



Performance Tuning Guide

Note

Before using this information and the product it supports, read the information in "Notices" on page 93.

Contents

Revision log.....	9
About this document	13
Tuning tip organization	13
IBM InfoSphere Enterprise Records and IBM FileNet Records Manager	13
Parameter settings	13
Assumptions	13
Recommended documentation	13
Application servers and Java Virtual Machine documentation	14
Database documentation	14
Directory servers documentation	15
Collaboration documentation	15
Accessing IBM FileNet documentation	16
Contacting customer support.....	16
Sending feedback	16
Operating systems	17
HP-UX.....	17
Java applications support	17
Solaris.....	17
Processing time	17
tcp_time_wait_interval and tcp_conn_req_max_q parameters: TCP connections	17
Decrease the tcp_time_wait_interval parameter.....	17
Increase the tcp_conn_req_max_q parameter	17
Windows	18
Processing time	18
TcpNumConnections, MaxUserPort, and TcpTimedWaitDelay registry parameters: TCP connections	18
TcpNumConnections registry parameter: Number of open TCP connections.....	18
MaxUserPort registry parameter: Port number assignment.....	18
TcpTimedWaitDelay registry parameter: Release closed connections	19
Windows 2008	19
CPU usage.....	19
Databases	20
All databases	20
Performance enhancing indexes	21
Indexing object value properties	23
Memory and disk I/O.....	23
Compute statistics for optimal performance	23
Reorganize indexes and tables	23
DB2.....	24
Performance enhancing indexes	24
Large object stores	24
DB2_WORKLOAD registry variable: Improve performance of Content Manager applications.....	24
DB2_OPTPROFILE	25
DB2_EVALUNCOMMITTED and DB2_SKIPINSERTED registry variables: Reduce deadlocking and improve application concurrency	25
APPLHEAPSZ parameter: Processing large SQL statements.....	26
MAXFILOP database parameter: Number of open file handles per database agent.....	26
Oracle	27
Indexing date properties	27
QUEUESIZE parameter: Maximum number of concurrent connections	27
Disable hash join to increase the response time of slow searches	28
Microsoft SQL Server	30
Large content uploads and downloads.....	30
Reclustering the docversion table.....	31
Application servers.....	32

WebSphere, WebLogic, and JBoss.....	32
Connection pools	32
Adjusting the connection pool maximum capacity setting.....	32
Thread pools	33
Adjusting the minimum and maximum thread pool settings.....	33
WebSphere.....	34
Lightweight third party authentication (LTPA) timeout: Token validity.....	34
Client-side security cache timeout: Cache refresh	34
Performance monitoring infrastructure PMI	35
Service logs	35
Local host lookups	35
WebLogic.....	36
Server and client JVM switches.....	36
Linux socket leak	36
JVM tuning.....	36
JVM tuning for LDAP connection pooling	36
Tuning for improved response time and CPU utilization	37
High-volume ingestion scenarios.....	38
JRockit tuning for improved response time and CPU utilization.....	38
Heap size	39
Native out-of-memory errors	39
JVM out-of-memory errors	40
IBM	40
JRockit.....	40
Directory servers	41
Tivoli	41
Optimize the database	41
.Net API client applications	42
Recycle Worker Process parameter (IIS 6.0 for .Net API)	42
Content Engine	43
Databases	43
Persisting multi-value properties in their own database tables.....	45
Force case-insensitive search check box: Case-sensitive queries.....	45
DB2	46
New (empty) object stores	46
DB2 tuning recommendations and maintenance tasks.....	49
Microsoft SQL Server.....	50
Full-text indexing for content-based retrieval.....	51
FileNet P8 parameters	51
Content Search Engine parameters.....	52
Stop words	52
Application servers	52
WebLogic.....	52
Group membership search.....	52
Operating systems.....	53
Windows 2003	53
Improving performance with “Maximize data throughput for networks applications”.....	53
Fixed content devices.....	53
EMC Centera	53
Recommendations for handling a large number of folders, documents, and objects	53
Folders and subfolders	53
Documents and custom objects.....	54
Recommendations	54
Avoiding index skew	55
Performance diagnosis.....	56
Skew solutions for DB2	56
Skew solutions for Oracle	56

Behavior on SQL Server	57
File storage area paths	58
Content streaming	58
Large content uploads and downloads	58
Inline Content Retrieval Limit and Put/Get Block Size parameters: Uploads and retrievals.....	58
Content upload destination	59
Auditing.....	60
Delete old records from the Event table	60
Object state snapshots.....	60
Improve audit logging and search performance	60
Business Process Management	61
Process Engine	61
Databases.....	61
Oracle – Sequence cache: Preallocated sequence numbers.....	61
Maximum database operations per transaction and maximum instructions per transaction parameter values.....	61
Client tuning – JVMs	62
ORB tuning with Process Engine client connection management.....	62
ORB tuning for IBM JAVA 5.0 JVM.....	63
ORB tuning for Sun JVM.....	65
ORB tuning for JRockit JVM	65
Process Task Manager.....	65
Minimum and maximum heap size for high volume Component Manager activities.....	65
Process Engine authentication memory allocation for workflows	65
VWK process optimization	66
Cached Entry Timeout parameter: User and group cache	67
Flush Granularity option: Cache refresh	67
WorkSpaceCacheSize parameter: Cache isolated region workspaces.....	68
Process Simulator (PS) and Process Analyzer (PA).....	68
Max Simulation Jobs parameter: Concurrent simulation runs	68
Java max memory	68
Tuning Process Analyzer for optimal performance.....	68
Publisher threads	68
Cache size.....	69
Databases.....	69
Publishing interval.....	69
Application Engine	70
Workplace.....	70
Databases.....	70
Folder browsing documents and custom objects	70
Java client cache settings.....	70
Image Viewer	70
eForms	71
Databases	71
High response times when rendering forms	71
Application servers	71
JVM tuning	71
Workplace XT.....	72
Databases	72
Application servers	72
JVM tuning	72
IBM InfoSphere Enterprise Records and IBM FileNet Records Manager	73
Databases	73
Performance enhancing indexes	73
Disposition Sweep.....	73
Disposition Sweep and Auto Destroy.....	74
Disposition Sweep, Hold Sweep, and Auto Destroy	77

Actions properties.....	77
Disposal instruction properties	77
Location properties.....	78
Record type properties.....	78
SupersedeBy property	79
Oracle.....	79
Prevent performance degradation during INSUBFOLDER searches	79
Application servers	79
JVM tuning	79
Thread stack size	79
Set the minimum and maximum heap size for the Dynamic Hold Sweep execution.....	80
Set the minimum and maximum heap size for Destroy Disposition Sweeps.....	80
Improve Sweep and Destroy throughput by tuning write threads.....	80
Content Federation Services.....	81
Federating documents.....	81
Chunk Size parameter: Allowable size of query results	81
Batch Size parameter: Number of documents batched by the exporter.....	81
Importer Threads and Dispatcher Wait Time parameters: Importer throughput.....	81
Maximum Number of (Exporter) Workers parameter: Exporter throughput	82
RMI bridges	82
Image Services.....	83
Cache Service Manager	83
Import agents, Batch Delay parameter, and Document Batch Size parameter.....	83
Import agents	83
Batch Delay parameter	83
Document Batch Size parameter	83
Database indexing	84
Content Manager on Demand	85
Federation throughput.....	85
I_EXTERNALIDENTI76 index.....	85
System logging and trace.....	85
Content Services	86
Simultaneous library requests per user for federation queries and exports.....	86
Cache parameters: Document federation with full-text indexing	86
Image Services.....	87
Cache Service Manager	87
Import agents, Batch Delay parameter, and Document Batch Size parameter.....	87
Import agents	87
Batch Delay parameter	87
Document Batch Size parameter	88
Database indexing	88
IBM FileNet Connectors for SharePoint.....	90
Thread count: Web Parts response times when archiving.....	90
Overlapped recycling for an application pool	90
Batch process and user operations.....	90
Document libraries.....	91
Separate application pool for DocLib Web service.....	91
QueryPageSize setting affects throughput	91
Collector schedules: Resource usage on idle systems with archiving	92
Notices	93
Trademarks	94
U.S. Patents Disclosure.....	95

Revision log

Date	Revision
11/20/09	<p>Added</p> <ul style="list-style-type: none"> • Support for DB2 9.7, as specified in the affected topics • Composite indexes for continuable searches • Reorganize tables and indexes • DB2_WORKLOAD registry variable: Improve performance of Content Manager applications • Indexing date properties in Oracle • Auto-unique naming and containment names • Skew solutions for DB2 • Single property index for records object stores in InfoSphere Enterprise Records <p>Updated</p> <ul style="list-style-type: none"> • Security folder id index • Relationship table index for INFOLDER and INSUBFOLDER searches for DB2 • Large object stores updates in DB2 • DB2_OPTPROFILE profiles • APPLHEAPSZ parameter note for DB2 9.5 and later • DB2_EVALUNCOMMITTED and DB2_SKIPINSERTED registry variables: Reduce deadlocking and improve application concurrency • MAXFILOP database parameter: Number of open file handles per database agent • Records Manager references to InfoSphere Enterprise Records, as appropriate <p>Removed</p> <ul style="list-style-type: none"> • DB2 8 content
07/27/09	<p>Added</p> <ul style="list-style-type: none"> • DB2_OPTPROFILE registry setting • JBoss connection pool maximum capacity • JVM tuning for LDAP connection pooling • Maximum connections for Centera • Content Manager On Demand federation throughput <p>Removed</p> <ul style="list-style-type: none"> • DB2_PINNED_BP
04/06/09	<p>Added</p> <ul style="list-style-type: none"> • Reference to the Lotus Quickr V8 tuning guide • Information about WebSphere DNS lookups • Index information to improve queries for multi-value properties when query has a wildcard value • Persisting multi-value properties in their own database tables • Examples of Content Engine clients • Information about the filenet.properties file • Content Federation Services exporter performance information

Date	Revision
11/24/08	<p>Added</p> <ul style="list-style-type: none"> • FileNet P8 4.5 release • Index on application name in the docversion table • DB2 9 large table spaces • Microsoft SQL Server large content uploads and downloads • Reclustering the docversion table on Microsoft SQL Server • WebSphere native memory errors associated with maximum heap size • Content Search Engine parameters and stop words • Object state snapshots for auditing Events table • Heap values for Component Manager activities • Workflows that have a large number of users or groups • Tuning Image Viewer • Image Services Cache Service Manager (CSM) <p>Updated</p> <ul style="list-style-type: none"> • Syntax of the –Xinitsh switch setting to –Xinitsh4m • DB2 tuning content • Records Manager indexing content
08/19/08	<p>Added</p> <ul style="list-style-type: none"> • Large jumps in CPU usage on Windows 2008 • Improving Content Engine performance on Windows 2003 • Multi-thread Write feature for Records Manager functionality • Content Federation Services-Image Services import agents, batch delay times, and document batch sizes
07/15/08	<p>Added</p> <ul style="list-style-type: none"> • JVM “-Xgc:gencon” tuning • Java heap recommendation • Enabling automatic statistics collection and manually gathering statistics using the DB2 Task Center • Recommendations for handling a large number of folders and objects (Maintained as a technical notice for 3.5) • Full-text indexing for content-based retrieval (CBR) • Indexing to improve the performance of multi-value property searches • Auditing performance • Process Engine ORB tuning and stand-alone Java clients • Sample stack trace error that indicates Process Engine ORB tuning would be beneficial • Process Engine cache refresh • eForms JVM tuning • Records Manager Destroy disposition sweeps • Content Federation Services Importer throughput when used with DB2 • RMI bridges for Content Federation Services

Date	Revision
	<ul style="list-style-type: none"> • Content Services document federation with full-text indexing • Content Federation Services-Image Services • Application pools for SharePoint • SharePoint and the QueryPageSize parameter setting affecting throughput <p>Removed</p> <ul style="list-style-type: none"> • 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.)
04/23/08	<p>Added</p> <ul style="list-style-type: none"> • Content Federation Services-Content Services simultaneous user connections for federation exports and queries • DB2 processing large SQL statements • DB2 reduce deadlock and improve application concurrency • Oracle increase the response time of searches that do not return a large number of records • File storage area paths • JRockit • IBM FileNet Connectors for SharePoint <p>Removed</p> <ul style="list-style-type: none"> • DB2 JDBC fullyMaterializeLOBData parameter • Records Manager Supercededby indexing recommendation
0/25/07	<p>Added</p> <ul style="list-style-type: none"> • DB2 JDBC fullyMaterializeLOBData parameter • Indexing for stored searches • Tivoli database optimization • Tivoli documentation reference • WebSphere 6.1 <ul style="list-style-type: none"> ○ ORB FragmentSize property for Process Engine client requests and replies ○ ORB ConnectionMultiplicity property for client load distribution ○ Client-side LTPA timeout and security settings • JVM tuning information for Process Engine clients <ul style="list-style-type: none"> ○ ORB tuning for IBM JAVA 5.0 ○ ORB tuning for Sun JVM ○ ORB tuning for JRockit JVM ○ ORB tuning with Process Engine Client Connection Management • Oracle Sequence Cache parameter for Process Engine • JVM tuning information for Records Manager • Oracle INSUBFOLDER search information for Records Manager • Content Federation Services section <p>Updated</p> <ul style="list-style-type: none"> • Percentage of silos with VWKs for process optimization • Common application server information

Date	Revision
	Removed Information about the Enumeration Batch Size and Default Query Batch Size parameters
06/29/07	Added <ul style="list-style-type: none">• Records Manager section• Client and server get and put content settings• Oracle connectivity parameters• Content Engine search for Microsoft SQL Server Content upload destination Updated <ul style="list-style-type: none">• WebLogic group membership search• Case insensitive queries
05/29/07	Initial document.

