




## Performance Tuning Guide





## Performance Tuning Guide

**Note**

Before using this information and the product it supports, read the information in "Notices" on page 61.

# Contents

Revision log.....	7
About this document .....	10
Tuning tip organization .....	10
Parameter settings .....	10
Assumptions .....	10
Recommended documentation .....	10
Application servers and Java Virtual Machine documentation .....	10
Database documentation .....	11
Access IBM FileNet documentation, compatibility matrices, and fix packs.....	12
Contact customer support .....	12
Feedback.....	12
Databases .....	13
Indexes .....	13
Memory and disk I/O .....	13
Compute statistics for optimal performance .....	14
Oracle .....	14
Disable hash join to speed up slow searches.....	14
QUEUESIZE parameter: Maximum number of concurrent connections .....	15
Microsoft SQL Server .....	16
Reclustering the docversion table.....	16
Application servers.....	17
JVM tuning.....	17
Heap size .....	17
Content Engine .....	18
Databases .....	18
Indexes – Overview .....	18
Indexes – Indexing multi-value property tables .....	19
Oracle and DB2 .....	21
DB2 .....	21
Full-text Indexes for content-based retrieval (CBR) .....	27
Indexing threads per collection parameter.....	27
Concurrent collections indexed threading parameter .....	27
Indexing Optimization Interval parameter .....	27
Auto-administration parameter .....	28
Centera FCD content migration throughput .....	31
Queries .....	31
Enumeration Batch Size parameter .....	31
Default Query Batch Size parameters .....	31
CBR searches.....	32
CBR temp table batching.....	32
Auditing.....	33
Delete old records from the Event table .....	33
Improve audit logging and search performance .....	33
Object stores .....	34
InitialCacheSize parameter.....	34
Folders, documents, and objects .....	34
Recommendations for handling large numbers of folders, documents, and objects.....	34
Business Process Management .....	40
Process Engine .....	40
Databases.....	40
JVM tuning .....	41
Process Task Manager.....	42
Use Loadstatus to monitor server load.....	45
PE Event Log purge best practice .....	45

Process Simulator (PS) and Process Analyzer (PA).....	46
Max Simulation Jobs parameter: Concurrent simulation runs.....	46
Tuning Process Analyzer for optimal performance.....	46
Databases.....	46
Publishing interval.....	47
Application Engine .....	48
Operating systems – Solaris .....	48
Processing time .....	48
Operating systems – Windows.....	49
Processing time .....	49
Workplace.....	50
Databases.....	50
Indexing common to all database.....	50
Searches that return large result sets.....	51
Folder browsing .....	51
Documents and custom objects.....	51
Java client cache settings.....	51
eForms .....	52
Rendering large templates to PDF.....	52
Email Manager .....	53
Email Manager .....	53
Databases .....	53
Indexing common to all databases .....	53
Records Manager.....	54
Databases .....	54
Indexing common to all database.....	54
Indexing for disposition and hold sweeps.....	55
Indexing for improved Property page performance .....	55
Set the minimum and maximum heap size for Dynamic Hold Sweeps .....	57
Build a DLL instead of VBScript for a Record Declaration event action module .....	58
Oracle.....	58
Content Federation Services (CFS).....	59
Federating documents.....	59
QueryBlockSize parameter: Size of query results.....	59
Batch Size parameter: Number of batched documents.....	59
CFS-IICE rules and application server memory .....	59
Web client time-out setting for viewing large federated documents.....	60
RMI bridges .....	60
JVM heap size for CFS-CS.....	60
Notices .....	61
Trademarks .....	62
U.S. Patents Disclosure.....	62

## Revision log

Date	Revision
04/06/09	<p><b>Added or updated</b></p> <ul style="list-style-type: none"> <li>• <a href="#">Suggested value for the enumeration batch size</a></li> <li>• <a href="#">Enumeration batch size and batch query size topic into two topics</a></li> <li>• <a href="#">Index information to improve queries for multi-value properties when query has a wildcard value</a></li> </ul>
12/10/08	<p><b>Added or updated</b></p> <ul style="list-style-type: none"> <li>• <a href="#">Reclustering the docversion table on Microsoft SQL Server topic</a></li> <li>• <a href="#">Full-text indexing for content-based retrieval (CBR)</a></li> <li>• <a href="#">DB2 tuning parameters</a></li> <li>• <a href="#">Managing more than 200 object stores in a single domain</a></li> <li>• <a href="#">RMI garbage collection triggers that affect PE</a></li> <li>• <a href="#">Heap values for Component Manager activities</a></li> </ul>
07/15/08	<p><b>Added or updated</b></p> <ul style="list-style-type: none"> <li>• <a href="#">Java heap recommendation</a></li> <li>• <a href="#">CBR search information</a></li> <li>• <a href="#">CBR temp table batching information</a></li> <li>• <a href="#">Index information about improving the performance of multi-value property searches</a></li> <li>• <a href="#">Information about improving performance related to Auditing</a></li> <li>• <a href="#">Recommendations for handling a large number of folders and objects (Formerly maintained as a technical notice)</a></li> <li>• <a href="#">Email Manager</a></li> <li>• <a href="#">Additional information about indexing for RM disposition sweeps</a></li> <li>• <a href="#">Web client timeout information associated with viewing large federated documents</a></li> </ul> <p><b>Removed</b></p> <ul style="list-style-type: none"> <li>• Reference to the now obsolete Recommendations for Handling a Large Number of Folders and Objects Technical Notice (The information in this technical notice has been merged into this guide.)</li> </ul>

Date	Revision
02/12/08	<p><b>Added</b></p> <ul style="list-style-type: none"> <li>• <a href="#">Centera migration tuning recommendation</a></li> <li>• <a href="#">Index information for stored searches</a></li> <li>• <a href="#">Index information for the applicationname column in the generic database table for Workplace performance</a></li> <li>• <a href="#">Index information for applicationname column in generic database table and uxy searchtype in the docversion table for Workplace performance</a></li> <li>• <a href="#">Index information for RM disposition and hold sweeps</a></li> <li>• <a href="#">RM INSUBFOLDER searches information</a></li> <li>• <a href="#">Heap size information for RM Destroy Disposition Sweeps</a></li> <li>• <a href="#">Content-based retrieval tuning tips</a></li> <li>• <a href="#">Content Federation Services section</a></li> </ul> <p><b>Updated</b></p> <ul style="list-style-type: none"> <li>• <a href="#">RM Supercededby indexing recommendation</a></li> <li>• <a href="#">Recommendation for the allocated number of VWJs per CPU</a></li> <li>• Document style template</li> </ul> <p><b>Removed</b></p> <ul style="list-style-type: none"> <li>• Note about BAM.</li> </ul>
12/31/06	Added RM sweep index tuning specific content. Added Oracle 9i specific information on disabling hash join in database query plans.
09/31/06	Updated document outline and formatting. Expanded list of required third-party documentation. Updated and expanded content throughout.
04/13/06	Added PE Tuning section, CFS section, and RM Sweep index tuning.
02/16/06	Minor tweak to "Indexing for Searches."
01/25/06	Updated the RM indexing section.
10/27/05	Updated PE Environment cache section.
09/28/05	Added DB2 generated columns info to support CE case-insensitive searches.
09/12/05	Updated the RM indexing section.
08/17/05	Updated JVM tuning related to TCM.



Date	Revision
06/22/05	Edits to the DB2 section.
06/07/05	Added indexing for stored searches section.
05/31/05	Added RM section and tweaked Indexes and Win TCP Parameters.
04/22/05	Added CE DB2 Tuning Info.
04/13/05	Updated the Process Engine section with feedback from reviewers.
03/31/05	Initial draft.

## About this document

This document provides tuning tips that can help improve the performance of your IBM® FileNet® P8 system.

## Tuning tip organization

If a tuning tip involves a third party product, and it applies to more than one of the FileNet P8 applications, the tuning tip will be in the section for the third party product, under a heading for the specific vendor. However, if a tuning tip involves a third party product, and it applies to only one FileNet P8 application, the tuning tip will be in the section for the FileNet P8 application, under a heading for the third party product and a sub-heading for the specific vendor.

## Parameter settings

Many tips in this document do not include suggested parameter settings, so that you can make an informed decision about the setting values should be. However, if the IBM Performance Lab has found certain settings beneficial, the settings are suggested. You can adjust the suggested settings to fit your environment.

## Assumptions

This document assumes you are a systems administrator and you have:

- Reviewed the IBM FileNet P8 Hardware and Software Requirements Guide, and have verified that FileNet P8 supports the hardware and software components in your current configuration.
- Installed FileNet P8 according to the IBM FileNet P8 Platform Installation and Upgrade Guide.
- Configured the FileNet P8 databases with sufficient processing power, memory, and a set of indexes appropriate to the application.

## Recommended documentation

The following IBM and third party reference documents are available for when you are tuning an environment:

- IBM FileNet P8 Platform Installation and Upgrade Guide
- IBM FileNet P8 Hardware and Software Requirements Guide
- IBM FileNet P8 Troubleshooting Guide

## Application servers and Java Virtual Machine documentation

### IBM WebSphere®

- WebSphere Application Server V6 Scalability and Performance Handbook  
<http://www.redbooks.ibm.com/redbooks/SG246392/wwhelp/wwhimpl/java/html/wwhelp.htm>
- WebSphere Application Server Documentation  
<http://www-306.ibm.com/software/webservers/appserv/was/library/>

- WebSphere Application Server Tuning Parameter Hotlist  
[http://publib.boulder.ibm.com/infocenter/wasinfo/v6r0/topic/com.ibm.websphere.nd.doc/info/ae/ae/rpf\\_hotparameters.html](http://publib.boulder.ibm.com/infocenter/wasinfo/v6r0/topic/com.ibm.websphere.nd.doc/info/ae/ae/rpf_hotparameters.html)
- WebSphere Information Integrator (IICE). (Applicable to CFS-CS+)  
[http://www.ibm.com/software/data/integration/db2ii/editions\\_content.html](http://www.ibm.com/software/data/integration/db2ii/editions_content.html)  
<http://publib.boulder.ibm.com/infocenter/wsiihelp/v8r3/index.jsp>

### **BEA WebLogic**

- General server documentation  
<http://edocs.bea.com/>
- WebLogic Server 8.1.x performance and tuning  
<http://edocs.bea.com/wls/docs81/perform/index.html>

### **JRockit**

- Tuning WebLogic JRockit 8.1 Java™ Virtual Machine (JVM)  
<http://e-docs.bea.com/jrockit/geninfo/confune/index.html>

### **Database documentation**

#### **DB2®**

- DB2 Information Center  
<http://publib.boulder.ibm.com/infocenter/db2luw/v8/index.jsp>
- Best practices for tuning DB2 UDB v8.1 and its databases  
<http://www.ibm.com/developerworks/db2/library/techarticle/dm-0404mcarthur/index.html>
- DB2 Tuning Tips for OLTP Applications  
<http://www.ibm.com/developerworks/db2/library/techarticle/anshum/0107anshum.html>

#### **Oracle**

- Oracle Database Performance Tuning Guide  
[http://download-west.oracle.com/docs/cd/B12037\\_01/server.101/b10752/toc.htm](http://download-west.oracle.com/docs/cd/B12037_01/server.101/b10752/toc.htm)

#### **SQL Server**

- Microsoft® SQL Server 2000 RDBMS Performance Tuning Guide for Data Warehousing  
<http://www.microsoft.com/technet/prodtechnol/sql/2000/maintain/rdbmspft.mspix>
- Improving SQL Server Performance  
<http://msdn2.microsoft.com/en-us/library/ms998577.aspx>
- Optimization Tips  
[http://www.databasejournal.com/features/mssql/article.php/1439801#part\\_2](http://www.databasejournal.com/features/mssql/article.php/1439801#part_2)

## Access IBM FileNet documentation, compatibility matrices, and fix packs

To access documentation, compatibility matrices, and fix packs for IBM FileNet products:

1. Navigate to the Product Documentation for FileNet P8 Platform support page.  
<http://www-1.ibm.com/support/docview.wss?rs=3247&uid=swg27010422>.
2. Select a PDF or a Doc Link, whichever is appropriate.

