


# Content Federation Services for Image Services (CFS-IS) V5.1.0

## Bi-directional replication performance tuning



© 2013 IBM Corporation

Today I will be presenting a course on performance tuning of CFS-IS bi-directional replication for IBM FileNet® Image Services and Content Engine. First, I would like to define some terms that I will be using. I will be referring to this feature using the term bi-di. The terms inbound and outbound refer to the direction data is traveling, and they are relative to the Content Engine. So inbound refers to data being federated from IS to CE, and outbound refers to replication from CE to IS.

## Introduction

- Course overview
  - Replication configuration parameters
  - Tuning guidelines
  - Troubleshooting
- Target audience:
  - Administrators, tech. support, consultants
- Prerequisites:
  - Knowledge of CFS-IS bi-directional synchronization
- Version release date: July 2011

In this course we will cover Replication Configuration Parameters, Tuning Guidelines, and troubleshooting. The target audience for this course are System Administrators, Technical Support personnel, and Consultants. The prerequisite for this course is Knowledge of CFS-IS bi-directional replication.

## Course objectives

After this course you will be able to:

- Adjust the replication configuration parameters to maximize performance
- Identify the causes of performance problems and take corrective action

Once this course is completed, you will be able to adjust the replication configuration parameters to maximize performance, identify the causes of performance problems, and take corrective action.

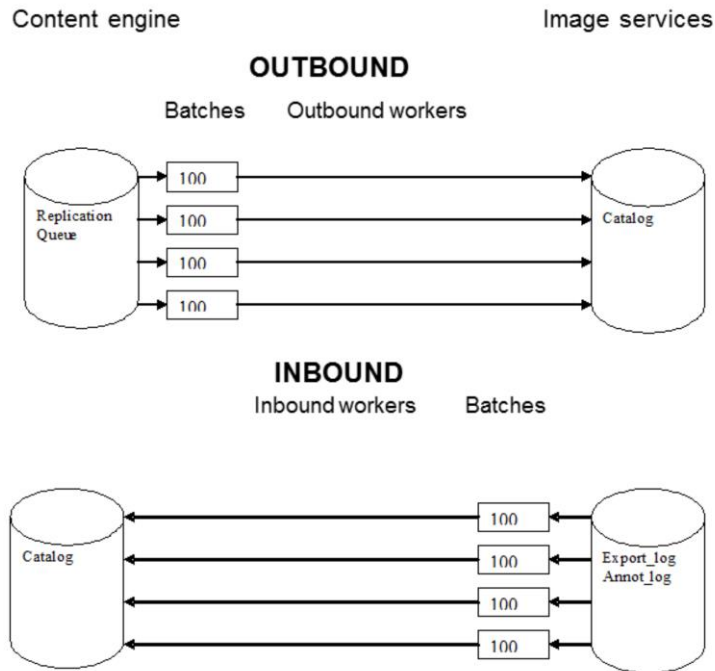
---

## Overview

- There is no formula or "cookbook"
- Each site must be uniquely tuned based on their load and mix of operations, using the guidelines to follow
- It is best to start with the default settings, and modify them as needed.

It is important to note that there is no formula or cookbook for performance tuning of bi-di. Each site must be uniquely tuned based on their load and mix of operations, using the guidelines to follow. Note that the guidelines are general; the person tuning a particular system is going to have to figure out what the best settings are for that system. It is best to start with the default settings and modify them as needed.

## Flow diagram



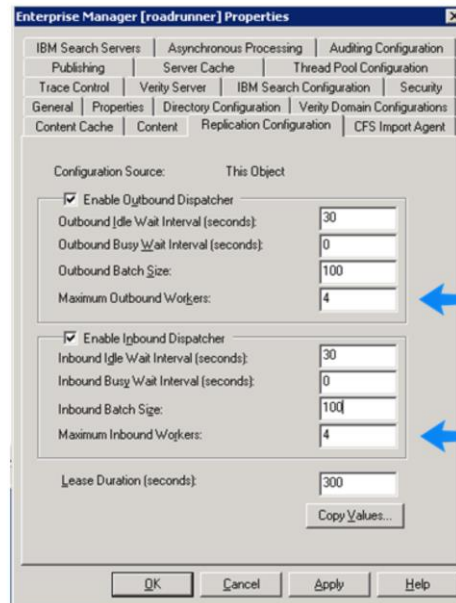
5

Bi-directional replication performance tuning

© 2013 IBM Corporation

This diagram shows the flow of bi-di, highlighting the parameters that can be used for performance tuning. These are the number of workers and the batch size. This shows the default settings, with 4 inbound and 4 outbound workers, both using a batch size of 100. Each worker will query the appropriate database table, retrieving records until there are none left or the batch is full, and then processing those records.