About this document

This document provides tuning tips that can help improve the performance of your IBM® FileNet® P8 system. Tuning tips in this document apply to both versions 4.0 and 4.5 unless otherwise indicated.

Tuning tip organization

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

IBM InfoSphere Enterprise Records and IBM FileNet Records Manager

IBM InfoSphere Enterprise Records is the new name for IBM FileNet Records Manager; you might see both names in the documentation as we transition to the new name.

Parameter settings

Many tips in this document do not include suggested parameter settings. 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 for your version of FileNet P8 and 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* for your version of FileNet P8.
- Installed any other software components, as appropriate.
- Configured the FileNet P8 databases with sufficient processing power, memory, and a set of indexes that are appropriate to the application.

Recommended documentation

The following topics list reference documents that are available from IBM and other associated software vendors.

- *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

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.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/a/e/ae/rprf_hotparameters.html
- WebSphere Information Integrator (IICE). (Applicable to Content Federation Services-Content Services)
http://www.ibm.com/software/data/integration/db2ii/editions_content.html
<http://publib.boulder.ibm.com/infocenter/wsiihelp/v8r3/index.jsp>

WebLogic

- General server documentation
<http://edocs.bea.com/>
- Top tuning recommendations for WebLogic Server 9.2
<http://edocs.bea.com/wls/docs92/perform/topten.html>
- WebLogic Server 8.1.x performance and tuning
<http://edocs.bea.com/wls/docs81/perform/index.html>

JRockit

- Tuning WebLogic JRockit Java™ Virtual Machine (JVM)
http://edocs.bea.com/jrockit/geninfo/diagnos/part_02.html

Database documentation

DB2

- DB2® Information Center
<http://www.ibm.com/support/docview.wss?rs=71&uid=swg27009474>
- 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

SQL Server

- Microsoft® SQL Server 2000 RDBMS Performance Tuning Guide for Data Warehousing
<http://www.microsoft.com/technet/prodtechnol/sql/2000/maintain/rdbmspft.msp>
- 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

Directory servers documentation

Tivoli

- Tivoli® Information Center – Performance Tuning and Capacity Planning Guide
<http://publib.boulder.ibm.com/infocenter/tivihelp/v2r1/index.jsp?toc=/com.ibm.IBMDS.doc/toc.xml>

Collaboration documentation

Lotus Quickr

- Lotus® Quickr™ V8.0 Tuning Guide
<http://www.ibm.com/support/docview.wss?uid=swg27010632&aid=1>

Accessing IBM FileNet documentation

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.

Contacting 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”.

Sending 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. 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).

Operating systems

HP-UX

Java applications support

The default HP-UX configuration, specifically kernel parameters and system settings that can be tuned, do not support certain types of Java applications well. Applications that typically need to be manually tuned are large, server-side Java applications that are long running, have large numbers of threads, and have a large ratio of client-to-server connects.

Java out-of-box is a stand-alone bundle that upon installation installs startup (RC) scripts, modifies kernel parameters, rebuilds the kernel, and restarts the system. During startup, the startup scripts adjust system settings that can be tuned, which provides better out-of-the-box behavior for Java.

For more information, see <http://www.hp.com/products1/unix/java/java2/outofbox/index.html>.

Solaris

Processing time

tcp_time_wait_interval and tcp_conn_req_max_q parameters: TCP connections

When Application Engine 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 meant for 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.

Decreasing the TCP wait time

1. 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.

Increasing the maximum number of queued completed connections

1. 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
```

Windows

Processing time

TcpNumConnections, MaxUserPort, and TcpTimedWaitDelay registry parameters: TCP connections

When Application Engine 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.

Adjusting the maximum number of open TCP connections

1. 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 Content Engine requires high throughput transaction processing.

Adjusting the port number that TCP can assign

1. 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 suggests that you first try adjusting the MaxUserPort and TcpNumConnections values.

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

1. 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

Windows 2008

CPU usage

By default, Windows® 2008 launches the Server Manager Console when the server starts, and refreshes the information displayed in the console every 2 minutes. These refreshes can cause large jumps in CPU usage.

If you notice large jumps in CPU usage at regular intervals, check the Server Manager Console configuration and consider the following modifications:

- Select “Do not show me this console at login” to stop the Server Manager Console from automatically starting when Windows 2008 starts.
- Click *Configure Refresh* to disable or modify the Server Manager Console refresh interval.

Databases

For additional information about application-specific databases, check the appropriate application-specific sections.

All databases

The database indexing tables in this document have information about performance enhancing database indexes that you can create on single or multiple database columns, which will be referred to from this point forward as “single property indexes” and “composite indexes”, respectively.

This information includes:

- **Class** – Class [path] in Enterprise Manager that you follow to access the property that is being indexed
NOTE If a class [path] is not listed, the property is a system property.
- **Database table name** – Location of the property or properties that are being indexed
- **Database column name** – Names of the FileNet P8 properties, as they appear in the database, that make up the index
- **FileNet P8 property name** – Symbolic names of the FileNet P8 properties that make up the index
- Reason to create the index

NOTES

- The “uxx” prefix represents a variable in the property name, and indicates that the property is a custom property that is from either an installed expansion product or user-defined.
- Use FileNet Enterprise Manager 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 index, use the respective database management software.
- In addition to these recommended indexes, you should also use the Microsoft SQL Profiler, the Oracle Statspack, Oracle ADDM, or the DB2 Design Advisor to identify properties (columns) that should be indexed.

Performance enhancing indexes

Consider creating one or more of these database indexes to improve system performance. These indexes are useful to both out-of-the-box applications that are ready to use and custom applications, which include but are not limited to Workplace, Workplace XT, Records Manager, SharePoint Web Parts, and Business Process Framework (BPF).

For more performance enhancing indexes, see:

- [Databases > DB2 > Performance enhancing indexes](#) on page 24.
- [IBM InfoSphere Enterprise Records and IBM FileNet Records Manager > Databases > Performance enhancing indexes](#) on page 73.

Class	Database column name	FileNet P8 property name	Reason to create index
Database table name			
Objectstore > Document Class docversion	object_class_id ----- uxy_searchtype ----- is_current	N/A	Create this composite index when both of the following conditions are true: <ul style="list-style-type: none"> • The number of records in the table exceeds 500,000. • Performance degrades across common functionality (browsing, searching, creating documents, et cetera.)
Objectstore > Document Class > Record docversion	mime_type	MIME Type	Create this single property index when the number of records in the table exceeds 1,000,000, or when both of the following conditions are true: <ul style="list-style-type: none"> • The number of items in a folder exceeds 2,000. • Performance degrades across common functionality (browsing, searching, creating documents, et cetera.)
Objectstore > Document Class > Access Role generic	uxx_applicationname	Application Name	Create this single property index when both of the following conditions are true: <ul style="list-style-type: none"> • The number of records in the table exceeds 500,000. • Performance degrades across common functionality (browsing, searching, creating documents, et cetera.) <p>NOTE 4.5 installations create an index on Application Name to improve Application Engine login performance and to improve retrieval of access roles. Before creating this index, verify that one has not already been created.</p>
Objectstore > Document Class > Preferences Doc and Objectstore > Document Class > Stored Search docversion			

Class	Database column name	FileNet P8 property name	Reason to create index
N/A	security_folder_id	N/A	<p>Before Content Engine 4.0.1</p> <p>Create this single property index if it does not exist in the database.</p> <p>Content Engine 4.0.1 and later</p> <p>Verify the index was created when the object store was created. If this single property index does not exist in the database, create it.</p>
docversion			
N/A			
generic			
N/A	security_policy_id	N/A	<p>Create a single property index in all of the listed locations when response time degrades or errors occur when deleting a security policy on an object store that has more than 1,000,000 documents.</p>
docversion			
N/A			
container			
N/A			
generic			

Composite indexes for continuable searches

To achieve the best performance during continuable (paged) searches on single-valued properties:

- Create a composite index on any selective column that is associated with the single-valued property that is searched on and object_id.

For example, if a continuable search on *myProperty* can use an index on its associated database column *uxy_myProperty*, that is the selectivity is good, then create the following composite index: (*uxy_myProperty*, *object_id*).

The new index must be created by using database tools because Enterprise Manager does not have the functionality to create composite indexes.

- Add an "ORDER BY *myProperty*" clause to the search.

By making these adjustments, the database SQL select statement that the Content Engine issues, which uses an "ORDER BY *uxy_myProperty*, *object_id*" clause can better achieve a database query plan that runs off the composite index and results in faster paging results.

Indexing object value properties

Consider creating single value indexes when you use object value properties (OVPs) that are associated with a large table, such as the docversion table.

Generally, if the association properties values are mostly null, there is little overhead in index maintenance, and little affect on document ingestion.

Class	Database column name	FileNet P8 property name	Reason to create index
Database table name			
class [path]	uxy_property name	Property Name	When object value properties are associated with a large table, they can benefit from indexes.
table name			

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

Reorganize indexes and tables

As large amounts of content is added, removed, or changed in a database, the associated tables and indexes need to be periodically reorganized for optimal performance.

See the database vendor documentation to determine if a REORG is necessary, and if so, how to complete the process.

DB2

Performance enhancing indexes

Consider creating one or more of these database indexes to improve system performance.

For more performance enhancing indexes, see:

- [Databases > Performance enhancing indexes](#) on page 21.
- [IBM InfoSphere Enterprise Records and IBM FileNet Records Manager > Databases > Performance enhancing indexes](#) on page 73.

Class	Database column name	FileNet P8 property name	Reason to create index
Database table name			
N/A	tail_id	N/A	<p>Creation of this composite index can improve the performance of INSUBFOLDER and INFOLDER searches.</p> <p>For example, when Workplace and Workplace XT use INSUBFOLDER and INFOLDER to retrieve stored searches or properties of stored searches.</p> <p>Records Manager also uses INSUBFOLDER and INFOLDER to retrieve stored searches.</p>
Relationship	----- head_id		

Large object stores

When you create object stores in DB2 9.1 or later, use the default setting of Large to create table spaces that use the new larger RID table space format.

When you upgrade existing object stores from DB2 version 8 to version 9.1 or later, consider converting your table spaces. For detailed instructions about this procedure, see:

http://publib.boulder.ibm.com/infocenter/db2luw/v9/index.jsp?topic=/com.ibm.db2.udb.apdv.sample.doc/doc/admin_scripts/s-largerid-db2.htm

DB2_WORKLOAD registry variable: Improve performance of Content Manager applications

(Content Engine 4.5.1 and later) DB2_WORKLOAD is a DB2 registry variable that represents pre-defined groups of (other) registry variables that have been specifically compiled to optimize performance of certain systems.

As of DB2 9.7 (and DB2 9.5 fix pack 5), "FILENET_CM" is a valid setting for the DB2_WORKLOAD registry value, and represents optimally set registry variables that have been compiled to increase the performance of applications that utilize the FileNet Content Manager.

IMPORTANT It is strongly recommended that you enable the DB2_WORKLOAD registry parameter with the FILENET_CM value on Content Engine 4.5.1 and later.

Enabling the DB2_WORKLOAD registry variables

Enable the DB2_WORKLOAD registry variable by entering the following command in the DB2 command-line processor:

```
db2set DB2_WORKLOAD=FILENET_CM
```

DB2_OPTPROFILE

All supported versions of DB2 9 before DB2 9.5 fix pack 5

Enable the DB2_OPTPROFILE registry variable by entering the following commands in the DB2 command-line processor:

```
db2set DB2_OPTPROFILE=ON
```

DB2 9.5 fix pack 5 and later DB2 9.7 and later

Enable the DB2_OPTPROFILE registry variable by entering the following commands in the DB2 command-line processor:

```
db2set DB2_WORKLOAD=FILENET_CM
```

For information about the DB2_WORKLOAD registry variable, see [DB2_WORKLOAD registry variable: Improve performance of Content Manager applications](#) on page 24.

DB2_EVALUNCOMMITTED and DB2_SKIPINSERTED registry variables: Reduce deadlocking and improve application concurrency

To improve concurrent query and ingestion workloads, consider enabling the DB2_EVALUNCOMMITTED and DB2_SKIPINSERTED registry variables. When these registry values are enabled, unnecessary locking is avoided, which can improve performance and, in some cases, prevent deadlock or query time-out errors.

IMPORTANT

- (Content Engine before 4.5.1) Enabling these parameters in pre-4.5.1 versions of Content Engine that use fixed content devices can, in rare case, cause data loss.
- If you use fixed content devices and system performance requires you to enable these parameters, you must upgrade to Content Engine 4.5.1
- These parameters are enabled at the instance level and apply to all databases that belong to the instance. However, these changes do not go into effect until the instance is restarted.

Enabling the DB2_EVALUNCOMMITTED and DB2_SKIPINSERTED registry variables

All supported versions of DB2 9 before DB2 9.5 fix pack 5

Enable the DB2_EVALUNCOMMITTED and DB2_SKIPINSERTED registry variables by entering the following commands in the DB2 command-line processor:

```
db2set DB2_EVALUNCOMMITTED=ON
```

```
db2set DB2_SKIPINSERTED=ON
```

**DB2 9.5 fix pack 5 and later
DB2 9.7 and later**

Enable the DB2_EVALUNCOMMITTED and DB2_SKIPINSERTED registry variables by entering the following command in the DB2 command-line processor:

```
db2set DB2_WORKLOAD=FILENET_CM
```

For information about the DB2_WORKLOAD registry variable, see [DB2_WORKLOAD registry variable: Improve performance of Content Manager applications](#) on page 24.

APPLHEAPSZ parameter: Processing large SQL statements

The APPLHEAPSZ parameter defines the number of private memory pages that can be used by the database manager on behalf of a specific agent or subagent. The minimum amount of heap is allocated when an agent or subagent is initialized for an application. As the agent or subagent requires more heap space to process larger SQL statements, the database manager allocates memory as needed, up to the maximum number of pages specified by this parameter.

Increase the APPLHEAPSZ parameter value if your applications receive an error indicating that there is not enough storage in the application heap.

Adjusting the APPLHEAPSZ parameter

The acceptable range of values for the APPLHEAPSZ parameter is 16 – 60,000. The default value is 128 (4 KB pages).

To find the current value for the APPLHEAPSZ parameter, issue the following command or view the Configure Database window of the DB2 Control Center:

```
GET DATABASE CONFIGURATION
```

To update the value of the APPLHEAPSZ parameter, issue the following command where “nnnn” is the parameter value. Begin with 2048 as the parameter value, and adjust it as necessary if more or less heap space is required.