## Contact customer support

For information about contacting customer support:

1. Navigate to the FileNet Product Family support page:  
<http://www-01.ibm.com/software/data/content-management/filenet-product-family/support.html>.
2. Click IBM FileNet Support Communications, or search for a particular support topic under “Enter search terms”.

## Feedback

Your feedback helps us to provide quality information. Send your comments about this publication or any other IBM FileNet documentation by e-mail to [comments@us.ibm.com](mailto:comments@us.ibm.com). Be sure to include the name of the product, the version number of the product, and the name and part number of the book (if applicable). If you are commenting on specific text, include the location of the text (for example, a chapter and section title, a table number, a page number, or a help topic title).

## Databases

The following topics apply to all databases unless otherwise specified.

### Indexes

The index tables in this document contain information about recommended non-clustered database indexes that you can create by using single or multiple database columns (which will be referred to from this point on as “single property indexes” and “composite properties indexes”, respectively).

These tables include:

- Database table name being indexed
- Database column names that make up the index
- FileNet P8 name (property name) for the database column
- The tool you need to use to create the index
- Description of the index
- Reason to create the index

#### NOTES

- The “uxy” prefix represents a variable in the property name, and indicates the property is a custom property that is from an installed expansion product or is user-defined.
- Use FileNet Enterprise Manager (EM) to create single property indexes, unless the index is for a FileNet P8 system property. If an index is for a FileNet P8 system property or composite properties index, use the respective RDBMS management software.
- In addition to these recommended indexes, you should also use Microsoft’s SQL Profiler, Oracle’s Statspack, Oracle ADDM, or DB2 Design Advisor to identify properties (columns) that should be indexed.

#### Indexing to improve stored searches

Database Table Name	Database Column Name	FileNet P8 Property Name	Tool	Description/ Reason to create index
Relationship	tail_id ----- head_id		RDBMS	Add a composite properties index to improve performance when retrieving stored searches or properties of stored searches.

## Memory and disk I/O

Memory and disk I/O limitations are the main contributors to database-related performance issues.

Use the appropriate operating system tools to monitor memory usage, CPU usage, and disk performance. Use vmstat and iostat on UNIX®-based systems.

## Compute statistics for optimal performance

As content is added, removed, or changed in your database system, it is important to periodically run a statistics update for optimal performance.

### Oracle

#### Disable hash join to speed up slow searches

*(Oracle 9i only)* When a search that is intended to return less than a few thousand objects is slow, and a gathered database query plan shows that a hash join is being used, you can potentially speed up the search by disabling the hash join.

#### NOTES

- The `hash_join` parameter is obsolete in Oracle 10g. The following procedure applies only to Oracle 9i.
- Disabling hash joins could adversely affect other searches or API calls that return very large rowsets (such as `GetContainees()` on a folder with a very large number of objects) and that rely on a hash join in the database query plan.
- The hash join type is best used in retrieving large rowsets, but can also be the result of certain complex queries. However, the hash join type is typically poor for transaction processing types of applications with fewer returned rows, or during searching, where the smaller response time of the initial batch returned is of highest importance.

#### Disable hash join

1. Set the `hash_join_enabled` parameter to “false” in the `initORACLE_SID.ora` file.
  - a. Edit the `initORACLE_SID.ora` file.

##### Windows®

```
ORACLE_HOMEDATABASE\INITORACLE_SID.ORA
```

##### UNIX

```
$ORACLE_HOME/dbs/initORACLE_SID.ora
```

- b. Add the line “`hash_join_enabled=false`” to the file, and save the file.

**NOTE** In some instances, an spfile might be used for Oracle startup. If this is the case for your setup, use the following procedure to update the `initORACLE_SID.ora` file using SQLPlus.

##### Windows

- c. Create the spfile by issuing the following command:

```
CREATE PFILE='ORACLE_HOMEDATABASE\INITORACLE_SID.ORA' FROM  
SPFILE='ORACLE_HOMEDATABASE\SPFILE ORACLE_SID.ORA';
```

- d. Add the line “`hash_join_enabled=false`” to the generated `init.ora` file.
- e. Start the database by issuing the following command:

```
STARTUP PFILE = ORACLE_HOMEDATABASE\INITORACLE_SID.ORA
```

## UNIX

- f. Create the pfile by issuing the following command:

```
CREATE PFILE='$ORACLE_HOME/dbs/initORACLE_SID.ora' FROM  
SPFILE='$ORACLE_HOME/dbs/spfileORACLE_SID.ora';
```

- g. Add the line “hash\_join\_enabled=false” to the generated init.ora file.  
h. Startup the database by issuing the following command:

```
STARTUP PFILE = $ORACLE_HOME/dbs/initORACLE_SID.ora
```

2. Restart CE.

- a. Verify that no users are logged in to the system.
- b. Stop the CE object store services on each machine that uses the database.
- c. Restart the database.
- d. Start the CE object store services.

## **QUEUESIZE parameter: Maximum number of concurrent connections**

The QUEUESIZE parameter value in the listener.ora file specifies the maximum number of concurrent connections that the listener can receive.

Increasing this parameter value helps prevent listener backlog full conditions, which generate the Oracle ORA – 12541: “No listener” and ORA-12535: “TNS:operation timed out.” errors.

The listener.ora file resides on the Oracle database server:

```
$ORACLE_HOME/network/admin
```

To set the QUEUESIZE parameter, edit the address line in the listener.ora file as follows:

```
(ADDRESS = (PROTOCOL = TCP)(HOST = servername)(PORT = portnumber)(QUEUESIZE=100))
```

For more information, see “Handling Large Volumes of Concurrent Connection Requests” in the Oracle online help.

[http://download-west.oracle.com/docs/cd/B10501\\_01/network.920/a96580/listener.htm#491633](http://download-west.oracle.com/docs/cd/B10501_01/network.920/a96580/listener.htm#491633)

## Microsoft SQL Server

### Reclustering the docversion table

Reclustering the docversion table on the object\_id column improves performance for searches that involve joins or that use an ORDER BY object\_id clause and result in a table scan and a large sort. When the table is reclustered, the rows are already sorted by object\_id, removing the SORT step in the query plan in certain cases.

To improve common searches, the IBM Performance Lab recommends changing the existing clustered index I\_DOCVERSION22 to a clustered index on the object\_id column.

The following searches can result in improved performance:

- CBR searches
- RM continuable searches
- Workplace searches with an ORDER BY object\_id clause
- Other searches with an ORDER BY object\_id clause

Reclustering is only necessary on SQL Server and for object stores originally created on the specific versions of CE as stated below.

#### Versions affected

The docversion table is clustered on object\_id when new object stores are created on CE 3.5.2-014 and later.

Therefore, reclustering is only necessary for CE 3.x object stores created before 3.5.2-014.

For details, see "Reclustering the FileNet Content Engine DocVersion Table on the SQL Server Database" technical notice at:

[http://www-01.ibm.com/support/docview.wss?rs=3278&context=SSNVNV&dc=DB560&dc=DB520&uid=swg21328193&loc=en\\_US&cs=UTF-8&lang=en&rss=ct3278db2](http://www-01.ibm.com/support/docview.wss?rs=3278&context=SSNVNV&dc=DB560&dc=DB520&uid=swg21328193&loc=en_US&cs=UTF-8&lang=en&rss=ct3278db2)



## Application servers

### JVM tuning

#### Heap size

The Java heap is a repository for live objects, dead objects, and free memory. The heap size is a configurable parameter that the user can control at startup. When the JVM runs out-of-memory in the heap all execution in the JVM stops while a Garbage Collection (GC) algorithm frees space that is no longer required by the application. Determining the heap size is important as it specifies how often, and for how long, the system executes garbage collections. High GC times negatively affect the performance of the system.

Two goals exist in tuning the heap size:

- Minimize the amount of time that is spent doing garbage collection
- Maximize the systems ability to handle a significant workload

Depending on your FileNet P8 workload patterns, you need to balance the size of your heap and the number of instances in your application server cluster to minimize the number of lengthy garbage collection pauses and frequency of full garbage collections caused by insufficient heap space.

The lowest acceptable value for the minimum heap size is 512MB, and the lowest acceptable value for the maximum heap size is 1GB.

#### Allocating heap space to system classes

In addition to objects allocated by the application, system classes are also stored in the heap. Like all other objects in the heap, if a system class is unused for some time, it will be de-allocated from the heap. However, it is a certainty that classes de-allocated will ultimately be required again. As a result, it is recommended that users allocate one-tenth of their heap to permanent objects, such as system classes.

*(Sun Hotspot only)* Allocating space to system classes will improve the stability of the system by reducing the likelihood of fragmenting the heap. Below is an example of the argument to pass to the JVM at startup if you were to allocate 128MB to the system classes:

```
-XX:MaxPermSize=128M
```

#### Thread pools

IBM allows clients to configure the number of threads available to the application server. A thread pool enables the application to reuse threads, eliminating the need to create new threads at run time. Creating new threads can be slow and time consuming. As a result, especially against a loaded system, increasing the number of threads in the pool might improve the performance of the system. By default a minimum of 10 threads are in the pool, with a maximum of 50. For applications handling a significant workload, IBM suggests increasing the number of threads in the pool. Furthermore, it is also recommended that clients apply the same value to the minimum and maximum size to eliminate the overhead of creating threads at run time.

# Content Engine

## Databases

### Indexes – Overview

The following topics apply to all CE databases unless otherwise specified.

For overview information, see “[Databases – Indexes](#)” on page 13.

#### Object Value Properties (OVP)

Database Table Name	Database Column Name	FileNet P8 Property Name	Tool	Description/ Reason to create index
<i>table name</i>	<i>uxy_property name</i>	<i>Property Name</i>	EM	<p>In general, if using OVPs where there is an association to a large table such as docversion, an index should be considered.</p> <p>Generally, if the association properties values are mostly null, there is little overhead in index maintenance, and little affect on document ingestion.</p>

#### Prevent MSDTC timeout problems when setting indexes

If you encounter a “Database SQL execution error has occurred” message when creating indexes with EM, it indicates that the time taken by the database to create the index has exceeded the default MSDTC timeout time of 60 seconds.

Try adjusting the MSDTC timeout value, or use the RDBMS management software instead.

#### Adjust the MSDTC timeout

Use the following procedure to enable index setting to proceed temporarily.

1. On the CE server, navigate to Start > Programs > Administrative Tools > Component Services.
2. Select Computers.
3. Right-click My Computer and select Properties.
4. Select Option.
5. Set Transaction Timeout to 0 (infinite) seconds.
6. Click OK.
7. Restart the CE server.

## Indexes – Indexing multi-value property tables

For overview information, see “[Databases – Indexes](#)” on page 13.

Consider the following indexes to enhance system performance [when users execute multi-value property searches](#).

Database Table Name	Database Column Name	FileNet P8 Property Name	Tool	Description/ Reason to create index
ListOfString	element_value -----	<i>Property Name</i>	RDBMS	Add a composite properties index to improve performance of multi-value property searches.
ListOfBoolean	parent_prop_id -----			
ListOfInteger32	parent_id			
ListOfFloat64				
ListOfDatetime				
ListOfId				
ListOfString	parent_id ----- parent_prop_id ----- element_value	<i>Property Name</i>	RDBMS	Add a composite properties index to improve performance of multi-value property searches that use a wildcard value, such as LIKE %value%.

Create these types of indexes only if the table has at least several hundred rows and has enough distinct values relative to the number of rows in the table. (You can determine the cardinality by reviewing the table statistics or by using a group by query, for example select count(\*), element\_value from ListOfString group by element\_value order by 1).

### Example: ListOfString index

(792350)For example, use the following CREATE INDEX statement to create an index that improves system performance during user searches for string multi-value properties:

```
create index I_MV on ListOfString (element_value, parent_prop_id, parent_id)
```

If the object store is case-insensitive, create the index using the LOWER function on (element\_value). For more information, see "[Force Case Insensitive Search check box: Case-sensitive queries](#)" on page 21.

### Example: ListOfString index for searches that begin with %wildcard% values

(792350)Use the following CREATE INDEX statement to create an index that improves system performance during user searches for string multi-value properties when the index on (element\_value, parent\_prop\_id, parent\_id) cannot be used because the search begins with a wild card value, such as LIKE %value%:

```
create index I_MV2 on ListOfString (parent_id, parent_ prop_id, element_value)
```

If the object store is case-insensitive, create the index using the LOWER function on (element\_value). For more information, see "[Force Case Insensitive Search check box: Case-sensitive queries](#)" on page 21.