## Maximum replication workers (1 of 2)



The first parameters we will discuss are Maximum Replication Workers. This is the screen in FEM where these parameters can be set. There are separate settings for Maximum Inbound Workers, and Maximum Outbound Workers, and as you can see, the default is 4.

## Maximum replication workers (2 of 2)

- Number of background tasks running simultaneously, dedicated to processing replication actions
- To increase throughput, increase maximum workers parameters
- If high volume in only one direction, increase only the corresponding parameter

Maximum Replication Workers are the number of background tasks running simultaneously, dedicated to processing replication actions. In general, to increase throughput, you are going to want to increase the maximum workers parameters. If you have a lot of volume in only one direction however, you only want to increase the corresponding parameter.

## Batch size (1 of 4)

Enterprise Manager [roadrunner] Properties

IBM Search Servers | Asynchronous Processing | Auditing Configuration  
Publishing | Server Cache | Thread Pool Configuration  
Trace Control | Verity Server | IBM Search Configuration | Security  
General | Properties | Directory Configuration | Verity Domain Configurations  
Content Cache | Content | Replication Configuration | CFS Import Agent

Configuration Source: This Object

Enable Outbound Dispatcher

Outbound Idle Wait Interval (seconds): 30  
Outbound Busy Wait Interval (seconds): 0  
Outbound Batch Size: 100  
Maximum Outbound Workers: 4

Enable Inbound Dispatcher

Inbound Idle Wait Interval (seconds): 30  
Inbound Busy Wait Interval (seconds): 0  
Inbound Batch Size: 100  
Maximum Inbound Workers: 4

Lease Duration (seconds): 300

Copy Values...

OK Cancel Apply Help

The next parameter we will discuss is batch size. There are separate settings for inbound batch size and outbound batch size. The default for both is 100.



## Batch size (2 of 4)

- Maximum number of records processed by a single worker at a time
- If mostly inserts and exports, use a larger batch size
  - Groups of inserts and exports are processed in a single database transaction
  - Best possible performance is inserts and exports only, with no re-exports and no errors

The batch size is the maximum number of records processed by a single worker at a time. If you are doing mostly inserts or exports, you are going to want to use a larger batch size. By inserts, we mean new documents. The reason for the larger batch size is that groups of inserts and exports are processed in a single database transaction. The best possible performance is inserts and exports only, with no re-exports and no errors. Re-exports are handled differently than exports; they are treated as updates, which are individually processed in separate transactions.

## Batch size (3 of 4)

- If there are a lot of update, delete, or dereplicate actions, use a smaller batch size
  - For outbound, each of the update, delete, and dereplicate actions is grouped and processed together up to a maximum of 50. Any more than that are processed separately. This can result in an extra operation to handle just a few more records.
  - For inbound, each update is handled as a separate database transaction. Consecutive deletes are handled in the same transaction.
  - Multiple transactions or operations per batch increase risk of a timeout

If there are a lot of update, delete, or dereplicate actions, you want to use a smaller batch size. For outbound, each of the update, delete, and dereplicate actions is grouped and processed together up to a maximum of 50 each. Any more than that are processed separately. This could result in an extra operation to handle just a few more records. So for outbound, you want to set a batch size so that you do not get more than 50 updates, or 50 deletes, or 50 dereplicates in the same batch. For inbound, each update is handled as a separate database transaction. Consecutive deletes are handled in the same transaction. However, if you have deletes scattered throughout the batch, those will all be handled separately. The purpose of having a smaller batch is to avoid timeout errors. Multiple transactions or operations per batch increase the risk of a timeout. You want to avoid timeouts as much as possible. Getting a timeout results in the entire operation being repeated, which has a substantial performance impact.