```
UPDATE DATABASE CONFIGURATION FOR db_name USING APPLHEAPSZ nnnn
```

(DB2 version 9.5 and later) The default value for the APPLHEAPSZ parameter is “Automatic”. When this parameter is set to “Automatic”, the application heap grows as needed, up to the appl_memory limit for the database providing that the total amount of memory allocated by the database partition does not exceed the instance_memory limit for the database partition. The entire application heap is freed when the application disconnects from the database.

MAXFILOP database parameter: Number of open file handles per database agent

The MAXFILOP database parameter specifies the maximum number of open file handles that a database agent can manage. If the number of open file handles exceeds the value of this setting, other file handles that are in use by the database handler are closed in order to remain within the limits of the MAXFILOP database parameter setting.

System performance can be adversely affected if file handles are repeatedly opened and then closed to accommodate the limits MAXFILOP database parameter setting.

You can determine if repeatedly opened and closed file handles is affecting system performance by monitoring the Database Files Closed counter that is included in the Get Snapshot for Database command output.

If the Database Files Closed counter is greater than 0, consider increasing the MAXFILOP database parameter. Start with a value of 128, and adjust the value according until Database Files Closed counter is equal to 0.

Oracle

Indexing date properties

For information about indexing date properties on Oracle, see the IBM FileNet “Enabling Oracle Date Index Use in the FileNet Content Engine” technote:

http://www-01.ibm.com/support/docview.wss?rs=3278&context=SSNVNV&context=SSTHRT&context=SSNVUD&context=SSS236&context=SSNW2F&dc=DB560&dc=DB520&uid=swg21397282&loc=en_US&cs=utf-8&lang=en

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:

Version 10g

http://download.oracle.com/docs/cd/B19306_01/network.102/b14212/listenercfg.htm#NETAG0101

Version 11g

http://download.oracle.com/docs/cd/B28359_01/network.111/b28316/listenercfg.htm#CHDEBAJE

Connectivity parameters

Increasing the parameters that are documented in the following table might prevent database connectivity issues when your system supports many simultaneous connections to Oracle.

The settings that you choose need to enable the Oracle database instance to sustain the total number of possible simultaneous connections for all object stores in the database instance.

For more information, see [“Connection pools”](#) on page 32.

Parameter name	Description
db_file_multiblock_read_count	Minimizes disk I/O during table scans by specifying the maximum number of blocks read in one I/O operation during a sequential scan.
job_queue_processes	Specifies the maximum number of processes that can be created for the execution of jobs.
parallel_max_servers	Specifies the maximum number of parallel execution processes and parallel recovery processes for an instance.
processes	Specifies the maximum number of operating system user processes that can simultaneously connect to an Oracle server. This value should allow for all background processes such as Job Queue (SNP) and parallel execution (Pnnn) processes.
sessions	Specifies the total number of user and system sessions. NOTE The default number is greater than PROCESSES to allow for recursive sessions. The sessions default is derived (1.1 * PROCESSES + 5).

Disable hash join to increase the response time of 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 effect other searches or API calls that return very large rowsets (such as GetContentees() on a folder with a very large number of objects) and that rely on a hash join in the database query plan.
- Enabling the hash join parameter has a positive effect on complex queries and queries that retrieve large rowsets. However, the hash join parameter is less effective when queries return smaller rowsets and when the response time of the initial batch is of the highest importance.

Disable hash join

1. Set the hash_join_enabled parameter to “false” in the initORACLE_SID.ora file.
 - a. Open the initORACLE_SID.ora file for editing.

The file is typically found here:

Windows

ORACLE_HOME\DATABASE\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.

2. Create the pfile by issuing the following command:

Windows

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

UNIX

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

3. Add the line “hash_join_enabled=false” to the generated init.ora file.
4. Start the database by issuing the following command:

Windows

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

UNIX

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

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

Microsoft SQL Server

Large content uploads and downloads

Uploading and downloading documents that are 300MB or larger can cause errors in Content Engine, Workplace, and other content-based application written to the Content Engine APIs.

To prevent or fix this problem, use the Microsoft SQL 2005 JDBC Driver v1.2 and add the following custom properties to the data source configuration:

Object Store DataSource non-XA:

responseBuffering=adaptive packetSize=32767

Object Store DataSource XA:

responseBuffering=adaptive packetSize=32767

WebSphere

1. From the WebSphere Administrative Console, select **Console > JDBC > Data Sources**, and then click on an object store data source.
2. Select **Custom Properties**.
3. From the Custom Properties page, click **New** to add the following two properties:
 - Name: responseBuffering, Value: adaptive;
 - Name: packetSize, Value: 32767

WebLogic

1. From the WebLogic Administrative Console, click on an object store data source.
2. Select **Connection Pool**.
3. Add the following two properties in the Properties field:
 - responseBuffering=adaptive
 - packetSize=32767

JBoss

1. Edit an object store data source XML file (your_datasource-ds.xml and your_datasource-xa-ds.xml) located in `./server/yourserver/deploy`.
2. Add the following two lines to the XML file:

XA

```
<responseBuffering>adaptive</responseBuffering>  
<packetSize>32767</packetSize>
```

Non-XA

```
<responseBuffering>adaptive</responseBuffering>  
<packetSize>32767</packetSize>
```

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
- Records Manager continuable searches
- Java API 4.x 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 Content Engine as stated below.

Versions affected

The docversion table is clustered on object_id when new object stores are created on:

- Content Engine 3.5.2-014 and later
- Content Engine 4.0.0-002 and later

Therefore, reclustering is only necessary for:

- Content Engine 3.x object stores created before 3.5.2-014
- Content Engine 4.x object stores created before 4.0.0-002
- Content Engine 4.x object stores upgraded from 3.x, when the 3.x object store was created prior to 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

Application servers

For details about the minimum supported levels, see the IBM FileNet P8 Hardware and Software Requirements Guide. To download this guide from the IBM support page, see [Access IBM FileNet Documentation](#) on page 16.

For additional information about application-specific application servers, check the appropriate application-specific sections.

WebSphere, WebLogic, and JBoss

Connection pools

WebSphere, WebLogic, and JBoss all use connection pooling in their data source configurations to minimize resource usage when connecting to the database.

If an error occurs during high-volume conditions, consider increasing the maximum capacity of the connection pools.

Use the following formulas as a general rule of thumb to determine the connection thread pool capacity.

- For Global Configuration Database (GCD) connection pools, both distributed (XA) and non-distributed (non-XA):
 - Maximum connections = $10 * (\# \text{ object stores}) + (\# \text{ execution threads}) + 10$
 - Minimum connections = $20 + (\# \text{ execution threads})$
- For each individual object store connection pools:
 - Maximum and minimum connections = $20 + (\# \text{ execution threads})$

NOTE Execution threads are the maximum number of expected concurrent client connections. In the GCD data source, execution threads are the maximum number of expected concurrent client connections for all object stores in a domain. In the object store data source, they are the maximum number of expected concurrent client connections for that object store.

Adjusting the connection pool maximum capacity setting

WebSphere 6.0.x

1. Navigate to Resources > JDBC Providers.
2. Select Connection Pool Name > Data Sources.
3. Select Data Source Name > Connection Pool Properties.
4. Adjust the Maximum Capacity setting.

WebSphere 6.1.x

1. Navigate to Resources > JDBC.
2. Adjust the Maximum Capacity setting.

WebLogic 8.1.x

1. Navigate to JDBC Connection Pool > Configuration > Connections.
2. Adjust the Maximum Capacity setting.

WebLogic 9.x

1. Navigate to Services > JDBC > Data Sources.
2. Select Data Source Name.
3. Select Connection Pool.
4. Adjust the Maximum Capacity setting.

JBoss 4.0.x

Edit the XA and non-XA datasource definitions under *JBOSS_DIR*\server\CE_SERVER\deploy:

DATASOURCE_NAME-ds.xml

DATASOURCE_NAME-xa-ds.xml

Add the following entries:

```
<min-pool-size>x</min-pool-size>
```

```
<max-pool-size>x</max-pool-size>
```

Thread pools

If you encounter errors when the application server handles a high volume of requests, or you encounter errors that are related to insufficient thread creation, consider increasing the minimum and maximum thread pool settings.

NOTE Increasing the minimum thread pool size enables the system to pre-allocate resources when the system is started.

Adjusting the minimum and maximum thread pool settings

WebSphere 6.0.x and 6.1

1. Navigate to Servers > Application Servers > ServerName > Additional Properties > Thread Pools > WebContainers.
2. Adjust either the Minimum setting or the Maximum setting, or adjust both.

WebLogic 8.1.x

1. Navigate to DomainName > Servers.
2. Right click ServerName and select View Execute Queues.
3. Select "weblogic.kernal.Default", the default execute queue (thread pool).
4. Adjust either the Thread Minimum setting or the Thread Maximum setting, or adjust both.

WebLogic 9.x

WebLogic 9.x has a single execute queue (thread pool) that is automatically tuned to achieve the highest level of performance. Thread count is based upon administrator-configured application-level parameters like fair-share and response time goals.

JBoss 4.0.x

Adjust the maximumpoolsize value in the jboss-service.xml file, which is located in the C:\JBoss405\server\your server\conf directory.

WebSphere

This section lists performance-tuning suggestions that apply only when FileNet P8 is running on a WebSphere application server.

Lightweight third party authentication (LTPA) timeout: Token validity

The LTPA “Timeout” setting specifies how long an LTPA token is valid (in minutes). The token contains this expiration time so that any server that receives the token can verify that the token is valid before proceeding further. When the defined amount of time has passed, all user tokens expire regardless of session activity.

The IBM Performance Lab suggests increasing this value if one or both of the following conditions occur:

- Batch processes or other operations run uninterrupted for longer than the currently configured LTPA timeout value.
- The following message is in the server.log file:
“org.omg.CORBA.NO_PERMISSION: Credentials have expired”

Adjusting the LTPA timeout setting – WebSphere 6.0.x

1. Navigate to Servers > Global Security.
2. Under Authentication, select Authentication Mechanisms > LTPA.

Adjusting the LTPA timeout setting – WebSphere 6.1.x

1. Navigate to Secure administration, applications, and infrastructure > Authentication and Authorization Service.
2. Under Authentication, select Authentication mechanisms and expiration.
3. Enter the security cache timeout value in the “Timeout value for forwarded credentials between servers” field.

Client-side security cache timeout: Cache refresh

The client-side security cache timeout setting specifies how often security-related caches are refreshed (in minutes). When the defined amount of time has passed, all cached information becomes invalid. Subsequent requests for the information result in a database lookup.

The default value is 600. Increasing this value to 3600 provides a modest improvement in server response times. (Increasing the value further will not provide additional gains.) If your environment does not require a security cache refresh every 600 seconds, consider increasing this value.

Adjusting the security cache timeout setting – WebSphere 6.0.x

1. Navigate to Servers > Global Security.
2. Click the Configuration tab.
3. Enter the security cache timeout value in the “Cache timeout” field.

Adjusting the security cache timeout setting – WebSphere 6.1.x

1. Navigate to Secure administration, applications, and infrastructure > Authentication and Authorization Service.
2. Under Authentication, click Authentication mechanisms and expiration.
3. Enter the security cache timeout value in the “Authentication cache timeout” field.

Performance monitoring infrastructure PMI

By default, WebSphere Performance Monitoring Infrastructure is enabled. IBM Performance Lab results show that disabling this option provides moderate response time and CPU utilization improvements. If your system is healthy and you are not troubleshooting performance or scalability issues, consider disabling this option.

You can disable PMI by navigating to **Monitoring and Tuning > Performance Monitoring Infrastructure (PMI)**. Select your server, and then deselect the **Enable PMI** option.

Service logs

Service logs consist of that are messages written to the System.out stream, as well as some special messages that contain extended service information that is normally not of interest, but can be important when analyzing problems.

On occasion these logs can cause excessive disk I/O even when your WebSphere system is performing as expected. If this I/O affects other applications that are co-located with your application server, consider moving them to a separate set of physical disks or disabling them entirely.

You can disable these logs by navigating to **Troubleshooting > Logs and Trace > ServerName > IBM Service Logs**, then deselect the “Enable service log” option.

Local host lookups

The following JVM property improves the performance of host name look-ups by caching the host name that is returned from the underlying operating system:

```
com.ibm.cacheLocalHost=true
```

The JVM always uses the cached host name, so this tuning tip might not be appropriate for situations where the local host name can change for a Java application.

Configuring the com.ibm.cacheLocalHost property

1. In the WebSphere Administrative Console, select **Application Servers > ServerName > Process Definition > Control > JVM Properties > Custom Properties > New**.
2. Enter the following information:
 - Name: com.ibm.cacheLocalHost
 - Value: true

WebLogic

This section lists performance-tuning suggestions that apply only when FileNet P8 is running on a WebLogic application server.

Server and client JVM switches

Sun Microsystems has a version of 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, which are suited to the distinctly different performance characteristics of clients and servers.

The server-side JVM has shown significant performance improvements over client-side JVM in IBM's highly scalable server environment. If your application server is running in production mode, the server compiler is already being used.

Linux socket leak

FileNet P8 applications might encounter out of memory errors when running WebLogic 9.x on Linux® RedHat. This is due to a WebLogic issue, which causes a socket leak.

To resolve this issue, add the following line to your WebLogic startup script:

```
-Dweblogic.MuxerClass=weblogic.socket.PosixSocketMuxer
```

JVM tuning

For additional information about application-specific JVM tuning, check the appropriate application-specific sections.

JVM tuning for LDAP connection pooling

By default, secure sockets layer (SSL) connections to lightweight directory access protocol (LDAP) servers are not pooled in the application server JVM, which can adversely affect Workplace logon response times when Content Engine 4.x is configured for SSL.

To improve Workplace logon response times when Content Engine 4.x is configured for SSL, enable LDAP connection pooling on the application server.

JBoss

Use the following procedure to enable LDAP connection pooling on a JBoss application server.

1. Stop JBoss.
2. Open run.bat (Windows) or run.sh (UNIX or Linux) from *JBoss_HOME/bin* folder.
3. Add the following system properties to the JAVA_OPTS variable:
 - Dcom.sun.jndi.ldap.connect.pool.initsize=20
 - Dcom.sun.jndi.ldap.connect.pool.maxsize=20
 - Dcom.sun.jndi.ldap.connect.pool.prefsiz=10
 - Dcom.sun.jndi.ldap.connect.pool.timeout=300000
 - Dcom.sun.jndi.ldap.connect.pool.protocol="plain ssl"

UNIX and Linux

Use escape characters for double quotes and put the JAVA_OPTS variable in double quotes when launching Java.

(For the bash shell, use a backslash (\) to escape double quotes.)

4. Tune the connection pool size according to the site requirements.
5. Start JBoss.

Other application servers

Refer to the application server documentation for information about changing JVM arguments.

Tuning for improved response time and CPU utilization

To improve overall throughput, the IBM Performance Lab suggests using a generational garbage collection strategy, which is also known as “generational concurrent garbage collection”.

The generational garbage collection strategy works best with workloads that seek to reduce or eliminate garbage collection pauses, and are willing to trade some memory throughput to accomplish the reduction or elimination. The applications that support this kind of workload can typically employ gigabyte-sized heaps. Following are some characteristics of workloads that could benefit from generational garbage collection:

- Allocates many short-lived objects
- Is transaction-based, where objects in the transaction are deleted after the transaction is committed
- Tend to develop a fragmented heap

NOTE To verify the effectiveness of the garbage collection strategy, run the application, and measure either response times or the throughput relative to garbage collection pause times, or both.

JVM switch	Description
WebSphere format -Xgcpolicy:gcType For example: -Xgcpolicy:gencon	Defines a specific garbage collector algorithm. The IBM Performance Lab suggests setting the garbage collector algorithm to “gencon”, a generational, concurrent garbage collector. With this type of garbage collector, objects are allocated in the young generation (the nursery). When the young generation is full, JVM stops all Java threads and moves the live objects in the young generation to the old generation.
JRockit format -Xgc:gcType For example: -Xgc:gencon	
Sun HotSpot Generational garbage collection is the default setting.	

High-volume ingestion scenarios

During high-volume ingestion scenarios, Content Engine creates a large amount of short-lived objects. The IBM Performance Lab suggests that you provide ample space to your new or tenure generation. Monitor your JVM by using VerboseGC or WebSphere Performance Management Infrastructure and adjust the size of these generations appropriately.

JRockit tuning for improved response time and CPU utilization

To improve response time and CPU utilization with WebLogic 9.2 MP2, the IBM Performance Lab suggests tuning the following JVM switches.

NOTES

- Start with the values provided in the examples and adjust them accordingly.
- See the BEA JRockit documentation for more information about these switches.

JVM switch	Description
-XXtlaSize	<p>Defines the size of thread local areas (TLAs), which are used for object allocation. The default value is 2KB.</p> <p>Format: -XXtlaSize:size[k K] [m M] [g G]</p> <p>For example: java -XXtlaSize:64k</p>
-XXlargeObjectLimit	<p>Defines the size of a large object. Objects that exceed the defined size limit are considered large and are not allocated in TLAs.</p> <p>Format: -XXlargeObjectLimit:size[k K] [m M] [g G]</p> <p>For Example: java -XXlargeObjectLimit:64k</p>
-Xns	<p>Defines the nursery size. The IBM Performance Lab suggests setting your JRockit nursery size to 25% of the total heap.</p> <p>Format -Xns:size[k K] [m M] [g G]</p> <p>For Example: java -Xns:300m</p>

Heap size

The JVM heap is a repository of active objects, inactive objects, and free memory. When the JVM heap runs out of memory, all processing in the JVM stops until garbage collection completes.

The JVM heap size is important because it controls how often, and for how long garbage collection runs. If garbage collection runs too infrequently, the extra time that is required to complete garbage collection can negatively affect FileNet P8 performance.

When you tune the JVM heap size, consider the following two objectives:

- Minimizing the amount of time that is required to complete garbage collection
- Maximizing FileNet P8 functionality to handle a significant workload

By balancing the JVM heap with the number of instances in the application server cluster, you can minimize the number of lengthy garbage collection pauses, and the frequency of full garbage collections that are caused by insufficient heap space.

The JVM heap size is a configurable parameter that you can set when you start the JVM. The lowest acceptable value for the minimum heap size is 512MB, and the lowest acceptable value for the maximum heap size is 1GB.

Native out-of-memory errors

Out-of-memory errors can occur when the JVM does not have an adequate amount of native memory space.

Confirm enough native memory is available to the JVM, which includes memory for the maximum heap size, and additional memory required by the JVM, such as the Just-In-Time compiler, and other native code. To be safe, the amount of native memory allocated for the JVM process should be no more than 60% of the per process memory limit as dictated by the operating system and application server platform.

The required amount of additional memory required by the JVM process varies for each environment. However, 25 percent of the maximum heap size is a reasonable guideline. For example, if the maximum heap size is set to 1 GB, allow the native memory for the JVM process to at least 1.25 GB.

By applying these rules, an example of a JVM running on 32-bit Windows with 2GB of memory available to any single process might use the following settings as the smallest minimum and largest maximum allowed:

`-Xms=512m`

`-Xmx=1228m`

JVM out-of-memory errors

IBM

Under certain conditions, the IBM Performance Lab observed that a FileNet P8 application server would encounter out-of-memory errors when the heap still had ample memory available. For example:

```
E_UNEXPECTED_EXCEPTION: An unexpected exception was detected. Caused by:  
java.lang.OutOfMemoryError: JVMCI015:OutOfMemoryError, cannot create more threads due to  
memory or resource constraints.
```

Situations similar this example can occur when the memory inside the heap becomes fragmented or is otherwise poorly allocated.

One way to avoid similar situations is to increase the minimum heap size to 50% of the maximum heap size assuring sufficient memory is allocated when the JVM is started. Another option is to ensure that an adequate amount of memory is set aside in the heap to load the JVM code. This area of memory can be called the class cache. The amount of memory for the class cache is controlled by the Xinitsh JVM parameter, which sets the initial system size of the heap.

For JVMs that run on 32-bit operating systems, consider a value of 4m. JVMs that run on 64-bit operating systems are set to 8m by default, which is typically large enough to avoid this problem. To use this option, add the following switch to the application server command line arguments: `-Xinitsh<size>`.

For example, `-Xinitsh4m`.

JRockit

Under certain conditions, the IBM Performance Lab observed that a FileNet P8 application server would throw out-of-memory errors when the heap still had ample memory available. An example of these errors is the following message: "Error starting thread: Not enough storage is available to process this command."

The cause of this message can be either improper sizing of the JRockit heap generations or memory fragmentation caused by system class placement in the heap, or both. Unlike the Sun Hotspot and IBM JVMs, JRockit does not have permanent system class heap space, which can cause fragmentation problems in certain scenarios.

To avoid this issue, the IBM Performance Lab suggests setting your JRockit nursery size to 25% of the total heap. Depending upon the workload, this setting might require some additional adjustments. You can monitor the nursery size by using VerboseGC collection and analysis.

Directory servers

Tivoli

Optimize the database

To improve performance and query speed, optimize the database by using the Optimize Database task in the IBM Tivoli Directory Server Configuration Tool. This tool uses the DB2 runstats command and updates statistical information that is used by the query optimizer for all the LDAP tables.

NOTE The IBM Performance Lab highly suggests that you optimize the directory server database periodically or after significant database updates.

Optimizing the database by using the Configuration Tool

1. Stop the directory server.
2. Start the Configuration Tool by typing `idsxcfg`
3. on the command line.
4. Click **Optimize Database** in the left navigation pane.
5. In the Optimize Database window, click **Optimize**.
6. When the optimization process finishes, exit the Configuration Tool to stop `idsxcfg`.
7. Restart the directory server.

.Net API client applications

Recycle Worker Process parameter (IIS 6.0 for .Net API)

Enabling the Recycle Worker Process setting for an application pool triggers automatic recycling of worker processes and can affect system performance.

The recycling interval is determined by taking the default recycling value and dividing it by the number of worker processes. For example, if you have 4 worker processes, using the default recycling value of 29 hours, the recycling interval is every 7.25 hours.

If you have batch processes or other user operations that require uninterrupted connectivity, such as the IBM FileNet SharePoint Document Library Archive task route, the IBM Performance Lab suggests that you set the recycling value to a higher value to prevent interruption of service.

Adjusting automatic worker process recycling

1. Open Internet Information Services Manager by clicking Start > Programs > Administrative Tools > Internet Information Services.
2. Expand Local Computer, and then expand the Application Pools folder.
3. Right-click on the application pool and select Properties.
4. Select the "Recycle worker process" check box, and enter a new (in minutes), or select the "Recycle worker processes at the following times" check box, and add an off-hours time.
5. Select the Health tab, and then adjust the "Startup time limit" and the "Shutdown time limit" values, if necessary.
6. Click Apply and then OK to exit the properties dialog.
7. Repeat these steps for all available application pools.

For more information, see:

<http://www.asp.net/faq/AspNetAndIIS6.aspx>

and

<http://www.microsoft.com/technet/prodtechnol/WindowsServer2003/Library/IIS/1652e79e-21f9-4e89-bc4b-c13f894a0cfe.mspx?mfr=true>.

Content Engine

Databases

For overview information, see [“Database – All databases”](#) on page 20.

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 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 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

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 45.

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

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 45.

NOTES

- Each database type has its own limit on index byte size, which can require extra work to create the index for string data types. By default, the ListOfString's element_value column is of character size 4000 (1333 on Oracle), 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 45.

Persisting multi-value properties in their own database tables

(Content Engine 4.5 only) You can improve the performance of multi-value property queries by assigning the Own Table persistence type to property templates.

Persistence types are assigned to property templates individually, and the Own Table persistence type isolates a multi-value property into its own database table so that only those values associated with the multi-value property are stored there.

Content Engine 4.0 and earlier stores multi-value property values in a set of single tables defined by data type, such as String, Float, and so on. You can still persist multi-value properties by data type by selecting System Table instead of Own Table when you configure a property template.

Setting a multi-value property to persist

1. Run the “Configure a property template” wizard.
2. Select the *More* tab.
3. Select the appropriate persistence type.

IMPORTANT Once you select a persistence type, it cannot be changed. For more information, see *System administration > Content Engine administration > Properties > View and modify property templates > More tab in the IBM FileNet P8 Help system.*

Oracle and DB2

Force case-insensitive search check box: Case-sensitive queries

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

Making ad hoc queries case insensitive

1. Log on to Enterprise Manager 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.

Maximizing 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 suggests that you replace these indexes with indexes that use the LOWER function to maximize query performance. (Indexes that use the LOWER function are function-based indexes in Oracle and generated columns in DB2.)

For more information about function-based indexes or generated columns, see the appropriate database 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 by 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, Content Engine 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 by 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 Content Engine server:
`User > object store name > Database`
- To view counts and durations for server operations select the RPC node for a given machine and Content Engine server.

For more information about FileNet System Manager, See “FileNet System Manager” overview in the FileNet P8 documentation.

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 In scenarios where an object store was just beginning to be populated, the IBM Performance Lab observed improvements when manually gathering statistics using the DB2 Task Center within the first 1,000,000 documents or the first 100,000 folders.

Manual statistics collection by 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 script 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>

Setting 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 always uses the indexes that are 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 approach works only in the cases where the index contains all the columns that are referenced or when the index is able to apply a predicate in the index scan.

Content Engine has defined some tables to be used as queues. Because the number rows in these tables are in constant change, setting the table VOLATILE is preferred over relying on current statistics. These tables are set to VOLATILE automatically during object store creation as of certain releases. Object stores that are created before these releases should be altered to have the VOLATILE setting on these tables.

The following tables are used as queues:

VOLATILE as of P8 Content Engine 4.0.0

CONTENTQUEUE

VOLATILE as of P8CE-4.0.0-002:

QUEUEITEM

PUBLISHREQUEST

INDEXREQUEST

WORKQUEUE

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

Defining 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
```