**NOTES**

- Each database type has its own limit on index byte size, which can require extra work to create the index for string datatypes. By default, the ListOfString element\_value column is of character size 4000 (1333 on Oracle), which might be longer than the maximum index length for the database.
- 4000 characters consume 8000 bytes during value insertion on MSSQL and 8000 bytes on column creation on DB2. On Oracle, the varchar2 allocation is tripled so it is 3999 bytes consumed for a 1333 character size column.

Use the guidelines in the following table. If the database does not allow the index to be created, reduce the size of the ListOfTable element\_value column. However, reducing this value might not work if the values stored in the column are too large. If only a few values are too large, consider changing the data in those columns to allow the index to be created.

Database	Index MaxBytes	Byte size of element_value	Index allowed?	Notes
Microsoft SQL 2000	900	8000	No	
Microsoft SQL 2005	900	8000	No	
Oracle 9i	75% block size	3999	Under certain conditions	An index cannot be created at the block size of 4k, but can be created at an 8k block size. The database block size is set during database creation. Oracle 10g has the same 75% block size limitation. For more information, see the Oracle 10g notes.
Oracle 10g	75% block size	3999	Under certain conditions	The Oracle 10g block size is also minus overhead, which is small.
DB2 8.2	1024	8000	No	The DB2 1024 block size is also minus overhead, which is small.
DB2 9	8k	8000	Yes	8k is one-quarter the table space size. The DB2 8k block size is also minus overhead, which is small.

**NOTE** For strings on Oracle and DB2, if Force Case Insensitivity is enabled, the index must be created on lower(element\_value) along with the other composite columns.

For more information, see “[Force Case Insensitive Search check box: Case-sensitive queries](#)” on page 21.

## Oracle and DB2

### Force Case Insensitive Search check box: Case-sensitive queries

When selected, the Force Case Insensitive Search checkbox causes all ad hoc queries to operate in a case insensitive manner. That is, selecting this checkbox causes CE to append the LOWER function to query statements, which lowercases all string column values used in query statements.

#### Make ad hoc queries case insensitive

1. Log on to EM as an object store administrator.
2. Select and right-click the object store and select Properties.
3. Go to the Query tab.
4. Select the Force Case Insensitive Search checkbox.
5. Click OK or Apply when you are done.

#### Maximize performance when using case insensitive queries

The performance of case insensitive queries is adversely affected when the string properties that are used in a query are indexed in the database. Therefore, the IBM Performance Lab recommends that you replace these indexes with indexes that use the LOWER function to maximize query performance. (Function-based indexes in Oracle and generated columns in DB2.)

For more information about function-based indexes or generated columns, see the appropriate RDBMS documentation.

## DB2

### New (empty) object stores

After you create an object store, it is important that you monitor the database. As new content is added to the object store, the DB2 optimizer continually requires current table and index statistics to select the most efficient access plan. Ensuring that the DB2 optimizer has current statistics is especially important during large ingestions at a high rate, such as migration. It is also important for proof-of-concept testing, stress testing, and testing custom applications.

This section provides information about:

- Monitoring your databases using the FileNet System Manager Dashboard
- Enabling automatic statistics
- Using the DB2 Task Center to gather statistics manually
- Setting tables as volatile

## System Manager Dashboard – View counts and durations

For each object store, CE provides counts and durations for the various database operations and server operations. This information can provide valuable insight into database performance.

You can find these counts and durations using the FileNet P8 System Manager Dashboard tool.

- To view counts and durations for database operations, select the following nodes for a given machine and CE server:  
*User > Object Store Name > Database*
- To view counts and durations for server operations select the RPC node for a given machine and CE server.

## Automatic statistics collection

- **DB2 version 9** – Automatic statistics collection is enabled by default.
- **DB2 version 8** – You can configure automatic statistics collection from the Health Center Monitoring tool by clicking *Selected > Configure > Automatic Maintenance* and following the steps provided by the wizard.

Click the following link for instructions on enabling automatic statistics from the command line:

<http://publib.boulder.ibm.com/infocenter/db2luw/v9/index.jsp?topic=/com.ibm.db2.udb.admin.doc/doc/t0011763.htm>

**NOTE** Enabling automatic statistics collection has shown to provide good results. However, in scenarios where you are just beginning to populate an object store, within the first 1,000,000 documents, better results have been observed by manually gathering statistics using the DB2 Task Center.

## Manual statistics collection using the DB2 Task Center

Use the DB2 Task Center to update statistics at regular intervals and send notifications. Scheduling information is stored in the DB2 tools Catalog database.

For example, something similar to the following would be listed on the Command Script tab of the Task properties. The following could then be scheduled to run every 20 minutes during a migration.

```
CONNECT TO OS1DB;

RUNSTATS ON TABLE CEUSER.CONTAINER ON ALL COLUMNS AND INDEXES ALL
ALLOW WRITE ACCESS ;

COMMIT WORK;

RUNSTATS ON TABLE CEUSER.CONTENT ON ALL COLUMNS AND INDEXES ALL ALLOW
WRITE ACCESS ;

COMMIT WORK;

RUNSTATS ON TABLE CEUSER.DOCVERSION ON ALL COLUMNS AND INDEXES ALL
ALLOW WRITE ACCESS ;

COMMIT WORK;

RUNSTATS ON TABLE CEUSER.GENERIC ON ALL COLUMNS AND INDEXES ALL ALLOW
WRITE ACCESS ;

COMMIT WORK;

RUNSTATS ON TABLE CEUSER.HIERARCHYINDEX ON ALL COLUMNS AND INDEXES ALL
ALLOW WRITE ACCESS ;
```

```
COMMIT WORK;  
  
RUNSTATS ON TABLE CEUSER.RELATIONSHIP ON ALL COLUMNS AND INDEXES ALL  
ALLOW WRITE ACCESS ;  
  
COMMIT WORK;  
  
RUNSTATS ON TABLE CEUSER.SECURITYDESC ON ALL COLUMNS AND INDEXES ALL  
ALLOW WRITE ACCESS ;  
  
COMMIT WORK;  
  
CONNECT RESET;
```

For more information about using the Task Center to schedule tasks, see the DB2 documentation:

<http://publib.boulder.ibm.com/infocenter/db2luw/v9/index.jsp?topic=/com.ibm.db2.udb.admin.doc/doc/t0022277.htm>

### **Set tables to volatile**

Ensuring that statistics are current is the preferred method after the initial migration and for production use. However, setting a table to volatile ensures that the DB2 optimizer will always use the indexes defined on the table. The advantage to tuning the DB2 optimizer in this way is the potential that the DB2 optimizer will make the correct optimization choice when the statistics are out of date, for example, when a table suddenly grows in size before the system statistics can be updated. However, this only works in the cases where the index contains all the columns referenced or when the index is able to apply a predicate in the index scan.

CE has defined some tables to be used as queues. Because the number rows in these tables are in constant change, setting the table to VOLATILE is preferred over relying on current statistics.

The following tables are used as queues:

QUEUEITEM

The following tables are not set as volatile, but they should be monitored for cases where a temporary setting would be beneficial:

CONTAINER

CONTENT

DOCVERSION

GENERIC

HIERARCHYINDEX

RELATIONSHIP

SECURITYDESC

### ***Define a table as volatile***

To define a table as volatile, run the following command from the DB2 command line processor:

```
db2 connect to database_name

db2 alter table schema.table_name volatile

db2 connect reset
```

### ***Remove the volatile setting from a table***

To remove the volatile setting from a table, run the following command from the DB2 command line processor:

```
db2 connect to database_name

db2 alter table schema.table_name not volatile

db2 connect reset
```

### ***List tables that are currently set to volatile***

To list tables that are currently set to volatile, run the following command from the DB2 command line processor:

```
db2 "SELECT NAME from sysibm.systables where volatile='C' order by NAME"
```

## **DB2 tuning recommendations and maintenance tasks**

The following table lists DB2 configuration and tuning parameters, and includes suggested starting values.

Use these settings, and perform these maintenance tasks to increase CE performance when CE is running on DB2.

**CAUTION** The IBM Performance Lab highly suggests that you reference the following IBM documents to perform maintenance tasks and to determine if the suggested parameter starting values are appropriate for your environment.

- Best practices for tuning DB2 UDB v8.1 and its databases
- DB2 Tuning Tips for OLTP Applications



**Configuration and tuning parameters**

Parameter	Suggested starting value	Notes
APPLHEAPSZ	2560  <b>NOTE</b> DB2 version 9.x: Set this value to automatic.	Setting this parameter value can help avoid CE installation errors.
DB2_PARALLEL_IO	Ensure that the value matches your table space and disk configuration.	The values of this registry parameter specifies how DB2 calculates the I/O parallelism of a table space and can which can have a significant impact disk I/O performance.
DB2_PINNED_BP	YES	The value in this registry parameter specifies whether the database global memory that is used (including buffer pool memory) is kept in system main memory.  Setting this parameter might help improve system performance.
DB2_USE_ALTERNATE_PAGE_CLEANING	ON	The value of this registry parameter determines if page cleaners are always on, or if they are initiated according to other registry parameter values.  Page cleaners improve the overall performance of a system because they perform I/O operations that would otherwise have to be handled by database agents. Additionally, when page cleaners are always on, they can improve system performance because:  They do not search the buffer pool for victim pages; they are informed of victim page locations.  Victim pages cannot accumulate, which can consume I/O when a large amount of victim pages are written to disk at one time.  Setting this parameter might help improve system performance.
LOCKLIST	MAXAPPLS multiplied by 14.  <b>NOTE</b> DB2 version 9.x: Set this value to automatic.	Specifies the maximum amount of memory allocated to the lock list, which holds the locks of applications that are concurrently connected to the database.
LOCKTIMEOUT	30	Specifies the number of seconds that an application waits to obtain a lock before it times out. Setting this value helps avoid global deadlocks and lock escalations.

Parameter	Suggested starting value	Notes
LOGFILSIZ	6000	Specifies the size of both primary and secondary log files.  <b>NOTE</b> The size of the log files should be in balance with the number of primary log files.
MAXAPPLS	25% to 50% of number of users who are actively using the system.  <b>NOTE</b> DB2 version 9.x: Set this value to automatic.	Specifies the maximum number of concurrent applications that can be connected (both local and remote) to the database.
PCTFREE	20% - 35%	Set the value in the following tables: CONTENT CONTAINER DOCVERSION PROPERTYDEFINITION RELATIONSHIP GENERIC LINK SECURITYDESC
SORTHEAP	2048  <b>NOTE</b> DB2 version 9.x: Set this value to automatic.	Specifies maximum number of private memory pages to be used for private sorts, or the maximum number of shared memory pages to be used for shared sorts.  If the number of overflows is greater than 10, gradually increase the value for SORTHEAP until the overflows are eliminated.
STMTHEAP	8192	Specifies the maximum amount of memory that can be allocated as a work space for the SQL compiler during compilation of an SQL statement.

## Full-text Indexes for content-based retrieval (CBR)

The following topics apply to all CE databases unless otherwise specified. They provide information and default values for a few of the Verity tunable parameters that are beneficial to CE full-text Indexes performance. All of these parameters are updated via a registry entry, so the corresponding key is also provided.

**NOTE** Out-of-the-box installations do not include most of these registry entries. Therefore, changes to the following parameters might require that a FileNet P8 administrator first create the registry keys.

### Indexing threads per collection parameter

The MtlIndexer Max Threads parameter controls the number of threads that concurrently index a collection. When it has a value that is greater than 1, multiple threads are launched to index a single batch (for a single collection). You can set the value of this parameter to equal the number of logical CPUs on the machine that hosts your filestore service for increased indexing throughput, however significant improvement are observed by leaving this parameter at its default of 1 and tuning the “Concurrent Collections Indexed Threading Parameter”.

Registry Key: HKLM - SOFTWARE\FileNET\ECM\Content RM\

DWORD: MtlIndexer Max Threads

Default Value: 1

### Concurrent collections indexed threading parameter

P8CE-3.5.2-016 introduces the Indexer Threads parameter that controls the number of collections that can be indexed simultaneously. The suggested value for this parameter is the value that is equal to the number of logical CPUs that are on the machine that hosts the filestore service.

