

# IBM® Connections Version 5.5

## Performance Tuning Guide

IBM® Collaborations Solutions Performance Team  
June 2016  
Document Version 1.00



## Table of Contents

<b>IBM® Connections Version 5.5</b> .....	<b>1</b>
<b>Introduction</b> .....	<b>5</b>
About this document.....	5
Document History.....	5
<b>What is new in Connections 5.5</b> .....	<b>6</b>
<b>Performance Tuning Overview</b> .....	<b>7</b>
Environment Recommendations.....	7
64-bit.....	7
Hardware Multi-threading (Hyper-threading).....	7
Virtualization.....	7
IBM® Connections Topologies.....	8
Deployment Options.....	9
Clustered Deployments.....	10
IBM® Connections and Horizontal Clustering.....	10
IBM® Connections and Vertical Clustering.....	10
Network File Share Considerations.....	10
Database Server Hardware.....	11
Database Server Architecture.....	11
Content Caching and Compression Overview.....	11
HTTP Server.....	12
Configuring file downloads through the HTTP Server.....	13
Configuring file uploads through the HTTP Server.....	14
Cognos Business Intelligence Server.....	14
Other Considerations.....	15
Search.....	15
Disable Servlet Caching for Files Server.....	15
Additional Resources.....	16
Benchmark System Used for Performance Measurements.....	17
Benchmark Workloads.....	17
Benchmark Deployments.....	17
Connection Content Manager (CCM).....	17
Benchmark Hardware & Operating Systems.....	18
<b>Tuning the Connections System</b> .....	<b>19</b>
Tuning the Application Server Environment.....	19
Tuning via the Integrated Solutions Console.....	19
Tuning the JVM.....	19
Tuning the Datasources.....	23
Web Container Thread Pools.....	25
Service Integration Bus Message Stores.....	25
Tuning the Web Server.....	27
Process & Thread Settings.....	27
Enabling Caching in IHS.....	28
Content Compression Directives.....	29
Monitoring HTTP Server Activity.....	29
HTTP Logging / Problem Determination.....	29



<b>Tuning the Database – DB2®</b> .....	<b>31</b>
General Tuning.....	31
Autonomic Features.....	31
DB2_MINIMIZE_LISTPREFETCH.....	31
DB2_SELECTIVITY.....	32
DB2 Data Compression.....	32
Enable compression and compress.....	33
<b>Tuning the Database – Oracle®</b> .....	<b>34</b>
Buffer Cache and Shared Memory Management.....	34
Automatic Shared Memory Management (ASMM).....	34
Automatic Memory Management (AMM).....	35
Oracle® SQL Profiles.....	37
Optimising the Execution Plan Using SQL Profiles.....	37
High CPU and Disk utilisation for Homepage queries due to sub-optimal Oracle® Execution Plans.....	37
Effects of SQL Profile.....	38
Creating and Accepting a SQL Profile.....	38
Automatically Accepting SQL Profiles.....	39
Additional Notes.....	39
Licensing Requirements for SQL Profiles.....	39
Oracle® Securefile LOBS.....	41
Oracle® Initialisation Parameters.....	42
Processes.....	42
Sessions.....	42
Open Cursors.....	42
DB Writer Processes.....	42
64-bit Linux Specific.....	42
Virtual Memory.....	42
Shared Memory.....	45
<b>Tuning the Operating System</b> .....	<b>47</b>
Windows.....	47
Network Tuning.....	47
Linux.....	48
Swapping.....	48
Network Tuning on Linux.....	48
Setting Shell Limits for the Oracle® User on the Database System.....	49
Setting Shell Limits for the Connection User on the Application Server Tier.....	50
Other Environments.....	51
<b>Tuning IBM® Connections Applications</b> .....	<b>52</b>
Profiles Login Attributes.....	52
Tuning the Social Analytics Indexer.....	52
Scheduled Jobs.....	53
Social Analytics Widgets.....	54
Public Blogs Page Cache-Control.....	54
Tuning Activities Subject Cache.....	54
Tuning the PushNotifications server for FileSync and Notifications Center.....	59
<b>System Monitoring</b> .....	<b>63</b>
Overview.....	63
Areas to Monitor.....	63
Metrics.....	63
Systems.....	63