Removing 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
```

Listing 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 Content Engine performance when Content Engine 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 NOTE DB2 version 9.x: Set this value to automatic.	Setting this parameter value can help avoid Content Engine installation errors. For more general information, see " APPLHEAPSZ parameter: Processing large SQL statements " on page 26.
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_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. NOTE 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.

Parameter	Suggested starting value	Notes
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.
LOGFILSIZ	6000	Specifies the size of both primary and secondary log files. NOTE 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. NOTE 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: <ul style="list-style-type: none"> • CONTENT • CONTAINER • DOCVERSION • PROPERTYDEFINITION • RELATIONSHIP • GENERIC • LINK • SECURITYDESC
SORTHEAP	2048 NOTE 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 NOTE DB2 version 9.5: Set this value to automatic.	Specifies the maximum amount of memory that can be allocated as a work space for the SQL compiler during compilation of an SQL statement.

Microsoft SQL Server

Large content uploads and downloads

(Microsoft SQL Server 2005 only) For information, see "[Large content uploads and downloads](#)" under the Databases - Microsoft SQL Server heading.

Indexing for Content Engine Search

(For objects stores created before Content Engine 4.0.1 only) Changing the clustered index on the SecurityDesc table from security_id to hash or removing the clustered index entirely could have a beneficial impact on your system.

Full-text indexing for content-based retrieval

FileNet P8 parameters

Adjusting one or more of the following Enterprise Manager parameters for Content Search Engine can help improve content-based retrieval (CBR) performance.

FileNet P8 domain parameter name	Description
Thread count	Specifies the number of background threads to process index requests. Increase this value on machines that have multiple CPUs, or on machines with configurations that distribute processing across multiple CSE servers. Typically, this is 2 threads per CSE server CPUs.
Optimization Interval	Specifies the number of document, annotation, folder (and so on) batches to index before optimizing on the index information.
Max Batch Size	Specifies the maximum number of objects to index for submitting to Content Search Engine in a single batch. The default value is 8000. The IBM Performance Lab suggests a setting of 4000 for smaller-sized documents and the English locale, and 1000 for larger-sized documents and the Uni locale.
Max Objects per Collection	Specifies the approximate number of objects to be placed in any Content Search Engine collection. The IBM Performance Lab suggests that you do not change this value from the default.
Lease Duration	Specifies the number of seconds that are required to complete an index request. If an index request takes longer than the lease time, the index request is considered abandoned and is re-executed. The IBM Performance Lab suggests that you do not change this value from the default.
Dispatcher Wait Interval	Specifies the number of seconds that the dispatcher waits before looking for work. Valid settings range from 5 to 600 seconds. If the value of this parameter is set too low, some content might not get migrated to the storage area by the time the Content Search Engine tries to index it.

Adjusting Content Search Engine parameters

1. Navigate to and open Enterprise Manager.
2. Right-click the Enterprise Manager domain root.
3. Select Properties, and then click on the Verity Server tab.

Content Search Engine parameters

The Asynchronous Thread field specifies the number of threads that are available for concurrent processing-intensive requests.

The IBM Performance Lab suggests that this value is two times the number of CPUs that will be used by the Content Search Engine server. When more than the allocated number of requests is made at one time, the excess are queued.

For information about how to tune this parameter, see the K2 Server Administration Guide.

NOTE Autonomy (formerly Verity) K2 software, which underlies the optional FileNet P8 Platform Content Search Engine component, installs with a large body of documentation that is not included in the general FileNet P8 documentation. For details about how to access this documentation, see the "Configure Content Engine for Content-Based Retrieval" topic in the IBM FileNet P8 Installation and Upgrade Guide.

Stop words

To improve Content Search Engine throughput by up to 30%, update the stop words list to include the most obvious case variations of a word. For example, the word "the" has three obvious variations: "The", "the", and "THE".

Application servers

For additional application server information that might apply to Content Engine and other IBM FileNet applications, see "[Application servers](#)" on page 32.

WebLogic

Group membership search

Because Content Engine uses WebLogic LDAP configuration for authentication but not authorization, the IBM Performance Lab suggests that you change the WebLogic "Group Membership Searching" setting from its default value of Unlimited to Limited, which will prevent WebLogic from performing unnecessary group membership searches that might cause high CPU usage on the LDAP repository server.

NOTE Verify the companion parameter "Group Membership Search Level" is set to its default: 0.

Adjusting the group membership searching setting

1. Logon to the WebLogic Administrative Console.
2. Select Security Realms.
3. Select myrealm.
4. Select Providers.
5. Select the [Provider Name].
6. Select Provider Specifics.
7. Click Lock and Edit.
8. Change the "Group Membership Searching" to Limited.
9. Click Activate Changes.

Operating systems

Windows 2003

Improving performance with “Maximize data throughput for networks applications”

To improve the performance of Content Engine transactions on a Windows 2003 operating system, select “Maximize data throughput for network applications” from the File and Printer Sharing for Microsoft Networks properties.

Adjusting the File and Printer Sharing properties

1. From the Start menu, select Settings > Control Panels > Network Connections.
2. Right-click the network connection, and then select Properties.
3. From the General tab, select File and Printer Sharing for Microsoft Network, and then click Properties.
4. Select Maximize data throughput for network applications.

Fixed content devices

EMC Centera

Set the maximum number of connections for a Centera fixed content device to a multiple of 15 according to the number of access nodes configured on the device. That is, if 2 access nodes are configured, set the Maximum Connections to 30.

Recommendations for handling a large number 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 methods that retrieve folder contents return a list of all folders and objects contained in the folder and store 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 Content Engine 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 suggests 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 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 suggests:

When filing in a folder, the application ensures that unique containment names are created for each object so that auto-unique naming is not overused. The Auto Unique setting when filing an object can seriously affect performance, and even cause timeouts, when it is used for a large number of objects that are filed in the same folder using a duplicate containment name.

- 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.
- 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.

Some operations 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 other methods to obtain the containees.
- 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 method that retrieves all containees.

- 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

Avoiding index skew

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 Content Engine that uses bound parameters.

NOTE 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 create an index skew, which can affect:

- Browsing operations
- INSUBFOLDER searches

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:

- Large numbers of subfolders directly under a parent
- Large numbers of documents or custom objects filed under one folder

Correspondingly, the two primary database optimizer issues encountered for those two cases are:

- 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

To diagnose slow browsing operations during full-table scans, obtain histograms as follows.

- Browsing operations on folders:

```
select count(*), parent_container_id from container group by parent_container_id
order by 1 desc
```

- Browsing operations on documents and custom objects:

```
select count(*), tail_id from relationship group by tail_id order by 1 desc
```

The 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.

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 DB2

For information, see the *Enabling DB2 Index Use with Skewed Data in the FileNet Content Engine* technote at:

http://www-01.ibm.com/support/docview.wss?rs=3278&context=SSNVNV&context=SSTHRT&context=SSNVUD&context=SSS236&context=SSNW2F&dc=DB560&dc=DB520&uid=swg21402554&loc=en_US&cs=utf-8&lang=en

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 (That is, 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".

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.

File storage area paths

Although Enterprise Manager has the option of creating file storage areas with local or UNC paths, it is recommended that you use a local path in all of the following configurations to help prevent performance degradation.

- A single Content Engine with a local or a SAN disk, or both
- Multiple Content Engines using GPFS (/vol/fs1 or F:\)
- Multiple Content Engines that use NFS and a consistent path on all UNIX systems (/vol/fs1)
- Multiple Content Engines that map UNC to a consistent drive on all Windows systems (F:\-> \\fs1\fs1)

Content streaming

Large content uploads and downloads

(*Microsoft SQL Server 2005 only*) For information, see “[Large content uploads and downloads](#)” under the Databases > Microsoft SQL Server heading.

Inline Content Retrieval Limit and Put/Get Block Size parameters: Uploads and retrievals

Inline Content Retrieval Limit (a Content Engine parameter that is set in Enterprise Manager) and FileNet.Content.GetBlockSizeKB (a client, such as Application Engine Workplace, Workplace XT, eForms and so on, parameter that is set in the filenet.properties file) work together to define the maximum amount of data that can be retrieved from a file store, database, or fixed content device in a single RPC.

The default setting for the Inline Content Retrieval Parameter is 1MB with a maximum possible value of 10MB. The FileNet.Content.GetBlockSizeKB setting cannot exceed the amount defined for the Inline Content Retrieval Limit parameter. For example, if the Inline Content Retrieval Limit is 5MB, GetBlockSizeKB cannot be greater than 5MB. If it is, the maximum amount of data received in each RPC will still be only 5MB.

FileNet.Content.PutBlockSizeKB is also a parameter in the filenet.properties file. It defines how much data can be sent to Content Engine in a single RPC. It has a preset maximum value of 10MB.

Increasing the values of these two filenet.properties parameters from their default setting (1024K) makes large content uploads and retrievals faster, but it also increases the memory footprint on the Content Engine server and client. Too many concurrent large creations and retrievals might cause memory contention.

Decreasing the values of these two filenet.properties parameters lowers the memory footprint but also causes extra round trips to the server and reduces performance. For example, if you set these values to 512K and you retrieve 1024K of content, the server will make two round trips instead of one.

Adjust these values according to your use case and workload.

Adjusting the Content Retrieval Limit parameter

1. Log on as object store administrator.
2. Expand the Sites node.
3. Right-click a site, and select Properties.
4. Go to the Content tab.
5. Adjust the Inline Content Retrieval Limit (kilobytes) field.
6. Click OK or Apply when you are done.

Adjusting the Put/Get Block Size parameters when an application server hosts the Content Engine client

1. Place these lines in the `filenet.properties` file:

```
FileNet.Content.PutBlockSizeKB=1024
```

```
FileNet.Content.GetBlockSizeKB=1024
```

2. Save the file to the following location, for example:

WebSphere – `C:\Program Files\IBM\WebSphere\AppServer\profiles\your_profile`

WebLogic – `C:\bea\user_projects\domains\your_domain`

NOTE This location corresponds with the directory in which you start the application server.

ABOUT THE FILENET.PROPERTIES FILE The `filenet.properties` file provides a way to edit certain preset configuration parameters. Because the configuration parameters are preset, the `filenet.properties` file does not exist until a preset configuration value that is managed by this file requires an adjustment. If the `filenet.properties` file has already been created, it will be located in one of the following locations. If it has not been created, place it in one of the following locations:

- The current directory (Java system property `user.dir`)
- The user's home directory (Java system property `user.home`)
- The Java JVM home directory (Java system property `java.home`)
- The classpath

Adjusting the Put/Get Block Size parameters when the Content Engine client is a stand-alone Java application

For stand-alone Java client applications, a switch needs to be applied to the application executable. For example: `java.exe -DfileNet.Content.PutBlockSizeKB=1024`.

Content upload destination

When content is stored in file stores and fixed stores, the content is uploaded into an inbound folder located in the same location as the local file directory, as specified for the file store. Each file store has only one inbound folder; and if a large volume of content uploads is expected, the inbound folder can become congested.

Better throughput can be achieved by providing multiple file stores, which enables content to be uploaded into multiple inbound folders. Multiple inbound folder effectively distribute the demands on the operating system administration of the file directory information and, if located on separate physical drives, the disk I/O. Multiple file storage areas can be made transparent by utilizing the Storage Policy feature. This feature provides load-balancing capabilities for content storage by transparently spreading the content across multiple storage areas.

When content is stored in database storage, content is uploaded to the folder that is designated for temporary file storage. By default, this is the same as the installation location of the Content Engine server. To lessen contention with other system processes and increase throughput, you can direct uploaded content to a dedicated portion of the file system. This can be done by using the Temporary Directory property on the Content Configuration tab in the object store properties.

Auditing

Delete old records from the Event table

Auditing events are stored in the Event table. Over time, 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 suggests that you first back up the table during a time when users are not accessing the system.

On Oracle systems, issue the truncate command to free any 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 will no longer have access to events information through the AuditedEvent 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 export file prior to performing the deletion or truncation.

Object state snapshots

Content Engine 4.5 introduces a feature where you can control what type of object state snapshots are retained for auditing events. These snapshots are accountable for the majority of the space consumed by the Event table.

To minimize the size of the Event table you can configure your audit definitions to reduce the number of snapshots captured. By default there are two snapshots enabled. One snapshot is for the original object state and another snapshot for the modified state.

Disabling object state snapshots

To disable one or both of the object state snapshots:

1. In Enterprise Manager, right-click on a Document Class Properties node.
2. Select the Audit Definitions tab.
3. Select one of the following values from the Object State Recording Level field:
 - None
 - Modified object only
 - Original and modified object

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)
```

Business Process Management

Process Engine

Databases

For additional database information that might apply to Process Engine and other IBM FileNet applications, see "[Databases](#)" on page 20.

Oracle – Sequence cache: Preallocated sequence numbers

The Sequence Cache parameter setting specifies the number of sequence values that are preallocated and kept in memory by the database for faster access.

By increasing this value for Process Engine sequences, you decrease the number of times that the database needs to be updated, which result in improved performance. The IBM Performance Lab found that a setting of 100 improved the performance of test systems.

Adjusting the sequence cache

1. Run the following command to adjust the sequence cache:

```
alter sequence f_sw.SequenceName cache 100;
```

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 that execute 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 through 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 about optimizing or increasing the amount of space that is 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.

Adjusting the parameter values

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

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

Client tuning – JVMs

For additional JVM tuning information that might apply to Process Engine, eForms, and other IBM FileNet applications, see "[JVM tuning](#)" on page 36 or the appropriate application-specific sections.

Use the following tuning tips, according to the application server, if Process Engine:

- Becomes unresponsive after extended durations of high workload conditions.
WebSphere – This failure manifests in the form of webcontainer threads in a 'Blocked State' due to SocketTimeouts.
Other applications servers – Drops in throughput might be observed.
- Shows poor performance under high workload conditions.

ORB tuning with Process Engine client connection management

To improve throughput and response times even further, P8PE-4.0.2-001 introduced a Process Engine client feature known as the Process Engine ORB pool. The Process Engine ORB pool is available to both application servers and stand-alone Java clients that use one of the following JVMs. Leveraging Process Engine ORB pool feature involves adding JVM-specific configuration settings to the appropriate startup files.

The minimum and maximum values shown are examples only. Appropriate values depend on the expected number of concurrent requests.