Registry Key: HKLM\SOFTWARE\FileNET\ECM\Content RM\

DWORD: Indexer Threads

Default Value: 3

### Indexing Optimization Interval parameter

P8CE-3.5.2-016 introduces the Indexing Optimization Interval parameter that controls the minimum amount of time (in minutes) between optimization operations for a collection. Increasing this parameter improves indexing throughput, but the drawback is slower searches for content that has only just been indexed since the last time optimization ran. Therefore, if you are indexing a large amount of content that will not be searched on, increasing the value of this parameter will improve indexing throughput.

Registry Key: HKLM\SOFTWARE\FileNET\ECM\Content RM\

DWORD: Indexer Optimization Interval

Default Value: 60

### Auto-administration parameter

P8CE-3.5.2-016 introduces the Verity 6.x Index With Optimize parameter that when enabled causes Verity 6.x to recover gracefully from terminated index jobs, crashes caused by power outages, and other unexpected system faults. Disabling this parameter might improve indexing throughput, but the drawback is slower search response times. However, due to the data integrity features associated with this parameter, the IBM Performance Lab suggests that this parameter value remain enabled.

Registry Key: HKLM\SOFTWARE\FileNET\ECM\Content Engine\Content Search\

DWORD: Verity Index With Optimize

Default Value: 1 (enabled)

### Secondary parameters

The following tables provide an additional list of configurable parameters that are available for Verity 6.x. However, the IBM Performance Lab has determined that these parameters are not likely to affect performance, and should be modified with caution.

#### NOTES

- Some of these parameters are explained in further detail in the sections following the tables.
- Some of these parameter names include the term “Verity”, which might have been changed by Autonomy after acquiring Verity.

### Content Search registry key

This table lists the parameters that are available under the Content Search registry key:

“ HKLM – SOFTWARE\FileNET\ECM\Content Engine\Content Search\ “

Parameter name	Default value	Description	Service
Verity Max Docs	1000	The maximum number of documents returned from a single query (VdkSearchNewArgRec maxDocs). The result set from a full-text query will never contain more than the specified value.	Object store service
Verity Max Memory	512000	The maximum memory (specified in bytes) that a Verity 6.x session could use (VdkSessionNewArgRecmaxMemory).	Object store and file store services
Verity Batch	100	The number of rows that are read in each batch during a query ((VdkDocReadBatchArgRecreadRange).	Store service
Verity Extra Locales	empty	A list of additional locale names supported by the Verity Developer Kit (VDK). (This has no effect on performance.)	Object store and file store services.
Verity Index with Optimize	0	If this parameter is set to a non-zero value, collections are opened for indexing at a service level VdkServiceType_Optimize.	File store service.

**Content Resource Manager registry key**

This table lists the parameters that are available under the Content Resource Manager registry key:

“ HKLM – SOFTWARE\FileNET\ECM\Content RM “

Parameter name	Default value	Description
Max Queued Index Bytes	512000	Specifies the maximum number of bytes to be queued before writing IQ files.
Max Queued Index Requests	500	Specifies the maximum number of requests to be queued before writing IQ files.
Max Waiting Index Workers	200	Specifies the maximum number of waiting worker objects that are blocked before writing IQ files.
IQ File Flush Seconds	60	Specifies the maximum number of seconds to wait before writing IQ files.
Retry Re-indexing Seconds	300	Specifies the minimum number of seconds to wait before attempting to re-index a batch that failed on a previous attempt.
Indexer Logging Level	0	Specifies the verbose indexing logging level. If this parameter is set to a non-zero value, each indexer thread will log activity to a file created in the IndexQueue directory:  Logging levels: 0 – None 1 – Errors 2 – Activity 3 – Verbose
Indexer Threads	3	Specifies the number of file store service threads controlling indexing. A file store is never assigned to more than one simultaneous controlling threads.
MtIndexer Min Files Per Thread	100	If multi-threaded indexing is enabled, this value is the minimum number of files per batch to cause a separate thread to be launched.
Re-index Batch Size	500 <sup>1</sup>	Specifies the batch size that is used during file store re-indexing. That is, the number of files to be processed before sending requests to the Index Queue.
Max IQ Files Per Merge	32	Specifies the maximum number of IQ files to be merged in a single file merge pass.  The maximum allowable value is 128, a value of 1 or less prevents IQ file merging.
IQ File Merge Max Requests	8000	If this parameter is set to a non-zero value, merging of IQ files is logged to a file named IQFileMerge.txt in the Index Queue directory.

<sup>1</sup> The default value of this parameter is the same value that is assigned to Max Queued Index Request.

### **Max Queued Index Bytes, Requests, Workers, and IQ File Flush Seconds parameters**

These four parameters control how often IQ (Index Queue) files are written from memory to disk by the file store service. The first limit reached triggers writing of IQ files. Increasing these values causes more requests to be queued in memory (backed by the transaction log files), with fewer but larger files written. Decreasing these values causes fewer requests to be queued in memory, with more, smaller IQ files being written.

Increasing the values is not recommended and is normally not necessary (IQ file merging will handle creating large index batches). In some cases, decreasing the values might be useful if there is a low volume of content ingestion and a need to have more immediate access to the content via full-text search.

**NOTE** The Re-index Batch Size parameter controls how many files are processed in a batch during re-indexing, but does not directly control the IQ file size (still controlled by the four limits listed above, and IQ file merging). Changing this value is not recommended.

### **Indexer Threads, and MtIndexer Min Files per Thread parameters**

These three parameters control how many file store service threads perform indexing. The Indexer Threads parameter specifies the maximum number of controlling threads. A file store is never assigned to more than one simultaneous controlling thread. The MtIndexer Max Threads parameter controls when multiple threads are launched to index a batch of files for a given Verity 6.x collection. By default, the value of the MtIndexer Max Threads parameter is 1, and the controlling thread does not launch any threads. Instead, the controlling thread indexes the batch itself. If the value of the MtIndexer Max Threads parameter is greater than 1, the controlling thread launches threads according to the parameter setting value to index the collection, and the controlling thread does not perform indexing.

The value of the Indexer Threads parameter can be increased if CPU resources are available and more than three file stores are managed by the file store service. The value can be decreased if CPU resources are limited and multiple file stores are indexed by the file store service.

The MtIndexer Min Files per Thread parameter prevents a very small number of files from being split into separate batches and indexed by separate threads. Normally, the default value of 100 is sufficient, and this parameter should remain at the default value.

### **Max IQ files Per Merge and IQ File Merge Max Requests parameters**

These two parameters control Index Queue file merging. File merging is an essential performance feature that enables the file store service to frequently clear the in-memory (log file backed) queue of index requests, while supporting large index batches (by merging small IQ files into combined files).

The Max IQ Files per Merge parameter controls only the maximum number of files the merger examines on a single merge pass. (This parameter should not have any reason to be modified.)

The IQ File Merge Max Requests parameter defines the maximum number of index requests that are contained in a single IQ file. The value of this parameter should not exceed 8000 because Verity 6.x will fail if more than 8000 files are passed to the VDK in an index batch.

**NOTE** If multi-threaded indexing is enabled, then in theory, a value larger than 8000 could be used, because the batch is divided up across the threads.

## Centera FCD content migration throughput

If a file store service machine has spare CPU cycles, and you want to increase the Centera FCD migration rates, you can increase the number of migration threads by adjusting the following registry value:

```
[HKEY_LOCAL_MACHINE\SOFTWARE\FileNet\ECM\Content Rm]  
"Fixed Content Migration Threads"=dword: 00000005
```

If this registry value is missing, then the system is using the default decimal value of 5. You can increase this number until you achieve a balance between CPU utilization and migration rates.

## Queries

### Enumeration Batch Size parameter

Use the Enumeration (Enum) Batch Size parameter to specify the number of enumerated items that CE fetches in one request

By default the displayed value is 0, which implies a server default of 20.

Keeping this parameter value at 20 or setting it to a value that is less than 20, maintains reasonable response times when users browse through an object store.

**Important** Setting this parameter to a value that is greater than 20 can initiate a database table scan when CE fetches the items for the request, which can have an adverse effect on performance.

### Modifying the Enumeration Batch Size parameter

1. In EM, expand the object stores node.
2. Right-click an object store, and then select Properties.
3. Select the Properties tab.
4. Locate this parameter in the displayed list of parameters, and then make the necessary modification.

### Default Query Batch Size parameters

Use the Default Query Batch Size parameter to specify the number of items CE fetches per roundtrip to the server when [users execute searches](#).

By default the displayed value is 0, which implies a server default of 20.

Increasing this parameter value slightly, can increase system response time when [users execute searches](#) in an object store. The IBM Performance Lab experienced improved performance when the value of the Default Query Batch Size parameter was increased to 50. However, when this parameter value was set to values higher than 50, system performance showed signs of degradation.

### Modifying the Query Batch Size parameter

1. In EM, expand the object stores node.
2. Right-click an object store, and then select Properties.
3. Select the Properties tab.
4. Locate this parameter in the displayed list of parameters, and then make the necessary modification.

## CBR searches

To increase search performance, configure searches so that they do not return data in the ContentSummary column.

## CBR temp table batching

P8CE-3.5.2-014 adds in the ability to batch inserts to a temp table that is used as part of CBR searching. Performance can be improved further by adjusting a registry entry for the batch size as follows below.

A CBR search has three internal steps that CE performs in order to process the request:

1. Search Verity 6.x collections for the search words.
2. Insert the Verity 6.x hits into a database temp table.
3. Issue a join query of the database tables such as docversion to the cbr temp table.

The inserts that are done in step 2 are batched as of P8CE-3.5.2-014 with the following default batch sizes.

## Default behavior

Following are the batch sizes that are used as of P8CE-3.5.2-014:

Database	Batch size with summary	Batch size without summary
SQL Server	5	10
Oracle	20	40
DB2	10	20

- **Batch size with summary** – When the CBR search SELECT includes the ContentSummary column
- **Batch size without summary** – When the CBR search SELECT does not include the ContentSummary column

The default batch sizes were defined as optimal in one test environment, but the optimal values might vary on other machines and depending on the database, network, and the CBR search statement that is used.

## Determine the optimal batch size

1. Configure a batch size registry entry as described in the following section.
2. Run CBR searches that retrieve a large number of records.
3. Adjust the registry batch size value until you find the fastest CBR search time.



## Configure the Batch size via a registry entry

On the object store servers, create and set a DWORD registry parameter:

```
HKLM\SOFTWARE\FileNET\ECM\Content Engine\Content Search\InsertBatchSize
```

If you set the value to 0, or a value is not present, InsertBatchSize uses the default listed in the previous table.

If you set the value larger than the maximum allowed by the OLEDB providers (these values vary by provider, statement size, and binding value total size), the batch size will be capped at the maximum the OLEDB provider allows.

**NOTE** The DWORD value defaults to HEX, so change it to decimal when setting a number.

## Auditing

### Delete old records from the Event table

Auditing events are stored in the Event table. Over time, too many records can accumulate in the Event table, making it very large. When the Event table has too many records, you can delete the rows from the table to improve performance. However, before you delete rows from the Event table, the IBM Performance Lab recommends that you first back up the table during a time when users are not accessing the system.

On Oracle systems, you can issue the truncate command, which also frees allocated extents. However, before you issue this command, verify that you have saved the events that customers need or might need.

**NOTE** After deletion or truncation, FileNet P8 applications no longer have access to events information through the Audited Event enumerated properties. You can save events for access through non-FileNet P8 applications by inserting the event rows into another table in another database or by exporting them to an exported file prior to performing the deletion or truncation.

### Improve audit logging and search performance

To improve audit logging and search performance when the Auditing feature is enabled, verify that the I\_Event48 index was created on the Event table using the columns in the following order:

```
source_object_id, create_date DESC
```

If it was not created with this order, perform the following steps:

#### Microsoft SQL (Connected to the database)

```
DROP INDEX Event.I_Event48
```

#### DB2 and Oracle (Logged in as the schema owner)

```
DROP INDEX I_Event48
```

#### All databases

```
CREATE INDEX I_EVENT48 ON Event (source_object_id, create_date DESC)
```

## Object stores

### InitialCacheSize parameter

CE-3.5.2 introduces the InitialCacheSize parameter that enables CE to manage more than 200 object stores in a single domain. Therefore, if you plan on configuring more than 200 object stores in a single domain, consider adding this InitialCacheSize parameter.