## Batch size (4 of 4)

- If 25 or more workers of one type are configured, adjust batch size in the opposite direction of the number of workers
  - When increasing number of workers, decrease batch size
  - When decreasing number of workers, increase batch size
  - This is done to avoid timeouts

If you have 25 or more workers of one type configured, you want to adjust the batch size in the opposite direction of the number of workers. If you are increasing the number of workers, you are going to want to decrease the batch size, and if you are decreasing the number of workers, you are going to want to increase the batch size. The reason for this, is when you have that many workers, there is a lot of contention on the database. The larger the batch size, the bigger the risk of a timeout, which, as we discussed earlier, you want to avoid.

## Wait intervals (1 of 2)

Enterprise Manager [roadrunner] Properties

IBM Search Servers | Asynchronous Processing | Auditing Configuration  
Publishing | Server Cache | Thread Pool Configuration  
Trace Control | Verity Server | IBM Search Configuration | Security  
General | Properties | Directory Configuration | Verity Domain Configurations  
Content Cache | Content | Replication Configuration | CFS Import Agent

Configuration Source: This Object

Enable Outbound Dispatcher

Outbound Idle Wait Interval (seconds): 30

Outbound Busy Wait Interval (seconds): 0

Outbound Batch Size: 100

Maximum Outbound Workers: 4

Enable Inbound Dispatcher

Inbound Idle Wait Interval (seconds): 30

Inbound Busy Wait Interval (seconds): 0

Inbound Batch Size: 100

Maximum Inbound Workers: 4

Lease Duration (seconds): 300

Copy Values...

OK Cancel Apply Help

Our next topic is wait intervals. There are two wait interval parameters, busy and idle, which can be set for both inbound and outbound. In both cases, the default for the idle wait interval is 30 seconds, and the busy wait interval default is 0 seconds.

## Wait intervals (2 of 2)

- Idle – number of seconds for a worker to wait before asking for more work, when previous request returned nothing
- Busy – number of seconds for a worker to wait before asking for more work, after processing a batch.
- In low-volume situations, trade-off between higher performance for the rest of Content Engine, and delays in synchronization
  - Increase interval value to reduce how often workers request work, allow batch to fill
  - Decrease interval value to minimize delay in synchronizing documents and annotations.

The idle wait interval is the number of seconds for a worker to wait before asking for more work, when the previous request returned no records. The busy wait interval is the number of seconds for a worker to wait before asking for more work, after it has just processed a batch. In low-volume situations, there is a trade-off between higher performance for the rest of Content Engine, and delays in synchronization. If you prefer better content engine performance, increase the interval value. This will reduce how often workers request data and allow the batch to fill. However, it will increase the delay in synchronizing data. If you prefer to reduce delays in synchronizing documents and annotation, you are going to want to decrease this value.

## Remote image services cache

- Improves image retrieval performance when Image Services and Content Engine are in separate locations, communicating by way of WAN
  - Effective for content retrieved from IS - bi-di docs originating on IS, Content Engine docs using IS as a Fixed Content Device
- Install an Image Services Cache Server at the same site as the Content Engine - can also be installed on the Content Engine server
- Improve performance for images that are retrieved multiple times
- For performance improvements on images that are retrieved once, use prefetch
- Make sure cache is large enough to hold all images to be retrieved locally

Now we are going to discuss the remote image services cache. Configuring this cache will improve image retrieval performance when Image Services and Content Engine are in separate locations, communication via WAN. It is effective for content retrieved from IS; these are bi-di docs originating on IS, and CE docs using IS as a Fixed content Device. It does not have any effect on bi-di docs originating on Content Engine, because the content for those docs is stored on CE. In order to configure this, an Image Services Cache Server must be installed at the same location as Content Engine. It can either be installed on its own server, or on the same server that CE is installed on. This will improve performance for images that are retrieved multiple times. The first time you retrieve an image, it will take some time to come across the WAN, but it will then be stored on the cache server, from which all future retrievals of that image will come. For performance improvements on images that are retrieved once, or if you want to improve the performance of retrieving an image the first time, you will need to use prefetch to load the Cache server with the images you will be using. You also need to make sure the cache is large enough to hold all of the images to be retrieved locally.

## MS SQL Server

- If the CEMP database for the object store is SQL Server, and there is high outbound volume resulting in a large backlog in the replication queue, a retrieval key (index) can be added to increase performance and reduce the backlog.
- Use SQL Server tools to perform these SQL command on the object store database:  

```
CREATE UNIQUE INDEX RQDES ON ReplicationQueue (destination_id, exclusion_id, seqnum_identity)
```

The next topic is performance tuning for MS SQL Server. If the Content Engine database for the object store is SQL Server, and there is high outbound replication volume resulting in a large backlog in the replication queue, a retrieval key, also known as an index, can be added to increase performance and reduce the backlog. What you are going to want to do is use the standard SQL Server tools to execute the SQL command shown on this slide to create the index. Note that this index is only effective for SQL Server. Our testing has shown that there is no gain if you use it with DB2® or Oracle.