The following line in an error stack trace is an indicator that the problem can be resolved by tuning the Process Engine ORB pool parameters

Caused by: java.lang.OutOfMemoryError: unable to create new native thread

IBM JVM

```
set JAVA_OPTS=%JAVA_OPTS% -Dfilenet.pe.peorb.pool.min=10 -Dfilenet.pe.peorb.pool.max=20
```

NOTES

- This feature does not apply to applications that use WebSphere.
- If IBM JVM 5.0 or above is in use, see [“ORB tuning for IBM JAVA 5.0”](#) on page 63 for additional tuning tips.

SUN JVM

```
set JAVA_OPTIONS=%JAVA_OPTIONS% -Dfilenet.pe.peorb.pool.min=10  
-Dfilenet.pe.peorb.pool.max=20
```

```
set JAVA_OPTIONS=%JAVA_OPTIONS% -Dfilenet.pe.peorb.pool.guard=true
```

If the application server is WebLogic, the Process Engine client application must be forced to use the Sun implementation of the ORB client. This is accomplished by setting the following additional options in the application server startup file:

```
set JAVA_OPTIONS=%JAVA_OPTIONS%  
-Dorg.omg.CORBA.ORBClass=com.sun.corba.se.impl.orb.ORBImpl
```

```
set JAVA_OPTIONS=%JAVA_OPTIONS%  
-Dorg.omg.CORBA.ORBSingletonClass=com.sun.corba.se.impl.orb.ORBSingleton
```

ORB tuning for IBM JAVA 5.0 JVM

JVM FragmentSize parameter

Adjust the com.ibm.CORBA.FragmentSize parameter to 0, so that WebSphere webcontainer threads will read the entire message before attempting to dispatch a worker thread.

WebSphere 6.0 and 6.1	JBoss or Standalone IBM JAVA 5.0
<ol style="list-style-type: none"> 1. In WebSphere, navigate to Servers > Application Servers. 2. Select ServerName > Container Services > ORB Service > Custom Properties. 3. Set the com.ibm.CORBA.FragmentSize parameter to 0. 	<ol style="list-style-type: none"> 1. Make the following adjustment in the orb.properties file: 2. Set com.ibm.CORBA.FragmentSize=0. <p>By default, the orb.properties file is found in the java_home\jre\lib directory for the JVM, for example \IBM\Java50\jre\lib\orb.properties.</p>

JVM ConnectionMultiplicity property

The com.ibm.CORBA.ConnectionMultiplicity custom property causes a client ORB to multiplex communications to each server ORB. Each server ORB cannot have more than n concurrent sockets at any one time, as defined by this parameter.

The default value is 1. Increasing this value, might increase throughput under certain circumstances, particularly when a long-running, multithreaded process is acting as a client. The number of parallel connections can never exceed the number of requesting threads. The number of concurrent threads is therefore a sensible upper limit for this property. 100 is the suggested starting value.

WebSphere 6.1	JBoss or Standalone IBM JAVA 5.0
<ol style="list-style-type: none"> 1. In WebSphere, navigate to Servers > Application Servers. 2. Select ServerName > Container Services > ORB Service > Custom Properties. 3. Set the com.ibm.CORBA.ConnectionMultiplicity parameter to an appropriate value. 	<ol style="list-style-type: none"> 1. Make the following adjustment in the orb.properties file: 2. Set com.ibm.CORBA.ConnectionMultiplicity=100 (or an appropriate value) <p>By default, the orb.properties file is found in the java_home\jre\lib directory for the JVM, for example \IBM\Java50\jre\lib\orb.properties.</p>

JVM thread pool

Adjust the size of the thread pool to accommodate the number of threads that are required by the particular workload.

WebSphere 6.1	JBoss
<ol style="list-style-type: none"> 1. In WebSphere, navigate to Servers > Application Servers. 2. Select ServerName > Container Services > ORB Service. 3. Define the general properties, and then Under Thread Pool Settings, select "Use the Thread Pool Settings directly associated with the ORB service". 4. Click Apply. 5. Continue defining the general properties, and then select "Allow thread allocation beyond maximum thread size". 6. Click Apply. 	<ol style="list-style-type: none"> 1. Using a text editor, edit the following .xml files: <ul style="list-style-type: none"> • jboss-minimal.xml • jboss-service.xml <p>By default, these files are typically found in \JBoss405\server\default\conf directory, where default is the name of the JBoss profile.</p> 2. Locate the section for the mbean: <pre><mbean code="org.jboss.util.threadpool.BasicThreadPool" name="jboss.system:service=ThreadPool"></pre> 3. Increase the MaximumPoolSize in both files. Adjust this value according to your business requirements. However, a suggested starting value is 300.

Process Engine ORB pool in JBoss

P8PE-4.0.2-001 introduced a Process Engine client feature known as the Process Engine ORB pool. For more information about using this feature with JBoss, see "[ORB tuning with Process Engine client connection management](#)" on page 62.

ORB tuning for Sun JVM

ORB implementation in WebLogic 8.x and 9.x

Under high workload conditions, performance might improve if the client is forced to use the Sun implementation of the ORB client. To force the use of the Sun implementation, set the following options in the application server startup file:

```
set JAVA_OPTIONS=%JAVA_OPTIONS%  
-Dorg.omg.CORBA.ORBClass=com.sun.corba.se.impl.orb.ORBImpl  
  
set JAVA_OPTIONS=%JAVA_OPTIONS%  
-Dorg.omg.CORBA.ORBSingletonClass=com.sun.corba.se.impl.orb.ORBSingleton
```

Process Engine ORB pool in WebLogic 8.x and 9.x

P8PE-4.0.2-001 introduced a Process Engine client feature known as the Process Engine ORB pool. For more information about using this feature with WebLogic, see [“ORB tuning with Process Engine client connection management”](#) on page 62.

Process Engine ORB pool in JBoss

P8PE-4.0.2-001 introduced a Process Engine client feature known as the Process Engine ORB pool. For more information about using this feature with JBoss, see [“ORB tuning with Process Engine client connection management”](#) on page 62.

ORB tuning for JRockit JVM

For information, see the previously mentioned topics: [“ORB implementation in WebLogic 8.x and 9.x”](#) and [“Process Engine ORB pool in WebLogic 8.x and 9.x”](#).

Process Task Manager

Minimum and maximum heap size for high volume Component Manager activities

Complete the following procedure to avoid out-of-memory errors during high volume Component Manager activities, such as Records Manager 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`.

Process Engine authentication memory allocation for workflows

Workflows that have a large number of users or groups can encounter processing errors if they do not have enough memory allocated to Process Engine authentication when they run.

Following is a sample error message: “Problem with group expansion”.

To prevent this error from occurring, use Process Task Manager to adjust the `-Xmx` setting that is included in the Java command that initiates `PEDirectoryServer`.

Adjusting memory allocation to Process Engine authentication

1. From the Advanced tab in Process Task Manager, specify a Property Name of "PEDirectoryServer.Xmx" and enter a value such as "-Xmx768M".
2. Restart Process Engine services for the change to take effect.

VWK process optimization

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

The default maximum number of VWKs is 5. The IBM Performance Lab suggests you allocate 2 to 5 VWKs per CPU.

Adjusting the number of VWKs

1. Open the Process Task Manager.
2. Right-click a Process Engine server node and select Properties.
3. Select the General tab.
4. Enter the new default number of VWKs in the "Number of VWKs Processes" field.

NOTES

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

When to add more silos

Individual VWKs coordinate access to shared memory resources within the Process Engine. If you significantly increase the number of VWKs, for example from the default of 5 to 10, you might find that the individual VWKs 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 VWK processes are waiting for memory interlock, run the vwspy program and use its "state" command. Look for lock or semaphore contention among the various VWKs. Run the "state" command several times and in rapid succession over several minutes when the system is busy. If you consistently see a VWKs process holding a lock and you see multiple other VWKs 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 VWKs do not block each other when they update this memory.

Adjusting 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 Process Engine server (or in the /fnsw/local/sd/ directory of a UNIX-based Process Engine server).
3. Restart the Process Engine 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 VWKs.

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.

Cached Entry Timeout parameter: User and group cache

(Process Engine 4.0.2 and earlier) The Cached Entry Timeout parameter specifies how often (in hours) Process Engine refreshes user and group cache. Caching user and group information causes Process Engine to make fewer calls to Content Engine (and the LDAP server) and provides better throughput. However, system response times might increase when the user and group cache refreshes. For example, caching systems that have roster security, 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 Process Engine 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 frequently, set the Cached Entry Timeout parameter to the number of hours that Process Engine should execute a refresh.

Adjusting the Cached Entry Timeout parameter

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

Manually refreshing the user and group cache

If the directory service rarely changes, and you want to update the cache only as necessary, set the Cached Memory Timeout to zero and manually update the cache as needed by using the flushusercache command.

Scheduling 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 i f m y command as a scheduled task.

Flush Granularity option: Cache refresh

(Process Engine 4.0.2 and earlier) Refreshing the user environment cache is part of the routine Process Engine server operations. You control the refresh period by setting the Cached Entry Timeout parameter in Process Task Manager.

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 can 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 Process Engine system is restarted.

The default value for flushGranularity is two seconds. The optimal setting for flushGranularity differs depending upon the 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.

Adjusting 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 Process Engine server's node properties, in Process Task Manager.

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

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 Process Analyzer is measured by the publishing rate. The higher the publishing 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

Process Analyzer 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 is taking too long, consider upgrading your disk hardware or reduce 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

Workplace

Databases

For additional database information that might apply to Workplace and other IBM FileNet applications, see “[Databases](#)” on page 20.

Folder browsing documents and custom objects

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

For more information, see “[Recommendations for handling a large number of folders, documents, and objects](#)” on page 53.

Java client cache settings

Business Process Manager (BPM) applets, such as the Process Designer, require all Process Engine 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.

Image Viewer

For information about tuning Image Viewer, see the ViewOne performance optimization white paper, which can be downloaded from the Daeja Image Systems Web site at:

<http://www.daeja.com/white/white.asp?id=5&c=1>.

eForms

Databases

For additional database information that might apply to eForms and other IBM FileNet applications, see [“Databases”](#) on page 20.

High response times when rendering forms

Class ----- Database table name	Database column name	FileNet P8 property name	Reason to create index
Objectstore > Document Class docversion	Formpolicylaunch ed workflow	N/A	Create this single property index when both of the following conditions are true: <ul style="list-style-type: none"> • The number of records in the table exceeds 1,000,000. • You encounter high response times when viewing form policies.

Application servers

JVM tuning

For JVM tuning information that might apply to eForms (prior to version P8eF-4.0.1-002) and other Process Engine clients, see [“Client tuning – JVMs”](#) on page 20.

Workplace XT

Databases

For additional database information that might apply to Workplace XT and other IBM FileNet applications, see "[Databases](#)" on page 20.

Application servers

For additional application server information that might apply to Workplace XT and other IBM FileNet applications, see "[Application servers](#)" on page 32.

JVM tuning

For additional JVM tuning information that might apply to Workplace XT and other IBM FileNet applications, see "[JVM tuning](#)" on page 36.

IBM InfoSphere Enterprise Records and IBM FileNet Records Manager

Databases

The following Records Manager database information applies to all supported databases.

For additional database information that might apply to InfoSphere Enterprise Records and other IBM FileNet applications, see [“Databases”](#) on page 20.

Performance enhancing indexes

Consider using one or more of the following database indexes to improve system performance.

For more performance enhancing indexes, see:

- [Databases > Performance enhancing indexes](#) on page 21.
- [Databases > DB2 > Performance enhancing indexes](#) on page 24.

NOTE The value of xx varies depending on the data model and database type that is being used.

Disposition Sweep

Disposition Sweep, Hold Sweep, and stored searches can rely on user-defined searches. For instance, Disposition Sweep is controlled by a schedule, which uses a trigger to build a SQL statement to select entities ready for disposition processing. In this case, the trigger contains one or more properties that are evaluated according to the rules defined by a user. In most cases, those properties can benefit greatly from being indexed.

However, if a search is not used on a regular basis, the overhead associated with indexing that search might not be appropriate for daily processing. If this is the case, it might make sense to index the search only when it will be used. The index can then be removed after the sweep processing finishes.

For performance reasons, it is important that the records administrator and the database administrator work together to ensure that all of the properties used in user-defined searches be properly indexed. Microsoft’s SQL Profiler, Oracle’s Statspack, Oracle ADDM, or DB2 Design Advisor is tools that they can use to help to identify the properties (columns) that should be indexed.

Class ----- Database table name	Database column name	FileNet P8 property name	Reason to create index
Objectstore > FPOS > Other Classes > Custom Object > AccessRole	uxx_ applicationname	Application Name	<i>(File plan object stores only)</i> Create this single property index to avoid timeouts when users try to log on. Records Manager 4.0 Create this index only if you have used Sweep.
Generic			

Class ----- Database table name	Database column name	FileNet P8 property name	Reason to create index
			<p>Records Manager 4.5</p> <p>Create this index only if you have used Sweep in any version of Records Manager before you upgraded to Records Manager 4.5.</p> <p>This index is not required on any new installation of Records Manager, even if you have used Sweep.</p>

Disposition Sweep and Auto Destroy

Class ----- Database table name	Database column name	FileNet P8 property name	Reason to create index
Objectstore > FPOS > Document Class > Record docversion	uxx_ cutoffdate	Cut Off Date	<p><i>(File plan object stores only)</i></p> <p>Create this single property index when performance degrades during the following actions:</p> <ul style="list-style-type: none"> • Sweep performs cutoff processing. • Sweep resets disposition. • The user has changed to a different schedule.
Objectstore > FPOS > Document Class > Record docversion	uxx_ cutoffinherited from	Cutoff Inherited From	<p><i>(File plan object stores only)</i></p> <p>Create this single property index when performance degrades during the following actions.</p> <ul style="list-style-type: none"> • Sweep performs cutoff processing. • Sweep resets disposition. • The user has changed to a different schedule.