Content uses this parameter when it is present. When it is not present, CE uses the default value.

Registry Key: HKLM\Software\Filenet\ECM\Content Engine\PerformanceProvider\

DWORD: InitialCacheSize

Default Value: 128K = 131584 bytes = approximately 200 object stores

Other Values: 32K = 32768 bytes = approximately 50 object stores

160k = 163840 bytes = approximately 250 object stores

## Folders, documents, and objects

### Recommendations for handling large numbers of folders, documents, and objects

This section has recommendations about handling large numbers of contained folders, documents, and custom objects in a FileNet P8 environment. Consider these recommendations when planning a FileNet P8 environment, developing efficient custom applications, and configuring Workplace.

**NOTE** These recommendations are not meant to replace any application-specific performance studies or recommendations provided by IBM for FileNet P8.

#### Folders and subfolders

Retrieving the contents of a folder can be a memory-intensive operation. The getContainees JAVA API method that retrieves folder contents returns a list of all folders and objects contained in the folder and stores that list in memory. If many containees exist, then the server can run out of memory and the application can fail or take an excessive amount of time to display the results. To avoid this problem, a custom application can use the search API methods to filter the list of sub-folders by using a query. For example, define a search query to return only those folder names that start with “Case 2”.

Workplace does not provide an option to limit the number of folders returned. Each time a user opens an object store or folder, Workplace retrieves a list of all the folders contained at that level along with several properties for each folder. For example, any time a user opens an object store, a list of all the top-level folders in that object store is retrieved from CE for display. If the object store or folder contains an excessive number of items, both performance and usability can be adversely affected. Requiring the user to browse through a large number of folders is not recommended for usability reasons.

While the FileNet P8 architecture does not impose a limit for the number of supported folders at a given level, you can optimize performance by using no more than 100-200 folders at a given level. For applications that require many folders, the IBM Performance Lab recommends using the search APIs and filtering to limit the results. In addition to performance considerations, limiting the number of folders at any given level enables the use of a more user-friendly interface by providing a browsing-style interface that does not require the user to enter a filter condition.

Although the number of folders displayed per page in Workplace is configurable via Site Preferences, the number of folders or subfolders retrieved by Workplace is not configurable. That is, any folder opened through Workplace will result in the retrieval of all subfolders at that folder level, regardless of whether they are displayed through the user interface.

## Documents and custom objects

As noted above, retrieving the contents of a folder can be a memory-intensive operation for custom applications and for Workplace. Each time a user opens a folder, Workplace retrieves the entire contents of the folder along with several properties for each item. If the folder contains an excessive number of documents or custom objects, system performance can be adversely affected.

Workplace does provide filtering options for items other than folders. You can use the Folder Filter options in Site Preferences and User Preferences to limit the number of documents and custom objects retrieved by Workplace. The Folder Filter options limit how many documents and/or custom objects are retrieved per folder access, with a default value of 100 per object type. If a folder contains more items than specified by the Folder Filter options, then the remaining items can be retrieved by entering criteria in the Browse page Filter folder field to retrieve only those objects that match the criteria. Overhead from the initial retrieval operation can be significantly reduced by applying folder filtering. For example, set the filter criteria so that only document names added by a specific user or documents containing specific words in the title will be returned, rather than all items in the folder.

## Recommendations

To minimize the affects on performance, the IBM Performance Lab recommends:

- When possible, limit the number of subfolders at every folder level. Create a hierarchy of subfolder levels rather than using many folders at the same level. Accessing additional levels of folders does not hinder the performance of the system, but accessing many folders at the same level does affect performance. In addition to performance considerations, the user experience can be improved by limiting the number of items displayed and by not requiring the user to enter filter criteria.
- Limit the number of server round trips by using EM to set the Enumeration Batch Size parameter and the Default Query Batch Size parameter to no more than 50 for each object store. This number reduces the number of round trips to the server, which should have a positive affect on performance.
- Use the FileNet P8 search APIs to filter the search results when the application requirements call for large numbers of folders or objects within folders.

getContainees() can return a large result set that affects performance. You can limit the result set size by requesting only the object types or properties that are needed, rather than requesting all items.

Ad hoc queries can be more efficient than using getContainees().

Use FROM and WHERE clauses to more precisely define the result set.

Use JOIN and SQL parameters to target the results.

Limit the number of results rows returned.

Use multiple queries. Several queries can be faster than a single getContainees() call.

Sample query:

```
SELECT Id, DateLastModified, IsCurrentVersion
FROM Document
WHERE Document.This INFOLDER '/wjc'
AND Id {...?...}
ORDER BY Id
```

- Use folder filtering whenever possible. Set the number of items to retrieve before filtering to be no higher than the default value of 100 for each type (documents and custom objects), and then test your system to find an appropriate maximum value. The default of 100 might be too conservative in some environment.

- Make filtering more efficient by assigning unique names to documents whenever possible. Avoid assigning similar names to multiple documents and/or custom objects.

For example, the document names in the left column of the following table might be more difficult to filter than the names in the right column if you are filtering documents based on the document title.

Loan_Origination_Application_John_D_Doe	John_D_Doe_application
Loan_Origination_Application_John_Q_Public	John_D_Doe_Loan_Terms_Complete
Loan_Terms_Complete_John D_Doe	John_Q_Public_application

### Avoid index skew

**NOTE** An index skew is an uneven distribution of values in an index. For example, if 1,000,000 documents are filed in the system, and half of those documents are filed in one folder, half of the rows in the index will have the same `tail_id` value.

When an index skew occurs, the database might not use the index, and instead might perform a full table scan because the logic behind the database optimizer is that if too many values are the same, it will take more time to scan the index and then lookup the rows than just reading the rows in the table. The database optimizer has this logic because it is not passed the values of the search, which is the case in CE that uses bound parameters.

Bound parameters, which allow for query plan re-use for queries re-executed with the same SQL statement but with different search values, are essential for transaction processing systems which frequently execute queries but cannot afford to have a query plan recompiled on each execution.

Because the optimizer uses a general plan for all values, it has no way to determine if the searched values are one of the few values that do not have a large amount of duplicates. If the search is in a folder that has a large number of documents, it is slightly better to use the full table scan to retrieve the documents. However, if the search is in a folder that has only a few documents, it is much worse to use the table scan. Therefore, the following solutions are presented to eliminate the table scans. A common use-case is a customer that knows that they should never browse a large folder.

Therefore, avoid placing a large number of objects in one folder regardless of whether they are subfolders, documents, or custom objects. A large numbers of objects in one folder relative to the other folders creates an index skew, which can affect:

The symptom of skew occurs when browsing a folder is too slow or times-out due to a query plan that uses table or index scans.

Specifically, the situations that can encounter poor performance while browsing are:

- Browsing operations
- INSUBFOLDER searches
- Large numbers of subfolders directly under a parent
- Large numbers of documents or custom objects filed under one folder
- Due to large numbers of subfolders directly within a folder, the index `I_CONTAINER24` on Container (`parent_container_id`, `name`) will be skewed
- Due to large numbers of objects filed in a folder, the index `I_RELATIONSHIP27` on Relationship (`tail_id`, `name`) will have skew

## Performance diagnosis

Obtain a catalog trace and identify the slow queries to see if the slow queries are either of the following then obtain the histogram as indicated:

### Browse getContainees() for documents or custom objects (replaced with FROM Generic for custom objects):

Form of query:

```
SELECT columns FROM DocVersion T1
WHERE T1.home_id IS NULL
and ( T1.object_id in (select r.head_id from Relationship r where r.tail_id = ?) )
and ( T1.uc1_isdeleted = 0 )
```

Alternate form of query:

```
SELECT columns FROM docversion t1
INNER JOIN relationship t2 ON (t1.object_id = t2.head_id )
WHERE t1.home_id IS NULL AND t2.home_id IS NULL AND ( (t2.tail_id = (:1))
AND ((t1.mime_type IS NULL) OR (t1.mime_type != :2)) )
ORDER BY t1.u2e_documenttitle
```

(variations might exist in the predicates for the above)

Obtain this histogram:

```
select count, tail_id from relationship group by tail_id order by 1 desc;
```

### Browse containees for subfolders:

```
SELECT object_id , object_class_id , epoch_id , home_id , security_id , creator ,
create_date , modify_user , modify_date , lock_token , lock_timeout , lock_owner ,
security_policy_id , parent_container_id , name , u1_ishiddencontainer ,
u12_containertype, lock_sid , parent_container_id , name FROM Container WHERE
parent_container_id = ?
```

Obtain this histogram:

```
select count, parent_container_id from container group by parent_container_id order by 1
desc;
```

The histogram output shows the counts of each group by column, thus allowing the skew to be observed. If one folder has more than half of the rows with the same tail\_id or parent\_container\_id value, that index is heavily skewed.

Obtain a query plan from the shared pool after the SQL is executed by the CE. Do not use Oracle SQL tools such as Explain Plan or TOAD or TOP\*SQL. Obtain the query plan directly from the v\$sql\_plan views as explained in the Oracle Query Analysis Document (aka DBA Guide for CE), because otherwise those tools misinterpret the GUID datatype (stored as a raw(16)) and show a full table scan rather than index usage, a tool behavior observed frequently.

If the query plan for the operation shows a full table scan rather than the use of the relevant index, then the cause is most likely from index skew.

## Skew solutions for Oracle

Below are some common approaches that enable Oracle to use an index. The best approaches are listed first.

### Histogram stats

Histogram statistics allow the optimizer's "peek binding" technology (when a query plan is first compiled, the data or "bind" values are peeked at) to use the values and choose the best plan for the peeked values because it stores the distribution with the histogram statistics.

While the following statistics solution is written for the skew on relationship.tail\_id, the skew on container.parent\_container\_id would be solved analogously.

First compute the number of distinct values:

```
select count(distinct(tail_id)) as NUM_DISTINCT from relationship;
```

If NUM\_DISTINCT is less than 254 (the Oracle 9i bucket maximum), then use value-based histograms by specifying the SIZE value to be greater than the number of distinct values (while allowing for growth).

For example, with 16 distinct values, the size could be SIZE=20.

Next, try the following solution to solve the skew, assuming less than 20 distinct values in tail\_id:

```
Exec dbms_stats.gather_table_stats( ownname= 'P8User', tabname= 'DocVersion' ,  
estimate_percent= 20, cascade= TRUE, method_opt= 'FOR ALL INDEXED COLUMNS SIZE 20');  
  
Exec dbms_stats.gather_table_stats( ownname= 'P8User', tabname= 'Relationship' ,  
estimate_percent= 20, cascade= TRUE, method_opt= 'FOR ALL INDEXED COLUMNS SIZE 20');
```

This solution keeps the statistics uniform across the DOCVERSION and RELATIONSHIP tables. If this solution does not work, consider an experimental approach to determine if the statistics will work, which is trying estimate\_percent=100 and FOR ALL COLUMNS. If one or both work, reduce the estimate\_percent to determine the required threshold.

The statistics take trial and error to refine settings for the best performance.

### NOTES

- Verify the Oracle Statistics Job, if configured, does not run. It can overwrite your experimental changes. Also, change the job to use the new statistics when they are found.
- Execute the query (that is, the folder browse) on a folder that does not have many documents or subfolders. Otherwise, the optimizer peeks (at compile time only) at the "bad" values – for the large folder that needs the table scan --- instead of at the "good" values for which on a fewer set of values the index lookup is best, and the bad plan will become cached and reused for the good values. By searching on the good values first, the good plan is cached and re-used for all queries of that form.
  - Never browse a large folder first after running the above statistics because the query plan can be marked as "Obsolete" or "Flushed" at any time, such as when new statistics are generated either by a background job or manually. Browsing a large folder first after the new statistics, would prime the cache with a plan that will be very slow for a small folder (e.g. a table scan).

## Behavior on SQL Server

Rather than "bind peeking", SQL Server has "parameter sniffing", which is very similar to "bind peeking". At query plan compile time, the optimizer examines the parameters. Subsequent parameterized query executions of the same query form, but with different data, do re-use the same query plan, which can be observed in a profile trace as a "cache hit".