Virtualised Environments.....	64
What to look for.....	64
<u>Appendix A - Reference Material.....</u>	<u>66</u>
<u>Appendix B - Oracle® Initialisation Parameters Summary.....</u>	<u>68</u>
<u>Appendix C – DB2 Compress Homepage.....</u>	<u>69</u>
<u>Appendix D - Credits and Feedback.....</u>	<u>87</u>
Trademarks.....	88
Terms and Conditions.....	88

## Figures

<u>Figure 1: Connections Performance Benchmark System Diagram.....</u>	<u>8</u>
<u>Figure 1: Manually Accepting a SQL Profile.....</u>	<u>38</u>
<u>Figure 2: Enable Automatic Acceptance of SQL Profiles.....</u>	<u>39</u>
<u>Figure 3: List Management Pack Access.....</u>	<u>40</u>

## Index of Tables

<u>Table 1: Base IBM® Connections Version 5.5 Maximum Heap Sizes.....</u>	<u>21</u>
<u>Table 2: Data Sources Tuning.....</u>	<u>24</u>
<u>Table 3: IBM® HTTP Server Tuning Directives.....</u>	<u>27</u>
<u>Table 8: Network Tuning Windows.....</u>	<u>47</u>
<u>Table 9: Network Tuning Linux.....</u>	<u>48</u>
<u>Table 10: Setting Shell Limits for the Oracle® User.....</u>	<u>49</u>
<u>Table 11: Setting Shell Limits for the Connections User.....</u>	<u>50</u>
<u>Table 12: Oracle® Option A:Automatic Shared Memory Management (ASMM).....</u>	<u>68</u>
<u>Table 13: Oracle® Option B: Automatic Memory Management (AMM).....</u>	<u>68</u>
<u>Table 14: Oracle® Initialisation Parameters.....</u>	<u>68</u>



# Introduction

## About this document

This white paper provides recommendations for performance tuning of IBM® Connections 5.5 based on the experience of the IBM® Collaboration Solutions Performance team.

Tuning and the resulting system capacity can be affected by many factors, including the characteristics of the load on the system and the hardware used, including servers, disk subsystems, network topology. Therefore the objective of this paper is not necessarily to recommend specific values, but rather to make you aware of the parameters used in the IBM® Connections Version 5.5 performance benchmarks.

This document does, however, provide tuning values specific to the hardware, topology and workloads used for IBM® benchmarks. With this in mind, administrators can use these values as an initial reference subject before conducting a formal sizing and tuning exercise specific to an individual customer's environment.

## Document History

- Version 1.0: Initial version



## What is new in Connections 5.5

- For *Medium* deployments, applications are now re-arranged on five application server clusters: Infra, Util, Apps, CCM, and Push.
- For *Large* deployments, the following applications are now separated out in their own dedicated AppserverCluster: Common, PushNotifications, RichtextEditor, WidgetContainer
- New file uploads through IBM HTTP Server feature
- The install manager now configures IHS if pre-installed before installing connections
- PushNotifications for Notification Center and File Synch
- IBM Connections 5.5 moved support to Oracle 12c
- IBM Edge Proxy is now End Of Life and is restricted by its 32-bit architecture. Our measurements were conducted with server side disk caching and compression enabled on IHS. For large deployments we recommend IHS on Linux to take advantage of its 64-bit architecture
- For more information on what is new access this link:  
[http://www.ibm.com/support/knowledgecenter/SSYGQH\\_5.5.0/admin/overview/i\\_ovr\\_r\\_whats\\_new.dita](http://www.ibm.com/support/knowledgecenter/SSYGQH_5.5.0/admin/overview/i_ovr_r_whats_new.dita)



## Performance Tuning Overview

Tuning the IBM® Connections environment involves tuning and configuring the various systems and components of the environment. This chapter discusses some general concepts and details the specifics of the configuration used in IBM®'s measurement environments. These specifics are as follows:

- Configuring the application server and the resources defined for that application server
- Tuning the databases and database server
- Tuning the web server and/or proxy server
- Tuning the operating system and network
- Tuning the IBM® Connections services

When tuning specific systems, it is important to begin with a baseline and monitor performance metrics to determine if any parameters should be changed. When a change is made, another measurement should be made to determine the effectiveness of the change.