## Troubleshooting (1 of 3)

- Monitor system vitals - memory consumption, processor usage, and so on.
  - As system resources are maxed out, increasing maximum workers does not increase performance and might decrease performance
  - If CPU usage is very high, try changing busy wait intervals from 0 to 1
- Monitor usage of IS stubs and increase limits if necessary
  - Increasing maximum workers increases the number of IS stubs needed to handle requests
  - Check IS error logs for messages indicating limits have been reached

Now we will discuss some troubleshooting tips. You want to monitor the system vitals: memory consumption, CPU usage, disk I/O, etc. As system resources are maxed out, increasing maximum workers does not increase performance and might actually decrease performance. Instead, you will need to resolve the resource issue by upgrading your CPU, increasing memory, adding disk, etc. One thing you can do if the CPU usage on CE is very high, try changing the busy wait intervals from 0 to 1. We have found that multiple workers on a busy system can max out the CPU. Forcing them to wait a second between processing batches resolves this problem. You also need to monitor the usage of IS stubs and increase the limits if necessary. Increasing the maximum workers increases the number of IS stubs needed to handle requests. You are going to want to check the IS error logs for messages indicating that these limits have been reached and use the PPMOI tool to increase them accordingly.



## Troubleshooting (2 of 3)

- Monitor JVM memory and increase maximum heap space if necessary
  - Increasing maximum workers and batch size increases JVM memory consumption
- Monitor error logs and resolve issues promptly
  - Some errors are retried
  - Failed records are handled in a separate database transaction
  - If timeouts occur, reduce batch size or increase timeout settings or both

You are also going to want to monitor JVM memory and increase the maximum heap space if necessary. Increasing the maximum workers and the batch size is going to increase your JVM memory consumption, and you do not want to run out. The other thing you need to do is monitor the error logs and resolve any issues promptly. Some errors are retried and are eventually successful. It is more efficient to have it work the first time. Also, failed records are handled in a separate database transaction, impacting performance. If timeouts occur, you can reduce the batch size and you can also try increasing the setting for the timeout that occurred.

## Troubleshooting (3 of 3)

- Custom applications must close input streams after retrieving content from IS
  - It takes up to 30 seconds for Content Engine to time out and close the connection to Image Services. This keeps Cache Services processes connected to CE, unable to process other requests, but still using resources
  - Use cormon to detect this - look for a large number of CSMs stubs in RCVMSG state
  - Contact application developer, have them review code for InputStreams not being closed

Another source of performance issues is from custom applications retrieving content. Custom applications must close their input streams after retrieving the content of federated documents from Content Engine, where the content resides on Image Services. If they do not, it takes up to 30 seconds to CE to time out and close the connection to IS. This keeps Cache Services processes connected to CE, unable to process other requests, but still using resources. Use cormon to detect this - look for a large number of CSMs stubs in the RCVMSG state. If you find this, you will need to contact the appropriate application developer and have them review their code for InputStreams not being closed.

---

## Course summary

You have completed this course and can:

- Adjust the replication configuration parameters to maximize performance
- Identify the causes of performance problems and take corrective action

You have completed this course and can now adjust the replication configuration parameters to maximize performance, identify the causes of performance problems, and take corrective action.

---

## Product documentation

- FileNet P8 V5.1 documentation in the IBM Knowledge Center

[http://www-01.ibm.com/support/knowledgecenter/SSNW2F\\_5.1.0/com.ibm.p8toc.doc/ic-homepage.html](http://www-01.ibm.com/support/knowledgecenter/SSNW2F_5.1.0/com.ibm.p8toc.doc/ic-homepage.html)

The link on this slide will bring you to the documentation for P8 5.1, which contains information on performance tuning for bi-di.

---

## Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, DB2, and FileNet are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of other IBM trademarks is available on the web at "[Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml)" at <http://www.ibm.com/legal/copytrade.shtml>

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

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS OR SOFTWARE.

© Copyright International Business Machines Corporation 2012. All rights reserved.