**NOTE** One slight difference between these two methods is that with variable length data such as varchar strings, there will be a recompiled query plan even with parameters for any data that is different in length for each execution (the OLEDB provider may request this internally, based on previous ODBC behavior). The behavior of the variable length data is the same as it is for literals, which can invoke auto-parameterization.

The SQL Server auto-statistics always generates histograms. Use the `dbcc show_statistics` command to see the chosen bucket ranges, as well as the sampled rows. It might be necessary to run the `create statistics` and `update statistics` commands manually to increase the sample size if skew problems occur (that is, query plans with table or index scans rather than index seeks). Examine the rows sampled from the `dbcc` command and increase the amount if query plans are not correct.

Following are two techniques that can be used in SQL Server while browsing a large folder either before or after statistics are out-dated. However, they result in a flush of the query plan then recompile on new execution, which can result in a poor plan being cached:

- Set the autostats to off
- Use the `norecompute` option in the `create statistics` command

Either choice essentially freezes the statistics in place, which might permit poor plans if data changes. Therefore, you need to monitor performance because newly added data could out-date the query plan, which might benefit from new statistics. The objective is to have the smaller folders browsed first.

For manual statistics gathering, SQL Server does have the option of specifically specifying the column to be analyzed, so `tail_id` could be histogrammed at a fine granularity specifically, thus avoiding an expensive statistics gathering on more columns than is necessary.

# Business Process Management

## Process Engine

### Databases

#### Insufficient Java heap size causing “out-of-memory” errors

On the AIX operating system, you might encounter Java out-of-memory errors due to too low Java heap size when working with Large Binary Objects (BLOBs) on Oracle or DB2 databases.

To accommodate for the memory requirement you must raise the Java heap size.

**NOTE** When setting the Java heap max value to 256 MB or greater on AIX you must set an environment variable that allows VWJs to make use of the AIX large address-space model.

The environment variable to set is:

```
LDR_CNTRL=MAXDATA=0xn0000000
```

where *n* is a value between 2 to 8.

By default this value is zero, which means that only 256 MB (1 data segment) is reserved for the Java process stack and data (Java heap). The memory available to the Java heap is less than 256MB if you don't use the LDR\_CNTRL=MAXDATA environment variable.

**CAUTION** By increasing the value of LDR\_CNTRL=MAXDATA=0xn0000000, you shrink the size of shared memory segments available for use by VWJs.

#### Maximum Database Operations per Transaction and Maximum Instructions per Transaction parameter values

Two process parameter values set limits on Instruction Sheet Interpreter (ISI) activity by determining when the ISI has to cease processing a work item:

- Maximum database operations per transaction – Default setting is 100.
- Maximum instructions per transaction – Default setting is 200.

#### Scenarios for changing the parameters

##### ***Workflow maps executing in an isolated region***

*Description:* Workflow maps executing in the isolated region contain 750 to 1000 system instructions between user-defined instructions.

*Recommendation:* To maximize the availability of the ISI and other server resources, decrease the value of the Maximum Instructions Per Transaction parameter.

##### ***InstructionSheetInterpreter system queue is filling up***

*Description:* The InstructionSheetInterpreter system queue is filling up.

*Recommendation:* Increase the value of the Maximum Instructions per Transaction parameter.

**NOTE** You can access the InstructionSheetInterpreter system queue via vwtool, APIs, or the Process Administrator application.



### ***Database audit trail size has been exceeded***

*Description:* The size of the database audit trail has been exceeded, preventing further database activity.

*Recommendation:*

- Decrease the value of the Maximum Database Operations Per Transaction parameter.
- See the database vendor's documentation for information on optimizing or increasing the amount of space available for the database audit trail.

### ***Step processor process time is excessive***

*Description:* The amount of time between one-step processor dispatching a work item and the next step processor finding and locking that work item is excessive.

*Recommendation:* Decrease both parameter values.

### **Adjust the parameter values**

1. Open the Process Configuration Console.
2. Click the Isolated Regions properties.

For most PE systems, these default settings are sufficient. However, certain scenarios might require adjustments to these settings.

## **JVM tuning**

### **Adjust RMI triggers for garbage collection**

**ATTENTION** This tuning tip is for IBM and Sun JVMs only, and applies to the PE 3.5 clients of all applications that use PE.

In Java 1.4.2 and 1.5, Remote Method Invocation (RMI) triggers full garbage collection once per minute. This frequency can cause excessive CPU consumption, heap fragmentation, and poor application response times.

To prevent this issue, pass the following custom JVM arguments to the application server, where *x* is the interval in milliseconds. The default value is 60000ms, a suggested starting value is 1800000, which equals 30 minutes.

-Dsun.rmi.dgc.server.gcInterval=*x*

-Dsun.rmi.dgc.client.gcInterval=*x*

### **Sun WASP WSI server and client JVM**

Sun Microsystems has a version of Java-based JVM for server-side performance and a version for client-side performance. These two solutions share the Java HotSpot runtime environment, but have different compilers suited for the distinctly different performance characteristics of clients and servers.

The server-side JVM has shown significant performance improvements over the client-side JVM in FileNet P8's highly scalable server environment.

## Process Task Manager

### Minimum and maximum heap size for high volume Component Manager activities

To avoid out-of-memory errors during high volume Component Manager activities, such as RM Destroy Disposition Sweeps and products or customizations that use component queues without additional modifications like CE\_Operations.

1. Open Process Task Manager.
2. Select Component Manager.
3. Click the *Advanced* tab.
4. Set the value of the JRE Parameters field to at least `-Xms1024m -Xmx1024m`.

### VWJ process optimization

PE is implemented as a collection of single threaded worker processes called VWJs. Having a sufficient number of worker processes available improves concurrency. All incoming requests are routed to an available process.

The default maximum number of VWJs is 5. The IBM Performance Lab suggests you allocate 1.5 to 2 VWJs per CPU when the total number of CPUs is less than 8 and 2 to 5 VWJs per CPU when the total number of CPUs is greater than 8.

### Adjust the number of VWJs

1. Open the Process Task Manager.
2. Select Component Manager from the.
3. Enter the new default number of VWJs in the “Number of VWJs Processes” field.

### NOTES

- If the existing pool of VWJs is not over-subscribed (that is, of the existing VWJs have minimal CPU time) making this change will not improve performance.
- VWJs can have a substantial memory footprint on the system. Monitor your system memory resources to avoid performance problems.

### When to add more silos

Individual VWJs coordinate access to shared memory resources within the PE. If you significantly increase the number of VWJs, for example from the default of 5 to 10, you might find that the individual VWJs are not able to do as much work as you expected, and you might need to increase the number of silos.

Additionally, if there are processes waiting for memory interlock, you might want to increase the number of silos. To determine if VWJ processes are waiting for memory interlock, run the `vwspy` program and use its “state” command. Look for lock or semaphore contention among the various VWJs. Run the “state” command several times and in rapid succession over several minutes when the system is busy. If you consistently see a VWJs process holding a lock and you see multiple other VWJs stacks waiting for that same lock, you probably need to increase your silos.

Having more silos breaks the shared memory into more separate pieces so that VWJs do not block each other when they update this memory.

### **Adjust the number of silos**

1. Create a simple text file called vwsilos.cfg that contains the new number of silos
2. Save the vwsilos.cfg file in the \fnsw\_loc\sd\ directory of a Windows-based PE server (or in the /fnsw/local/sd/ directory of a UNIX-based PE server).
3. Restart the PE server.

### **When to not add more silos**

In general, if vwspy's "state" command does not show any lock contention, do not change the silos. If you change the silos, it is generally best to keep the number of silos within 75% of the number of VWJs.

**CAUTION** More silos can lead to memory fragmentation and other problems. Do not increase the number of silos without carefully evaluating the new setting to confirm that the new setting is an improvement over the previous setting.

### **Java RMI Backlog optimization**

The Java RMI Backlog represents the maximum queue length for incoming Java RMI requests. If a request arrives when the queue is full, the connection is refused.

Users should edit the Java RMI Backlog parameter in two situations:

- When experiencing lengthened response times under high loads.
- When the message could not locate available server processes appears in the Process Task Manager.

The default value for Java RMI Backlog is 100 and should be increased in increments of 100.

**CAUTION** Increasing the Java RMI Backlog is only suitable to deal with short-term spikes in response times. That is, if the system is consistently performing slowly, increasing the Java RMI Backlog is not likely to improve the systems performance.

### **To adjust the Java RMI Backlog**

1. On the Process Engine server, launch the Process Task Manager.
2. Select PPM.
3. Select the *Advanced* tab.
4. Enter a new value in the RMI Backlog field.
5. Stop and restart PPM to ensure the new value takes affect.

### **Cached Entry Timeout parameter: User and group cache**

The Cached Entry Timeout parameter specifies how often (in hours) PE refreshes user and group cache. Caching this information causes PE to make fewer calls to CE (and the LDAP server) and provides better throughput. However, system response times might increase when the user environment cache refreshes. For example, caching systems that have roster security or that contain queues with complex security specifications or step definitions that refer to large user groups (several hundred members or more) could affect system performance. Therefore, because user and group cache refreshes might degrade system performance, controlling the frequency of refreshes also maximizes PE performance.

The default value is 4 hours. Base this setting on how often the directory service changes (additions and deletions of users and groups, or modifications to group memberships).

If changes rarely occur on the Directory Service, set the Cached Entry Timeout to 0 and manually update the cache as needed. If changes occur quite frequently, set the Cached Entry Timeout parameter to the number of hours by which PE should execute a refresh.

### **Adjust the Cached Entry Timeout parameter**

The frequency by which PE executes a refresh is specified by the Cached Entry Timeout parameter value in the Advanced tab, in the PE server's node properties, in Process Task Manager.

### **Manually refresh the user and group cache**

You can manually refresh user and group cache by using vwtool's "*env if m y*" command or by setting it up to refresh automatically.

### **Schedule cache refresh**

If you want cache refreshes to run during certain off-peak hours, set the Cached Entry Timeout parameter to 0 and use the appropriate operating system tools to run the vwtool *env if m y* command as a scheduled task.

### **Flush Granularity option: Cache refresh**

As a part of routine PE server operations, the user environment cache is periodically refreshed. You can control the refresh period by using the Process Engine Task Manager option "Cached Entry Timeout".

During the cache refresh period, a difference in system performance might be noticeable; throughput might drop and response times might rise. This situation can occur if, for example, the system contains queues with complex security specifications that refer to large groups (several hundred members or more). This situation might also occur with roster security, and in general when dealing with large groups in security specifications or step definitions

The administrator can add that setting to the vwserver.ini file (flushGranularity= number of seconds) and the engine uses it after the PE system is restarted.

The default value for flushGranularity is two seconds. The optimal setting for flushGranularity differs depending upon complexity of ACLs on the queues, speed of the directory service, the setting for the user environment cache timeout (in the Process Engine Task Manager), and system workload.

- For a predictable but potentially more severe effect on the system, use a low value for flushGranularity.
- For a sporadic but dampened effect on the system, use a higher value for flushGranularity.

### **WorkSpaceCacheSize parameter: Cache isolated region workspaces**

Each isolated region caches workspaces. By default, the cache holds 10 workspaces, which is enough in most situations.

However, you can monitor the cache size by using the loadstatus command to determine how frequently the cache size is exceeded.

### **Adjust the workspace cache size**

If necessary, you can increase the size of the workspace cache by specifying a WorkSpaceCacheSize parameter value, which is in the Advanced tab, in the PE server's node properties, in Process Task Manager.

**NOTE** If the new value is too high, the PE server will exceed available memory and fail. To prevent the PE server from failing, lower the size of the workspace cache.

## Use Loadstatus to monitor server load

The loadstatus command in vwtool outputs statistics on the load being placed on the server for all isolated regions. Use loadstatus periodically to get a feel for the rate at which requests or errors normally occur on the system. Once you have determined this rate, execute loadstatus when response time is slow to see if any of the error counters are going up at a rapid rate.

The loadstatus command outputs two types of statistics:

- **Remote procedure call (RPC) counters** – This covers the number of RPCs of various types that have been processed on the server.
- **Error counters** – This covers how many times a function must be retried due to an error on the server.