In addition to the tuning changes made in IBM®'s measurement environments, there are some additional tuning options available, which can improve performance in certain circumstances; these are discussed in the [Other Considerations](#) section.

## Environment Recommendations

Before installing IBM® Connections, you should consider how the environment will be used to achieve ideal performance.

### ***64-bit***

All measurements were run with 64-bit operating systems running 64-bit versions of WebSphere. This is the recommended configuration for production deployments.

### ***Hardware Multi-threading (Hyper-threading)***

Many modern processor architectures support hardware multi-threading. This is known as Hyper-Threading (HT) on Intel processors and Simultaneous Multi-threading (SMT) on Power-series processors. IBM®'s experience is that using hardware multi-threading provides an improvement in capacity in all of the scenarios and platforms IBM® has measured, so we would recommend its use on platforms where this is an option.

### ***Virtualization***

When virtualizing IBM® Connections, it is important to ensure that the virtual machines are allocated enough resources to perform optimally. To reach capacity on a virtual machine (VM), it may be necessary to ensure that the virtual resources map one-to-one with physical resources, especially CPU and memory. The memory used by Java-based applications, such as IBM® Connections, should not cause excessive disk swapping; otherwise, the application can suffer a significant performance degradation.

Running IBM® Connections on a VM whose host is over committed does not achieve optimal performance. Consider dedicating CPUs and memory to the IBM® Connections VMs. In addition, ensure that the network bandwidth to the host is sufficient for all VMs. Depending on requirements, the IBM® Connections VM may require a dedicated network interface on the host.



## IBM® Connections Topologies

IBM® Connections supports wide-ranging business requirements with a variety of deployment topologies.

Typical deployments use a three-tier configuration:

- Proxy and HTTP server(s)
- Application server(s)
- Database and directory servers(s)

The primary benefit of having a multi-tiered configuration is to avoid resource contention caused by multiple databases and applications residing on a single server. For example, if the database server shares a node with the application server, the combined resource contention would negatively impact the amount of throughput that can be achieved. On the other hand, a small deployment may have sufficiently small resource requirements that some of these servers could be deployed on a single node.

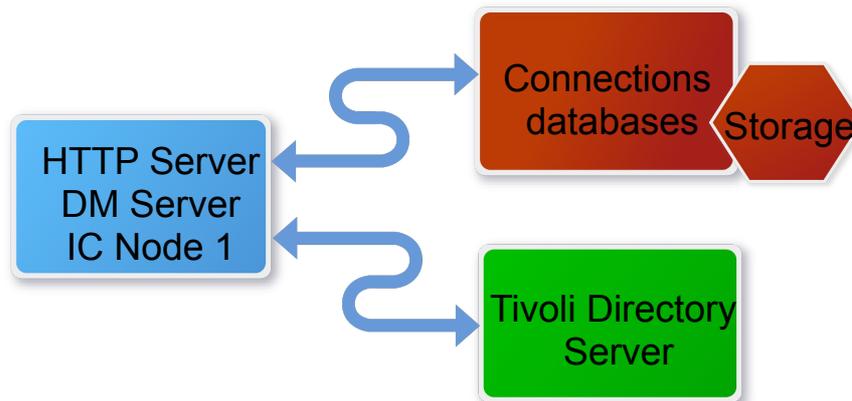
A multi-tiered configuration also allows a single component to be scaled up by adding extra clustered servers. IBM® Connections cluster server nodes, for instance can be added to increase capacity without also requiring a new database installation and configuration.

### Single-Server Topology

For smaller deployments, some of these tiers may be run on a single system. For example, a common configuration is to use a single node to run the HTTP server and the application server, while the database and directory servers are run on separate servers. This is the configuration we have used for most performance benchmarks on Windows and Linux. On newer, more powerful Windows systems an outboard 64-bit Linux based HTTP server is required as the Windows version of IHS is only 32-bit.

Figure 1 illustrates a typical system diagram used for the performance benchmark tests.