Class ----- Database table name	Database column name	FileNet P8 property name	Reason to create index
Objectstore > FPOS > Document Class > Record docversion	uxx_currentphaseaction	Current Phase Action	<p><i>(File plan object stores only)</i></p> <p>Create this single property index when performance degrades during the following actions:</p> <ul style="list-style-type: none"> • Sweep performs a phase transition • Sweep launches a cutoff workflow • Sweep performs complex phase processing (such as during a change in a phase or during alternate retentions)
Objectstore > FPOS > Document Class > Record docversion	uxx_currentphasexecutiondate	Current Phase Execution Date	<p><i>(File plan object stores only)</i></p> <p>Create this single property index when performance degrades during the following actions:</p> <ul style="list-style-type: none"> • Sweep performs a phase transition • Sweep launches a cutoff workflow • Sweep performs complex phase processing (such as during a change in a phase or during alternate retentions)
Objectstore > FPOS > Document Class > Record docversion	uxx_exportstatus	Export Status	<p><i>(File plan object stores only)</i></p> <p>Create this single property index when performance degrades during the following actions:</p> <ul style="list-style-type: none"> • Sweep performs a phase transition • Sweep launches a cutoff workflow • Sweep performs complex phase processing (such as during a change in a phase or during alternate retentions)

Class ----- Database table name	Database column name	FileNet P8 property name	Reason to create index
Objectstore > FPOS > Document Class > Record docversion	uxx_isdeleted	Is Deleted	<p><i>(File plan object stores only)</i></p> <p>Create this single property index when performance degrades during the following actions:</p> <ul style="list-style-type: none"> • Sweep performs a phase transition • Sweep launches a cutoff workflow • Sweep performs complex phase processing (such as during a change in a phase or during alternate retentions)
Objectstore > FPOS > Document Class > Record docversion	uxx_onhold	On Hold	<p><i>(File plan object stores only)</i></p> <p>Create this single property index when performance degrades during the following actions:</p> <ul style="list-style-type: none"> • Sweep performs a phase transition • Sweep launches a cutoff workflow • Sweep performs complex phase processing (such as during a change in a phase or during alternate retentions)
Objectstore > FPOS > Document Class docversion	uxx_record information	Record Information	<p><i>(File plan object stores only)</i></p> <p>Create this single property index when performance degrades during the following actions:</p> <ul style="list-style-type: none"> • Sweep performs a phase transition • Sweep launches a cutoff workflow • Sweep performs complex phase processing (such as during a change in a phase or during alternate retentions) <p>NOTE This property is automatically indexed in Records Manager 4.0, but is not automatically indexed in Records Manager 4.5.</p>

Class ----- Database table name	Database column name	FileNet P8 property name	Reason to create index
Objectstore > ROS > Document Class Docversion	uxx_record information	Record Information	<p><i>(Record object stores only)</i></p> <p>For record object stores, create this single property index during the following actions:</p> <ul style="list-style-type: none"> • CBR searches from Hold Sweep • CBR searches from the Administration Web application. • Selecting the Get action for the Record Information property from Workplace or Workplace XT.

Disposition Sweep, Hold Sweep, and Auto Destroy

(For Content Engine versions prior to 4.0.1 that use a Microsoft SQL Server database) 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.

Actions properties

Class ----- Database table name	Database column name	FileNet P8 property name	Reason to create index
N/A docversion	uxx_vitalrecordreviewaction ----- uxx_currentphaseaction	N/A	Create this composite index to improve the performance of the Get Info actions.

Disposal instruction properties

Class ----- Database table name	Database column name	FileNet P8 property name	Reason to create index
Objectstore > FPOS > Other Classes > Folder docversion		Disposition Instruction (DoD and Base) or Disposition Schedule (PRO)	Create this single property index to Improve response time when navigating disposal instructions property pages.

Location properties

Class ----- Database table name	Database column name	FileNet P8 property name	Reason to create index
N/A	uxx_ barcodeid	N/A	Create this composite index to improve retrieval of physical records. NOTE This index is beneficial only when you store physical records. The index is not needed when you store only electronic records.
docversion	----- -- uxx_ rmentitydescription ----- -- uxx_rmentitytype		
Objectstore > FPOS > Document Class > Record	uxx_location	Location	Create this composite index to improve retrieval of physical records. NOTE This index is beneficial only when you store physical records. The index is not needed when you store only electronic records.
docversion			
Objectstore > FPOS > Document Class > Record > Marker Record	uxx_homelocation	Home Location	Create this composite index to improve retrieval of physical records. NOTE This index is beneficial only when you store physical records. The index is not needed when you store only electronic records.
docversion			
Objectstore > FPOS > Other Classes > Folder > RM Folder > Record Folder	uxx_location	Location	Create this composite index to improve retrieval of physical records. NOTE This index is beneficial only when you store physical records. The index is not needed when you store only electronic records.
docversion			

Record type properties

Class ----- Database table name	Database column name	FileNet P8 property name	Reason to create index
Objectstore > FPOS > Document Class > Record		Type (DoD and Base)	Create this single property index to improve response time when navigating through record type property pages.
docversion		or Record Type (PRO)	

SupersedeBy property

Class ----- Database table name	Database column name	FileNet P8 property name	Reason to create index
Objectstore > FPOS > Document Class docversion	uxx_ supersededdat	Superseded By	Create this single property index to improve the performance of the user interface retrieving values from the enumeration property "Superseded Records". NOTE This is an object value property (OVP). For more information, see " Indexing object value properties " on page 23 .

Oracle

Prevent performance degradation during INSUBFOLDER searches

In general, do not use the dbms_stats statistics collection mechanism with its default settings without first reviewing and understanding the impact of the custom setting options.

With Records Manager, certain default settings result in poor query plans that degrade the performance of INSUBFOLDER searches. At a minimum, the IBM Performance Lab suggests 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 values that produce the best results.

Application servers

For additional application server information that might apply to Records Manager and other IBM FileNet applications, see "[Application servers](#)" on page 32.

JVM tuning

For additional JVM tuning information that might apply to Records Manager and other IBM FileNet applications, see "[JVM tuning](#)" on page 36.

Thread stack size

The "-Xss" switch value in the startweblogic.cmd specifies how much memory is allocated to the thread stack in kilobytes. Increasing this switch value can help prevent stack overflow errors that cause system fails when threads are computation intensive and need extra memory.

Over time, Records Manager user interface operations might fail because of a `java.lang.StackOverflowError` exception. If you encounter this problem, consider increasing the thread stack size, for example:

```
MEM_ARGS= -Xss2m
```

Set the minimum and maximum heap size for the Dynamic Hold Sweep execution

To avoid out of memory errors while performing bulk operations in a batch during Dynamic Hold Sweeps, increase the heap size for the Records Manager 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

1. Identify the heap size entries.

In the `RecordsManagerSweep` script, search for:

- o `RecordsManagerSweep.bat` :

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

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

2. Set the minimum and maximum heap size to at least 1024MB for sweep execution, for example:

- o `RecordsManagerSweep.bat` :

```
java -Xms1024m -Xmx1024m -jar -Dwaspl.location=./lib ./lib/%JarFile%
```
- o `RecordsManagerSweep.sh` :

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

3. Save the file.

Set the minimum and maximum heap size for Destroy Disposition Sweeps

For information, see “[Minimum and maximum heap size for high volume Component Manager activities](#)” on page 65.

Improve Sweep and Destroy throughput by tuning write threads

The IBM Performance Lab suggests that the Multi-threaded Write feature thread count for each of the following processes equal the number of logical CPUs on the Content Engine server:

- Disposition Sweep
- Dynamic Holds Sweep
- Auto Destroy (Available in Records Manager 4.5.)

NOTE Use the Disposition Sweep console, the Dynamic Holds Sweep console, and the Auto Destroy console respectively, to set the thread count for the Multi-threaded Write feature.

Content Federation Services

Federating documents

Chunk Size parameter: Allowable size of query results

The Federation Administration Chunk Size parameter value specifies the allowable size of query results. This parameter works with the Content Federation Services Chunking feature, which enables a large number of documents to be federated by using a single rule. Large query results are divided into smaller, more manageable “chunks” in order to prevent out-of-memory errors.

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

Adjusting the Chunk Size parameter

Adjust the Chunk Size parameter in the General Configuration section of the Federation Administration user interface.

Batch Size parameter: Number of documents batched by the exporter

The Federation Administration Batch Size parameter specifies the number of documents that are 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 content is received from the source repository, and increases the performance of the exporter.

Adjusting the Batch Size parameter

Adjust the Batch Size parameter in the General Configuration section of the Federation Administration user interface.

Importer Threads and Dispatcher Wait Time parameters: Importer throughput

To increase the importer rate, consider increasing the Importer Threads parameter and decreasing the Dispatcher Wait Time Interval parameter.

USAGE CONSIDERATION Adjust these values to optimize your system resources.

Adjusting the Import Threads and Wait Time Interval parameters

1. In Enterprise Manager, right-click the domain and select Properties.
2. Select the CFS Import Agent tab.
3. Increase the Importer Threads value.
4. Decrease the Dispatcher Wait Time Interval value.

Maximum Number of (Exporter) Workers parameter: Exporter throughput

To increase performance of the exporter, consider increasing the value of the Federation Administration Maximum Number of (Exporter) Workers parameter and using multiple instances of the exporter.

Content Services exception If Content Services is the source repository, reduce the value of the Maximum Number of (Exporter) Workers parameter to 1 and use only a single exporter instance.

Usage consideration Adjust these values to optimize your system resources.

Adjusting the Maximum Number of (Exporter) Workers

1. In the General Configuration screen of the Federation Configuration user interface, scroll down to the Export Workers section.
2. Adjust the value of the Maximum Number of Workers parameter.

Using multiple exporter instances

For information about using more than one exporter instance, see "Set up multiple Content Federation Services exporters" in the IBM FileNet P8 Content Federation Services Installation and Upgrade Guide. To download this guide from the IBM support page, see [Access IBM FileNet Documentation](#) on page 16.

RMI bridges

The IBM Performance Lab suggests running the RMI bridges on a remote server to prevent any reduction of throughput due to resource constraints. Additionally, using multiple RMI bridges can increase the performance of the Content Federation Services exporter.

Configuring RMI bridges on a remote server

1. Install just the IICE connectors on a remote server.
2. Launch the RMIBridge.bat file on the same remote server.
3. Launch the IICE Administration Tool on the Content Federation Services/IICE server.
4. Select the connector that will be used to connect to the external repository.
5. Edit the "RMI Proxy Connector URL" settings.
6. Change the value from "LOCALHOST" to the actual hostname of the server that is running the RMI bridge.
7. Verify that the port number matches the port number from the running RMI bridge.
8. Save the changes and verify that the connection can be established.
9. Log in to the IICE Administration Tool on the IICE server.
10. Select the connector that is being used to connect to the external repository.
11. Enter the "RMI Proxy Connector URL" to add additional RMI entries.

Configuring multiple RMI bridges

For information, see the section about configuring multiple RMI bridges in the IBM FileNet P8 Content Federation Services Installation and Upgrade Guide. To download this document from the IBM support site, see "[Access IBM FileNet documentation](#)" on page 16.

Image Services

Cache Service Manager

Consider creating an Image Services cache on the Content Engine server to improve image retrieval performance when the Image Services server is in a remote location from the Content Engine server in a Content Federation Services-Image Services configuration.

When the Image Services server is not in the same physical location as the Content Engine server, creating an Image Services cache on the Content Engine server can provide a significant savings in time and resources by avoiding multiple round trips across the WAN for images. If your configuration has multiple Content Engine servers, you can create the Image Services cache on any of the Content Engine servers as long as they are all in the same physical location. If you have multiple Content Engine cluster member nodes in different geographic locations, consider creating an Image Services cache in each location. However, if you do create an Image Services cache in each location, you can create only one Image Services cache per location.

Import agents, Batch Delay parameter, and Document Batch Size parameter

Import agents

Increasing the number of import agents increases the overall throughput. However, increasing the number of import agents also:

- Increases the number of connections consumed on the Image Services server
Each import agent has two connections. One connection is for documents, and the other connection is for annotations. For example, four import agents on a Content Engine server use eight Image Services connections.
- Increases the JVM memory consumed on the application server
Monitor the application server heap size memory usage from the application server administration console, and increase the maximum heap size if possible.
- Increases the contention for each individual federation
Although throughput might increase, the time it takes to federate each document can decrease slightly due to database, memory, and CPU contention to create each document in the object store.

Batch Delay parameter

The Batch Delay parameter specifies the number of milliseconds that the import agent waits between batch requests before retrieving the next batch from Image Services.

Decreasing the batch delay processes more requests in a given amount of time, and is an approach that can be used as an alternative to adding more import agents.

Document Batch Size parameter

The Document Batch Size parameter specifies the number of documents that are federated in a batch.

Increasing the Document Batch Size parameter causes each import agent to federate a larger number of Image Services documents at a time, thus requiring fewer round trips between the Image Services server and Content Engine server.

However, making the batch size too large increases the chance that an import agent will not have enough time to federate an entire batch within the Doc services timeout period. If a batch is not fully federated within the timeout period, it can be re-processed by another Image Services import agent. A second Image Services import agent reprocessing a batch that is not fully federated results in duplication errors, which are automatically rolled back, and needlessly increases the CPU and database contention because of the duplicate commit actions and subsequent rollback actions.