## PE Event Log purge best practice

Over time, accumulated event log records in the PE Event Logging table can have an adverse effect on performance. Therefore, the IBM Performance Lab recommends that you use the vwlog utility to schedule a reoccurring purge job to keep the records in the PE Event Logging table down to a reasonable amount. (Less than 10 million rows are suggested.)

For information about scheduling a purge job using the vwlog utility, see the following help topic in the FileNet P8 online help: FileNet P8 Documentation > Process Engine Administration > Administrative Tools > vwlog > Purging data from the Event Logging table.

To download documentation from the IBM support site, see “[Access IBM FileNet documentation](#)” on page 12.

## Process Simulator (PS) and Process Analyzer (PA)

### Max Simulation Jobs parameter: Concurrent simulation runs

The FileNet P8 default settings might not be optimal for your hardware setup. For example, if your server hardware contains more processors than are configured by the default settings, you can increase the Max Simulation Jobs value so that the server can handle a larger number of concurrent simulation runs.

**NOTE** Increasing this value will most likely require an increase in the Java max memory.

### Java max memory

Increase this value if you experience memory errors when running a simulation that has a high number of active work items at any moment or when concurrently running many simulations concurrently.

### Tuning Process Analyzer for optimal performance

The performance of the PA is measured by the publishing rate. The higher the rate, the better the performance. Factors affecting the publishing rate are number of publisher threads, cache size, database, and publishing interval. Each factor is described below.

#### Publisher threads

The PA uses multiple threads to publish event data. On a multiple CPU system, increasing the number of publisher threads can improve the publishing rate, as long as the disk I/O and the database are not bottlenecks, the CPU utilization is not approaching 100%, and sufficient memory is available.

**NOTE** Increasing the number of publisher threads will most likely require an increase in the Java max memory.

#### Cache size

A low cache hit ratio can affect the publishing rate. If the cache-hit ratio is low (that is, 50%), increase the size of the corresponding cache until the ratio improves or until the memory limit is reached.

#### Databases

PA monitors and tracks the average time for updating the PA database, processing the work item cube, and processing the workflow cube. PA updates the database and processes the work item and workflow cubes (that is, the “work in progress” cubes) at the end of every publishing interval (default: every 5 minutes). If any of these three operations takes longer than one minute, you should monitor the PA database and tune it as necessary.

The processing of historical cubes can take a very long time (hours). You should configure the historical cube processing to occur at the earliest off-hour time to allow sufficient time to process the cubes; the default processing time for the historical cubes is 11:00 PM.

If the historical cube processing takes too long, consider upgrading your disk hardware or reducing the size of the historical fact tables by using the PA Compression Wizard.

## Publishing interval

In general, the PA processes events for the amount of time specified in the publishing interval; the default is every 5 minutes. After the publishing interval is reached, the PA spends time updating the PA database and process the “work in progress” cubes. When the updating and processing completes the PA switches back into its event-processing mode and the process repeats.

If PA has a large backlog of events to process and you have already tried the other performance tuning options, consider increasing the publishing interval to give PA more time to process events.

**NOTE** Increasing the publishing interval directly affects how often the “work in progress” cubes are updated and processed. Configure the publishing interval relative to your needs for having up-to-date “work in progress” cubes. For example, if it is sufficient to have the “work in progress” cubes updated every hour, then change the publishing interval to 60 minutes.

# Application Engine

## Operating systems – Solaris

### Processing time

#### tcp\_time\_wait\_interval and tcp\_conn\_req\_max\_q parameters: TCP connections

When Application Engine (AE) processes a large number of requests, you might receive “Java.net.BindException: Address already in use” errors. These errors occur when the Java process has run out of available TCP connections. You can prevent these errors and improve TCP socket performance by decreasing the `tcp_time_wait_interval` parameter value or increasing the `tcp_conn_req_max_q` parameter value, or by adjusting both parameters.

#### Decrease the `tcp_time_wait_interval` parameter

The `tcp_time_wait_interval` parameter value specifies how long a TCP port is unavailable after both the client and the server have closed a connection. Temporarily disabling the port prevents new programs from inadvertently receiving packets that were meant for the previous program.

The parameter default is 240,000 milliseconds (4 minutes). By decreasing this value, TCP connections can be recycled more quickly.

**CAUTION** Changing the `tcp_time_wait_interval` to a value lower than its default value creates a TCP protocol violation.

#### Decrease the TCP wait time

Use the following command to adjust the TCP wait time:

```
ndd -set /dev/tcp tcp_time_wait_interval 240000
```

#### Increase the `tcp_conn_req_max_q` parameter

The `tcp_conn_req_max_q` parameter value specifies the maximum number of completed connections that can be queued while waiting to return from an accept call.

The parameter default is 128. By increasing this value, Java can use more TCP connections.

**CAUTION** Each open TCP connection occupies some memory (usually both bookkeeping tables and buffer space), and fewer TCP connections are available to other processes.

**NOTE** The memory footprint is usually only a minor concern for enterprise-class servers with large amounts of available RAM.

#### Increase the maximum number of queued completed connections

Use the following command to adjust the maximum number of allowable completed connections that can be queued:

```
ndd -set /dev/tcp tcp_conn_req_max_q 128
```



## Operating systems – Windows

### Processing time

#### TcpNumConnections, MaxUserPort, and TcpTimedWaitDelay registry parameters: TCP connections

When AE processes a large number of requests, you might receive “Java.net.BindException: Address already in use” errors. These errors occur when the Java process has run out of available TCP connections. You can prevent these errors and improve TCP socket performance by increasing the MaxUserPort and TcpNumConnections parameter values and/or by decreasing the TcpTimedWaitDelay parameter value in the registry.

#### TcpNumConnections registry parameter: Number of open TCP connections

The TcpNumConnections parameter value specifies the maximum number of connections that TCP can have open simultaneously.

**CAUTION** By increasing this parameter value, each open TCP connection occupies some memory (usually both bookkeeping tables and buffer space).

**NOTE** The memory footprint is usually only a minor concern for enterprise-class servers with large amounts of available RAM.

#### Adjust the maximum number of open TCP connections

Use the TcpNumConnections parameter to adjust the number of allowable open TCP connections. You can find this parameter in the following registry path:

```
HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters
```

#### MaxUserPort registry parameter: Port number assignment

The MaxUserPort parameter value specifies the highest port number that TCP can assign when an application requests an available user port from the system. Increase this value when CE requires high throughput transaction processing.

#### Adjust the port number that TCP can assign

Use the MaxUserPort parameter to adjust the number of allowable open TCP connections. You can find this parameter in the following registry path:

```
HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters
```

#### TcpTimedWaitDelay registry parameter: Release closed connections

The TcpTimedWaitDelay parameter value specifies the amount of time that must elapse before TCP can release a closed connection (TIME\_WAIT state) and reuse its resources.

Decreasing this value decreases the amount of time that a connection stays in the TIME\_WAIT state.

**CAUTION** Decreasing the default TcpTimedWaitDelay value creates a TCP protocol violation. The IBM Performance Lab highly recommends that you first try adjusting the MaxUserPort and TcpNumConnections values.

#### Adjust the amount of elapsed time before TCP can release a closed connection

Use the TcpTimedWaitDelay parameter to adjust the amount of elapsed time before TCP can release a closed connection. You can find this parameter in the following registry path:

```
HKLM\System\CurrentControlSet\Services\Tcpip\Parameters
```

## Workplace

### Databases

#### Indexing common to all database

To improve Workplace performance, consider the following indexes.

For overview information, see "[Databases – Indexes](#)" on page 13.

Database Table Name	Database Column Name	FileNet P8 Property Name	Tool	Description/ Reason to create index
docversion	object_class_id ----- uxy_searchtype ----- is_current	N/A	RDBMS	Create a composite properties index when:  The number of records in the table exceeds 500,000.  Performance degradation is occurring across common functionality (browsing, searching, creating documents, et cetera).
docversion	mime_type	MIME Type	RDBMS	Create a single property index when:  The number of records in the table exceeds 1,000,000.  The number of items in a folder exceeds 2,000.  Performance degradation is occurring across common functionality (browsing, searching, creating documents, et cetera.)
generic	applicationname		RDBMS	Create a composite properties index when:  The number of records in the table exceeds 500,000.  Performance degradation is occurring across common functionality (browsing, searching, and log on)
docversion ----- container ----- generic	security_policy_id		RDBMS	Create a single property index when:  You encounter slow performance or errors when attempting to delete a security policy on an object store containing many files (1 million or more).

## Searches that return large result sets

If the Workplace application receives many search requests that return 50 or more results, users might experience slow response times and/or timeout errors even if they are accessing non-search functionality. In the IBM Performance Lab, these failures begin to surface at around 20 searches per minute per AE server, but can occur at lower numbers of searches depending on hardware.

The symptoms of this problem are a high thread-queue buildup on the AE and an AE CPU usage that drops to zero even though there are active requests waiting to be executed.

**NOTE** This problem is specific to SQL Server and can be corrected by applying the patch described in the following Microsoft Knowledge Base article: <http://support.microsoft.com/kb/884554>.

You must contact Microsoft directly for this fix because it is not available (but will be included in SQL Server SP4).

## Folder browsing

### Workplace path control

When using the Workplace application, the number of folders in a single level of your folder hierarchy will have an impact on browsing performance. If you experience slow browse times when navigating a tree that has many subfolders at the same level, you can improve performance by disabling the Workplace Path Control in the Administrator site.

#### To disable Workplace path control

1. Sign on to Workplace with an administrative account.
2. Select the Admin tab.
3. Click Site Preferences.
4. Navigate to the General Settings > Application section.
5. Change the “Enable browse path tree control” option to No.
6. Click Apply, and then click Exit.

## Documents and custom objects

When using the Workplace application, the maximum number of documents and custom objects retrieved when opening a folder is configurable via the Site Preferences.

For more information, see “[Recommendations for handling a large number of folders, documents, and objects](#)” on page 34.

## Java client cache settings

Business Process Manager (BPM) applets, such as the Process Designer, require all PE jar files to be cached on the client machine, which can consume as much as 35 MB of disk space. Java cache settings on the client machine might cause these applets to be purged, causing the applet download process to reinitiate, and slowing end-user response times.

By tuning your client-side Java cache settings, you can minimize performance degradation and the number of times these applets are downloaded.

**NOTE** Some applets, such as the Process Tracker and Process Administrator, have been known to experience problems if insufficient Java memory is available on the client.

## eForms

### Rendering large templates to PDF

Rendering large templates to PDF can cause the print server to fail and to return an error message similar to the following: "Maximum request length exceeded received."

If you expect to render large templates (4MB or larger) to PDF, add the following content to the Web.config file, which is located on the Windows server that hosts the Print Services.

**NOTE** Add this content to the latter part of the file, near the Globalization setting.

```
!-- HTTP RUNTIME SETTINGS

This setting is used to increase the Maximum amount of data that can be contained within a
request
Object. It is specified in number of KB

--

httpRuntime
maxRequestLength="10240"
/

Cycle IIS
```

# Email Manager

## Databases

### Indexing common to all databases

For overview information, see “[Databases – Indexing](#)” on page 13.

Consider the following indexes to improve query performance.

Database Table Name	Database Column Name	FileNet P8 Property Name	Tool	Description/ Reason to create index
N/A	N/A	From	EM	Add a single-value property index to enhance query performance.
ListOfString	element_value ----- parent_prop_id ----- parent_id	To	RDBMS	Add a composite properties index to enhance query performance.  For more information, see “ <a href="#">Databases – Indexing multi-value property tables</a> ” on page 12.
ListOfString	element_value ----- parent_prop_id ----- parent_id	CC	RDBMS	

**NOTE** Run database statistics after creating the indexes.

### Force case insensitive

For strings on Oracle and DB2, if Force Case Insensitivity is enabled, the index must be created on lower(element\_value) function.

For more information, see “[Force Case Insensitive Search check box: Case-sensitive queries](#)” on page 21.

# Records Manager

## Databases

### Indexing common to all database

To improve Records Manager (RM), consider the following indexes.

For overview information, see "[Databases – Common to all databases – Indexing](#)" on page 13.

