Naming Directory Interface (JNDI). After a work area is established, data can be placed there for future use in any subsequent method calls to both remote and local resources.

You can utilize a work area when a large number of methods require common information or if information is only needed by a method that is significantly further down the call graph. The former avoids the need for complex parameter passing models where the number of arguments passed becomes excessive and hard to maintain. You can improve application function by placing the information in a work area and subsequently accessing it independently in each method, eliminating the need to pass these parameters from method to method. The latter case also avoids unnecessary parameter passing and helps to improve performance by reducing the cost of marshalling and de-marshalling these parameters over the Object Request Broker (ORB) when they are only needed occasionally throughout the call graph.

When attempting to maximize performance by using a work area, cache the UserWorkArea partition that is retrieved from JNDI wherever it is accessed. You can reduce the time spent looking up information in JNDI by retrieving it once and keeping a reference for the future. JNDI lookup takes time and can be costly.

Additional caching mechanisms available to a user-defined partition are defined by the configuration property, "Deferred Attribute Serialization". This mechanism attempts to minimize the number of serialization and deserialization calls. Refer to the Work area partition service article for further explanation of this configuration attribute.

The `maxSendSize` and `maxReceiveSize` configuration parameters can affect the performance of the work area. Setting these two values to 0 (zero) effectively turns off the policing of the size of context that can be sent in a work area. This action can enhance performance, depending on the number of nested work areas an application uses. In applications that use only one work area, the performance enhancement might be negligible. In applications that have a large number of nested work areas, there might be a performance enhancement. However, a user must note that by turning off this policing it is possible that an extremely large amount of data might be sent to a server.

Performance is degraded if you use a work area as a direct replacement to passing a single parameter over a single method call. The reason is that you incur more overhead than just passing that parameter between method calls. Although the degradation is usually within acceptable tolerances and scales similarly to passing parameters with regard to object size, consider degradation a potential problem before utilizing the service. As with most functional services, intelligent use of the work areas yields the best results.

The work area service is a tool to simplify the job of passing information from resource to resource, and in some cases can improve performance by reducing the overhead that is associated with a parameter passing when the information is only sparsely accessed within the call graph. Caching the instance retrieved from JNDI is important to effectively maximize performance during runtime.

Appendix. Directory conventions

References in product information to *app_server_root*, *profile_root*, and other directories imply specific default directory locations. This topic describes the conventions in use for WebSphere Application Server.

Default product locations - IBM i

These file paths are default locations. You can install the product and other components in any directory where you have write access. You can create profiles in any valid directory where you have write access. Multiple installations of WebSphere Application Server products or components require multiple locations.

app_client_root

The default installation root directory for the Application Client for IBM WebSphere Application Server is the `/QIBM/ProdData/WebSphere/AppClient/V8/client` directory.

app_client_user_data_root

The default Application Client for IBM WebSphere Application Server user data root is the `/QIBM/UserData/WebSphere/AppClient/V8/client` directory.

app_client_profile_root

The default Application Client for IBM WebSphere Application Server profile root is the `/QIBM/UserData/WebSphere/AppClient/V8/client/profiles/profile_name` directory.

app_server_root

The default installation root directory for WebSphere Application Server Network Deployment is the `/QIBM/ProdData/WebSphere/AppServer/V8/ND` directory.

java_home

Table 14. Root directories for supported Java Virtual Machines.

This table shows the root directories for all supported Java Virtual Machines (JVMs).

JVM	Directory
32-bit IBM Technology for Java	<code>/QOpenSys/QIBM/ProdData/JavaVM/jdk60/32bit</code>
64-bit IBM Technology for Java	<code>/QOpenSys/QIBM/ProdData/JavaVM/jdk60/64bit</code>

plugins_profile_root

The default Web Server Plug-ins profile root is the `/QIBM/UserData/WebSphere/Plugins/V8/webserver/profiles/profile_name` directory.

plugins_root

The default installation root directory for Web Server Plug-ins is the `/QIBM/ProdData/WebSphere/Plugins/V8/webserver` directory.

plugins_user_data_root

The default Web Server Plug-ins user data root is the `/QIBM/UserData/WebSphere/Plugins/V8/webserver` directory.

product_library

product_lib

This is the product library for the installed product. The product library for each Version 8.0 installation on the system contains the program and service program objects (similar to `.exe`, `.dll`, `.so` objects) for the installed product. The product library name is `QWAS8x` (where `x` is A, B, C, and so on). The product library for the first WebSphere Application Server Version 8.0 product installed on the system is `QWAS8A`. The `app_server_root/properties/product.properties` file contains the value for the product library of the installation, was `.install.library`, and is located under the `app_server_root` directory.

profile_root

The default directory for a profile named *profile_name* for WebSphere Application Server Network Deployment is the `/QIBM/UserData/WebSphere/AppServer/V8/ND/profiles/profile_name` directory.

shared_product_library

The shared product library, which contains all of the objects shared by all installations on the system, is QWAS8. This library contains objects such as the product definition, the subsystem description, the job description, and the job queue.

user_data_root

The default user data directory for WebSphere Application Server Network Deployment is the `/QIBM/UserData/WebSphere/AppServer/V8/ND` directory.

The profiles and profileRegistry subdirectories are created under this directory when you install the product.

web_server_root

The default web server path is `/www/web_server_name`.

Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, is the user's responsibility.

APACHE INFORMATION. This information may include all or portions of information which IBM obtained under the terms and conditions of the Apache License Version 2.0, January 2004. The information may also consist of voluntary contributions made by many individuals to the Apache Software Foundation. For more information on the Apache Software Foundation, please see <http://www.apache.org>. You may obtain a copy of the Apache License at <http://www.apache.org/licenses/LICENSE-2.0>.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Intellectual Property & Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
USA

Trademarks and service marks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. For a current list of IBM trademarks, visit the IBM Copyright and trademark information Web site (www.ibm.com/legal/copytrade.shtml).

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

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.

Other company, product, or service names may be trademarks or service marks of others.

Index

C

commands
 command reference
 PropFilePasswordEncoder 82

D

data sources
 tuning 7
directory
 installation
 conventions 12, 117

E

encoding algorithms
 changing OS/400 to XOR 85
 custom OS/400 password encoding 83
Enterprise JavaBeans (EJB)
 tuning 20

F

files
 security 78

H

HTTP session management
 tuning 100
HTTP sessions
 best practices 102
 tuning 100

J

JPA
 tuning 23

O

objects
 security 78
 validation list replacement 85
 validation list restoration 85
ORB
 tuning 41

P

password encoding 79
 manual encoding 81
password encryption
 disablement 88
 enablement 86
passwords
 encoding 76
 encryption 76
 securing passwords 76
properties files
 manual password encoding 81

S

security
 configuration tuning 67
 performance tuning 73
security configurations
 back-up files 89
 hardening considerations 67
 maintenance considerations 67
 tuning considerations 67
security hardening
 configurations 73
 enablement 74
 migration 74
SIP
 servlets 91
SSL
 performance tips 71

U

URL cache
 tuning 99
URL invocation cache
 tuning 99

W

web services security
 tuning 107
 Version 5.x applications 109