Adjusting import agents, the Batch Delay parameter, and the Document Batch Size parameter

1. Launch Enterprise Manager.
2. Right-click the Domain Name and select Properties.
3. Click the IS Import Agents tab.
4. Adjust the following values, and then click OK:
 - Number of Import Agents
 - Batch Delay (milliseconds)
 - Document per Batch

Database indexing

The Image Services Catalog Export Application works mainly with the DOCTABA table, which holds the metadata for all of the documents in an Image Manager system. The DOCTABA table is typically a very large table with a small number of indexes defined for standard Image Manager operations.

The Image Services Catalog Export Application federates metadata from existing documents in Image Services to Content Engine. This is typically a one-time operation. The Image Services Catalog Export Application analyzes the document classes of all documents to determine if the documents are candidates for federation. Due to the nature of the query, no viable indexes exist that the database can use to facilitate the processing of the query issued without performing a full table scan. This process can be very time consuming.

The performance of this process can be improved by creating a database index on DOCTABA as described below. Creating this index helps the Image Services Catalog Export Application complete its processing. However creating this indexes also causes:

- A requirement for more space to hold the index data
- Potentially reduced performance on inserts into and deletes from the table

Because the Image Services Catalog Export Application reads only from DOCTABA, its performance improves only by the newly created index. The amount of time and space required to build the index varies based on each specific customer.

The IBM Performance Lab supports creating this specific index for the sole purpose of the export phase used in conjunction with the Image Services Catalog Export Application. However, once all documents have been federated to the appropriate object stores, the IBM Performance Lab suggests removing the index. This index has no added benefit for everyday Image Manager production use.

NOTE If this index is created for the export, the database administrator should closely monitor the system before and after the index creation to determine any negative effects resulting from the index. None are expected, but both performance and resource utilization should be monitored.

Connect to the database as f_sw.

DB2

```
create index f_docn_doccn_ix on doctaba  
("F_DOCCLASSNUMBER","F_DOCNUMBER")
```

Oracle

```
create index f_docn_doccn_ix on doctaba  
("F_DOCCLASSNUMBER","F_DOCNUMBER")  
tablespace table_space>;
```

SQL Server

```
create index f_docn_doccn_ix on doctaba  
("F_DOCCLASSNUMBER","F_DOCNUMBER")  
on file_group
```

Content Manager on Demand

Federation throughput

I_EXTERNALIDENTITY76 index

During periods of high volume federation, you can improve federation throughput on Content Manager On Demand by removing the I_EXTERNALIDENTITY76 index that is on the EXTERNALIDENTITY table.

Important Remove this index only for situations that involve high volume federation. When federation is finished, recreate the index. Removing the index permanently can degrade other operations.

```
CREATE UNIQUE INDEX I_EXTERNALIDENTITY76 ON EXTERNALIDENTITY"  
    ("EXTERNAL_ID" ASC,  
     "REPOSITORY_ID" ASC)  
ALLOW REVERSE SCANS;
```

System logging and trace

By default, Content Manager on Demand enables the following system and trace logging. Disabling all logging and trace on the Content Manager on Demand source repository dramatically improves the retrieval performance for federated documents.

Disabling logging

1. From the OnDemand Admin user interface, right-click the server icon and select System Parameters.
2. In the System Logging section disable:
 - Login
 - Logoff
 - Failed login
 - Application group messages
3. Click Update to save the changes.

Disabling tracing

1. From the OnDemand Admin user interface, right-click the server icon and select Trace Parameters.
2. Verify that none of the traces are selected.
3. Click Update to save the changes.

Content Services

Simultaneous library requests per user for federation queries and exports

By default, Content Services can process up to 10 simultaneous library requests per user credential. Content Federation Services use 10 export workers and 4 additional run workers by default, which causes an error during the query and exporting phases of federation.

To prevent this error, increase the number of simultaneous library requests per user to approximately 20.

Adjusting the number of simultaneous library requests per user

1. On the IDMWS server, select Start > Programs > FileNet > IDM > Configure.
2. Click Additional Preferences.
3. Locate the Content Services library and select Libraries.
4. Increase the value for Active Requests per User to 20.
5. Restart the RMI bridge.
6. Rerun the federation rule.

Cache parameters: Document federation with full-text indexing

Two cache parameters in the IDM settings affect Content Services document federation with full-text indexing:

- **Maximum Cache Size parameter** – Specifies the cache directory size that triggers the deletion of documents stored in cache.
- **Maximum Cached Documents parameter** – Specifies the number of documents that can be cached.

However, if the number of cached documents is reached before the directory cache size is reached, documents will no longer be cached, and Content Federation Services can no longer retrieve any documents. As a result, the following error will occur in the Content Engine trace logs: "CBR - ERROR Updating indexing request after failure..."

To prevent this error, adjust the maximum cache size parameter to a smaller value and adjust the maximum Cached Documents parameter to the largest possible value. The maximum value allowed is 100,000. The default cache directory is: c:\FileNet\IDM\Cache.

Adjusting the cache parameters

1. On the IDMWS server, select Start > Programs > FileNet > IDM > Configure.
2. Locate the Content Services library and select Additional Preferences.
3. Expand Directories and Files > Local Caching.
4. Set the "Maximum cache size".

NOTE This value should be greater than or equal to the largest document size in the source repository. For example, if the largest document in the source repository is 20MB, then set this value at 20MB or higher.

5. Increase the "Maximum cached documents" value to the maximum value (100,000).
6. Apply the changes.
7. Restart the IDM Web server.

Image Services

Cache Service Manager

Consider creating an Image Services cache on the Content Engine server to improve image retrieval performance when the Image Services server is in a remote location from the Content Engine server in a Content Federation Services-Image Services configuration.

When the Image Services server is not in the same physical location as the Content Engine server, creating an Image Services cache on the Content Engine server can provide a significant savings in time and resources by avoiding multiple round trips across the WAN for images. If your configuration has multiple Content Engine servers, you can create the Image Services cache on any of the Content Engine servers as long as they are all in the same physical location. If you have multiple Content Engine cluster member nodes in different geographic locations, consider creating an Image Services cache in each location. However, if you do create an Image Services cache in each location, you can create only one Image Services cache per location.

Import agents, Batch Delay parameter, and Document Batch Size parameter

Import agents

Increasing the number of import agents increases the overall throughput. However, increasing the number of import agents also:

- Increases the number of connections consumed on the Image Services server
Each import agent has two connections. One connection is for documents, and the other connection is for annotations. For example, four import agents on a Content Engine server use eight Image Services connections.
- Increases the JVM memory consumed on the application server
Monitor the application server heap size memory usage from the application server administration console, and increase the maximum heap size if possible.
- Increases the contention for each individual federation
Although throughput might increase, the time it takes to federate each document can decrease slightly due to database, memory, and CPU contention to create each document in the object store.

Batch Delay parameter

The Batch Delay parameter specifies the number of milliseconds that the import agent waits between batch requests before retrieving the next batch from Image Services.

Decreasing the batch delay processes more requests in a given amount of time, and is an approach that can be used as an alternative to adding more import agents.

Document Batch Size parameter

The Document Batch Size parameter specifies the number of documents that are federated in a batch.

Increasing the Document Batch Size parameter causes each import agent to federate a larger number of Image Services documents at a time, thus requiring fewer round trips between the Image Services server and Content Engine server.

However, making the batch size too large increases the chance that an import agent will not have enough time to federate an entire batch within the Doc services timeout period. If a batch is not fully federated within the timeout period, it can be re-processed by another Image Services import agent. A second Image Services import agent reprocessing a batch that is not fully federated results in duplication errors, which are automatically rolled back, and needlessly increases the CPU and database contention because of the duplicate commit actions and subsequent rollback actions.

Adjusting import agents, the Batch Delay parameter, and the Document Batch Size parameter

1. Launch Enterprise Manager.
2. Right-click the Domain Name and select **Properties**.
3. Click the **IS Import Agents** tab.
4. Adjust the following values, and then click **OK**:
 - Number of Import Agents
 - Batch Delay (milliseconds)
 - Document per Batch

Database indexing

The Image Services Catalog Export Application works mainly with the DOCTABA table, which holds the metadata for all of the documents in an Image Manager system. The DOCTABA table is typically a very large table with a small number of indexes defined for standard Image Manager operations.

The Image Services Catalog Export Application federates metadata from existing documents in Image Services to Content Engine. This is typically a one-time operation. The Image Services Catalog Export Application analyzes the document classes of all documents to determine if the documents are candidates for federation. Due to the nature of the query, no viable indexes exist that the database can use to facilitate the processing of the query issued without performing a full table scan. This process can be very time consuming.

The performance of this process can be improved by creating a database index on DOCTABA as described below. Creating this index helps the Image Services Catalog Export Application complete its processing. However creating this indexes also causes:

- A requirement for more space to hold the index data
- Potentially reduced performance on inserts into and deletes from the table

Because the Image Services Catalog Export Application reads only from DOCTABA, its performance improves only by the newly created index. The amount of time and space required to build the index varies based on each specific customer.

The IBM Performance Lab supports creating this specific index for the sole purpose of the export phase used in conjunction with the Image Services Catalog Export Application. However, once all documents have been federated to the appropriate object stores, the IBM Performance Lab suggests removing the index. This index has no added benefit for everyday Image Manager production use.

NOTE If this index is created for the export, the database administrator should closely monitor the system before and after the index creation to determine any negative effects resulting from the index. None are expected, but both performance and resource utilization should be monitored.

Connect to the database as f_sw.

DB2

```
create index f_docn_doccn_ix on doctaba  
("F_DOCCLASSNUMBER","F_DOCNUMBER")
```

Oracle

```
create index f_docn_doccn_ix on doctaba  
("F_DOCCLASSNUMBER","F_DOCNUMBER")  
tablespace table_space>;
```

SQL Server

```
create index f_docn_doccn_ix on doctaba  
("F_DOCCLASSNUMBER","F_DOCNUMBER")  
on file_group
```

IBM FileNet Connectors for SharePoint

Thread count: Web Parts response times when archiving

To prevent slower response times from the Web Parts user interface when you archive documents from the SharePoint document libraries, consider decreasing the thread count that is allocated to the Task Route service.

By default, the Task Route service thread count is 16. Reducing this thread count to 4 reduces the rate at which records are archived, but it can also reduce overall CPU usage by up to 50%.

NOTE This information assumes that the Connector Administration services are not located on a remote server machine. To benefit from the suggested adjustment, the Connector Administration services must be collocated on the same server machine that hosts the SharePoint services and the SharePoint Connector services.

Adjusting the Task Route service thread count

1. Launch the Configuration Manager for the IBM FileNet SharePoint Connector for document libraries.
2. From the Tools menu, select Task Route Service Configuration > Thread Count.
3. Adjust the thread count value.

Overlapped recycling for an application pool

To maintain a healthy memory usage pattern, consider enabling overlapped recycling for the IIS application pools running WebParts and the DocLib Connector Web service.

The IBM Performance Lab observed a short time of service interruption (<10 seconds) when a worker process recycle occurs. Therefore, the IBM Performance Lab suggests that you set the recycling schedule during off-hours when the workload is at a minimum.

Scheduling recycling for an application pool

For information, see "[Recycle Worker Process parameter \(IIS 6.0 for .Net API\)](#)" on page 42.

Batch process and user operations

For information, see "[Recycle Worker Process parameter \(IIS 6.0 for .Net API\)](#)" on page 42.

Document libraries

Separate application pool for DocLib Web service

To increase the overall performance of SharePoint and to provide better error isolation and security, consider creating a separate application pool for the Web service in the FileNet SharePoint connector for Document Library 2.2.0.

Creating a new application pool

1. Open the Internet Information Services Manager by selecting Start > Programs > Administrative Tools > Internet Information Services.
2. Expand the local computer, right-click Application Pools, point to New, and then click Application Pool.
3. If the ID that displays for Application Pool ID is not the correct ID, enter a new ID.
4. Under Application pool settings, click the appropriate setting. If you click Use existing application pool as a template, then select the template application pool from the Application Pool Name list.

Assigning FnSPWebservices Web site to an application pool

1. Open the Internet Information Services Manager by selecting Start > Programs > Administrative Tools > Internet Information Services.
2. Expand the WebSites folder, right-click FnSpWebServices, and then click Properties.
3. Click the Home Directory tab.
4. Select the application pool name from the Application Pool list.

QueryPageSize setting affects throughput

In the FileNet SharePoint connector for Document Library 2.2.0, adjusting the QueryPageSize in the FnSpWebServices web.config file can produce better throughput.

The QueryPageSize is used for document crawling queries issued by the DocLib collector. This entry controls the maximum number of documents that will be batched and returned per round-trip. The default setting is 800.

To increase the potential throughput, consider increasing the QueryPageSize. Verify that sufficient system resources are present. For example, available CPU, fast disk storage, and a high network bandwidth. Then increase the QueryPageSize in increments of 400 observing the impact on throughput and system resource utilization until available resources are maximized.

Adjusting the QueryPageSize

1. Edit the web.config for the FnSpWebservice. The default location set by the installer is at C:\inetpub\wwwroot\FnSpWebServices on the SharePoint sever.
2. Find the section `<add key="QueryPageSize" value="800" />`
3. Change the value as necessary.

Collector schedules: Resource usage on idle systems with archiving

To prevent resource usage overhead due to the frequent polling and crawling that is associated with archiving, consider adjusting the SharePoint Collector schedule from Run Always to a specific time interval such as 15 minutes.

The IBM Performance Lab has determined that adjusting this setting can prevent CPU usage overhead that is associated with archiving by up to 15% on an idle system.

Adjusting the Collector schedule

1. Launch the Configuration Manager for the IBM FileNet SharePoint Connector for document libraries.
2. Select **Collectors** from the **Navigation** panel, and select the affected collector configuration.
3. Select the **Schedule** tab to change the collector schedule.

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:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at 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.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

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

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

GC31-5483-08