Database table name	Database column name	FileNet P8 property name	Tool	Description/ reason to create index
docversion	<i>uxy_searchtype</i>	N/A	RDBMS	Create a composite properties index when both of the following conditions are true:  The number of records in the table exceeds 500,000.  Performance degradation is occurring across common functionality (browsing, searching, creating documents, et cetera.)
generic	applicationname		RDBMS	Create a composite properties index when both of the following conditions are true:  The number of records in the table exceeds 500,000.  Performance degradation is occurring across common functionality (browsing, searching, and log on)

### Indexing SupersededBy property for improved performance

**NOTE** This tuning tip applies to RM versions that were released prior to version 3.7.0.

If RM takes a long time obtaining the values of the RM enumeration property "Superseded Records", and obtaining such values is done frequently enough to necessitate improved performance, you can create the following single index by using EM.

Class	Property symbolic name	Tool	Description or reason to index
Objectstore/FPOS/Document Class	Superseded By	EM	Single.

### Indexing for improved sweep performance

To increase retention calculation (sweep) performance, IBM strongly recommends that you use your RDBMS management software to add an index to the the “symbolic\_name” column of the “ClassDefinition” table in your file plan object store (FPOS) database.

Table	Symbolic name	Tool	Description or reason to index
clasdefinition	Symbolic name	EM	Single

#### NOTES

- You do need to index a property that is only involved in conditions that use the 'like' operation.
- Generally, indexing a property that is used in a condition will have a positive impact on sweep performance. There are, however, cases where indexing the property can have a negative impact on performance.
- Before setting up an index you should consider the following conditions for your environment:
  - The size of the data held by the property
  - The frequency that the property will be involved in write operations such as inserts, updates, and deletes
  - If the property is involved in frequent write operations and the amount of data being indexed is significant, then indexing the property can cause overall performance degradation.

### Indexing for disposition and hold sweeps

*(MS SQL only)* Changing the clustered index on the docversion table in the file plan object store from version\_series\_id, major\_version\_number, and minor\_version\_number to object\_id significantly increases document retrieval performance during a disposition sweep or a hold sweep.

*(File plan object stores only)* For optimal performance, add single-property indexes on propertyname and applicationname in the generic table in a file plan object store, as shown in the following table. Indexing these properties can increase search performance and avoid time-out when running disposition sweep in a large-sized database.

Table	Property display name	Property symbolic name	Columns	Tool	Description or reason to index
generic	N/A	Property Name	uxx_ propertyname	RDBMS	Single
generic	N/A	Application Name	uxx_ applicationname	RDBMS	Single

### Indexing for improved Property page performance

If you have a large number of records in your repository you might experience degraded performance when navigating various RM property pages. To improve property page performance consider adding the following indices.

### Actions Property page

Add composite indexes to the following docversion table by using these properties in the order the rows are listed below.

Table	Property display name	Property symbolic name	Columns	Tool	Description or reason to index
docversion	Vital Record Review Action, Current Phase Action	VitalRecordReview Action, CurrentPhaseAction	uxx_vitalrecordreviewaction, uxx_currentphaseaction	RDBMS	Composite

**NOTE** The value xx will vary depending on the data model/database type being used.

### Location Property page

Use your RDBMS management software to add a composite index to the *DocVersion* table by using these properties in the order the rows are listed below.

Table	Property display name	Property symbolic name	Columns	Tool	Description or reason to index
docversion	Barcode, RM Entity Description, RM Entity Type	BarcodeID, RMENTITYDescription, RMENTITYType	uxx_barcodeid, uxx_rmentitydescription, uxx_rmentitytype	RDBMS	Composite

**NOTE** The value xx will vary depending on the data model/database type being used.

Use EM to create single indexes on the object store class properties in the order the rows are listed below.

Class	Property/symbolic name	Tool	Description or reason to index
Objectstore/FPOS/Document Class/Record	Location	EM	Single
Objectstore/FPOS/Document Class/Record/Marker Record	Home location	EM	Single
Objectstore/FPOS/Other Classes/Folder/RM Folder/Record Folder	Location	EM	Single
Objectstore/FPOS/Other Classes/Folder/RM Folder/Record Folder/Hybrid Record Folder	Location	EM	Single
Objectstore/FPOS/Other Classes/Folder/RM Folder/Record Folder/ Physical Record Folder	Location	EM	Single



Class	Property/symbolic name	Tool	Description or reason to index
Objectstore/FPOS/Other Classes/Folder/RM Folder/Record Folder/ Physical Container	Location	EM	Single

### Record Type Property page

Use EM to create a single index on the object store class property in the order the rows are listed below.

Class	Property symbolic name	Tool	Description or reason to index
Objectstore/FPOS/Document Class/Record	Type ( <i>DoD and Base</i> ) or Record Type ( <i>PRO</i> )	EM	Single

### Disposal Instruction Property page

Use EM to create a single index on the object store class property in the order the rows are listed below.

Class	Property Symbolic Name	Tool	Description or reason to index
Objectstore/FPOS/Other Classes/Folder	Disposition Instructions ( <i>DoD and Base</i> ) or Disposition Schedule ( <i>PRO</i> )	EM	Single

### Set the minimum and maximum heap size for Dynamic Hold Sweeps

To avoid “out-of-memory” errors while performing bulk operations in a batch during a Dynamic Holds Sweep, increase the heap size for the RM sweeps to 1GB. These values are set in the RecordsManagerSweep.bat/sh file.

#### NOTES

- The minimum and maximum heap sizes are not set by default in the script.
- The following example is for a server with 2GB available RAM. Substitute the number 1024 for the amount of RAM that you want to allocate for your setup.

### Set the minimum and maximum heap sizes for Dynamic Holds Sweeps

1. Identify the heap size entries.

In the RecordsManagerSweep script, search for:

- RecordsManagerSweep.bat :  

```
java -jar -Dwasp.location=./lib ./lib/%JarFile%
```
- RecordsManagerSweep.sh :  

```
java -jar -Dwasp.location=./lib ./lib/$JarFile
```

2. Set the minimum and maximum heap size to 1024MB for sweep execution.

o RecordsManagerSweep.bat :

```
java -Xms1024m -Xmx1024m -jar -Dwasp.location=./lib ./lib/%JarFile%
```

o RecordsManagerSweep.sh :

```
java -Xms1024m -Xmx1024m -jar -Dwasp.location=./lib ./lib/$JarFile
```

3. Save the file.

## Build a DLL instead of VBScript for a Record Declaration event action module

The following recommendation applies to users planning to use the RM event declaration sample script. This includes CFS customers as well as out-of-the-box RM customers who choose this declaration method.

If you are developing an event action module to declare records managed by RM, The IBM Performance Lab recommends building it as a DLL. Building an event action module as a DLL provides better performance than using a VBScript, with significant benefits noticeable when the action module performs complex operations.

**CAUTION** See the following topic in the FileNet P8 Developer Help for important information about developing event action procedures. In the online help, navigate to [Developer Help > Content Engine > COM API > COM API Guide > Working with Content Engine Objects > Event and Auditing Subsystem > Restrictions on Event Procedure Code](#). To download IBM FileNet documentation from the IBM support site, see [“Access IBM FileNet documentation”](#) on page 12.

## Oracle

### Prevent performance degradation when running Disposition or Dynamic Hold Sweep

In general, the dbms\_stats statistics collection mechanism should not be used with its default settings without first reviewing and understanding the impact of the custom setting options.

With RM, certain default settings result in poor query plans that degrade the performance of INSUBFOLDER searches, which are used for the Disposition and Dynamic Hold Sweep operations. At a minimum, the IBM Performance Lab recommends that both the dbms\_stats.gather\_schema\_stats command and the dbms\_stats.gather\_table\_stats command specify custom values for estimate\_percent and method\_opt, for example:

```
exec dbms_stats.gather_schema_stats(ownname = 'SCHEMA', estimate_percent = 20, method_opt = 'for all columns size 1', cascade = TRUE);
```

One of the following optional modifications might also be needed to improve performance:

- method\_opt = “for all indexed columns”
- method\_opt = “for all columns size *N*”
- method\_opt = “for all indexed columns size *N*”
- Increase the value of *M* in the estimate\_percent

**NOTE** Experiment with the estimate\_percent and method\_opt, and use the smallest value(s) that produce the best results.

## Content Federation Services (CFS)

### Federating documents

#### QueryBlockSize parameter: Size of query results

The queryblocksize parameter value specifies the allowable size of query results. This parameter works with the CFS “Chunking” feature, which enables large amounts of documents to be federated using a single rule by dividing large query results into smaller, more manageable size “chunks” in order to prevent out-of-memory errors.

The default value of the queryblocksize parameter is 2000, which should be sufficient for federating documents. If you encounter out-of-memory errors, try lowering this parameter value.

#### Adjust the queryblocksize parameter

Adjust the queryblocksize parameter in the “General Configuration” section of the FedAdmin user interface.

#### Batch Size parameter: Number of batched documents

The Batch Size parameter specifies the number of documents that is batched and processed by the exporter at one time.

The default value of the Batch Size parameter is 500. Increasing the Batch Size parameter value reduces the number of roundtrips between the exporter and the IICE RMIBridge when retrieving content from the source repository, and increases the performance of the exporter.

#### Adjust the Batch Size parameter

Adjust the Batch Size parameter in the “General Configuration” section of the FedAdmin user interface.

### CFS-IICE rules and application server memory

WebSphere and WebLogic have default timeout settings, and rules that Federate large amounts of content can exceed these timeouts. In general, it's best to create rules that will complete Export from your external repository in less than 30 minutes.

If you do encounter timeout errors in your SystemOut.log such as the ones below, you can try increasing the value for ‘Unused Timeout’ under Resources, JDBC providers, Cloudscape JDBC Provider, Data Sources, IICEServer.server1-iiceSIB, Connection Pools.

```
[1/31/06 14:38:03:984 PST] 0000000e ThreadMonitor W   WSVR0605W: Thread "ORB.thread.pool : 1"
(00000044) has been active for 738187 milliseconds and may be hung.  There is/are 1 thread(s)
in total in the server that may be hung.
```

```
[1/31/06 14:55:45:391 PST] 0000005c SystemOut      O WARN   - COEC0561W: Releasing overdue user
session (exceed maximum checkout time of 1,800,000ms).
```

## Web client time-out setting for viewing large federated documents

Increase the Web client timeout setting to prevent a system time-out when viewing large native contents of federated documents.

The Web client time-out parameter is a registry setting on the CE server.

### Configure the Web client time-out parameter

1. Locate the following registry key:

HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Services\CFS Importer Service.

2. Create a new DWORD registry value named "IICE Timeout".
3. Set the value of IICE Timeout to -1. (This value represents the time-out value in milliseconds. A -1 setting sets the value to an infinite amount of time.)

## RMI bridges

Using multiple RMI bridges increases the performance of the exporter.

### Configure multiple RMI bridges

1. Log in to the IICE Administration Tool on the IICE server.
2. Select the connector that is being used to connect to the external repository.
3. Enter the "RMI Proxy Connector URL" to add additional RMI entries.

## JVM heap size for CFS-CS

When using an RMI bridge with CFS-CS, setting the JVM heap size too high can cause federation to fail. The default JVM heap size for the RMI bridge is `-ms64m` and `-mx256m`, which should be sufficient for federating documents from Content Services (CS) to CE.

### Configure the JVM heap size for an RMI bridge and CFS-CS

1. Open the RMIBridge.bat file in a text editor.
2. Adjust the JVM heap size and save your changes.

**NOTE** Federation failure happens with the IICE CS connector HF0013 and older.

## Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation  
Licensing  
2-31 Roppongi 3-chome, Minato-ku  
Tokyo 106-0032, Japan

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:** INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation  
J46A/G4  
555 Bailey Avenue  
San Jose, CA 95141-1003  
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

#### **COPYRIGHT LICENSE:**

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

## **Trademarks**

IBM, the IBM logo, and [ibm.com](http://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. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, and service names may be trademarks or service marks of others.

## **U.S. Patents Disclosure**

This product incorporates technology covered by one or more of the following patents: U.S. Patent Numbers: 6,094,505; 5,768,416; 5,625,465; 5,369,508; 5,258,855.





Program Number: 5724-R76, 5724-R80, 5724-R81, 5724-S19

Printed in USA

GC31-5585-03