Figure 1: Connections Performance Benchmark System Diagram



## ***Deployment Options***

Choosing the most effective deployment topology requires evaluating security requirements, network infrastructure, plus considering management complexity and integration with other products. And in all cases, performance should be a consideration.

This section describes types of deployment topologies, and strategies for obtaining high performance when building the IBM® Connections environment.

Three deployment options are presented by the IBM® Installation Manager when installing IBM® Connections Version 5.5. Note that **Medium** and **Large** deployments offer considerably more scaling capabilities, and support multiple clusters.

- In the **Small** deployment option, all applications are run within a single Java Virtual Machine (JVM). This is typically used for low-volume deployments of IBM® Connections, such as proof-of-concept deployments.
- The **Medium** deployment has changed in IBM® Connections 5.5; applications have been re-balanced in 5 clusters instead of 4:
  - Infra: Common, Communities, ConnProxy, Help, News, Profiles, Search, URLPreview, Oauth, WidgetContainer
  - Util: Homepage, Moderation, RichTextEditors
  - Apps: Activities, Blogs, Bookmarks (Dogear), Files, Forums, Metrics, Mobile, Mobile Admin, Wikis
  - CCM: Extensions, FileNetEngine, Navigator
  - Push: PushNotifications
- The **Large** deployment option provides a predefined cluster name for each application. This option provides the best performance in terms of scalability and availability options, but also requires more system resources. In IBM® Connections 5.5 the following Applications are separated out in their own clusters: Common, PushNotification, RichtextEditor, WidgetContainer.

Note that this configuration was used for IBM®'s performance benchmarks.

**Note:** If you expect users to frequently download multiple big files in a single download session, we strongly recommend a *Large* deployment. See also the suggestion to [Disable Servlet Caching for Files Server](#)

For a more in-depth overview and deployment diagrams, see the IBM® Connections Version 5.5 wiki at: [http://www-01.ibm.com/support/knowledgecenter/SSYGQH\\_5.5.0/admin/plan/c\\_planning\\_the\\_installation.dita](http://www-01.ibm.com/support/knowledgecenter/SSYGQH_5.5.0/admin/plan/c_planning_the_installation.dita)

You can also review the "Other Considerations" section in this white paper.



## ***Clustered Deployments***

A clustered IBM® Connections deployment exploits the clustering capability of WebSphere Application Server to run applications on two or more cluster members – the JVMs, which run the applications. The cluster members can be deployed on the same node (server machine), or can be spread across multiple nodes. Running multiple cluster members on the same node is known as a *vertical clustering*, while running cluster members on separate nodes is known as a *horizontal cluster*.

## ***IBM® Connections and Horizontal Clustering***

Horizontal clustering has two main purposes: increasing capacity and increasing reliability. Since horizontal clusters run on multiple nodes, more resources are available, increasing the total number of users that can be supported by the cluster. A horizontal cluster provides greater reliability, as the failure of a single node does not cause a failure of the entire cluster.

Production deployment strategies are discussed in more detail in the Planning section of the IBM® Connections Version 5.5 product documentation, at:

[http://www-01.ibm.com/support/knowledgecenter/SSYGQH\\_5.5.0/admin/plan/c\\_installation\\_overview.dita](http://www-01.ibm.com/support/knowledgecenter/SSYGQH_5.5.0/admin/plan/c_installation_overview.dita)

All performance benchmarks were run on single-node environments.

**Key point:** Horizontal clustering of IBM® Connections provides increased capacity and reliability compared to single-node deployments.

## ***IBM® Connections and Vertical Clustering***

**Key point:** In IBM® Connections, vertical clustering is not required and not supported.

Vertical clustering is used for some Web applications where the application is not able to exploit all the system resources with a single application server instance. For these applications, several application servers are deployed on each node in the cluster. In the case of IBM® Connections, we recommend using a Medium or Large deployment model; these will deploy multiple application servers to each node, which will typically make good use of the resources on the host node. IBM® Connections does not support vertical clustering.

## ***Network File Share Considerations***

In a clustered environment, IBM® Connections uses a shared file system for storing file content, attachments in Activities, message stores for the WebSphere Application Server's Service Integration Bus (WAS SIB), and other content.

### **Notes:**

- The search index must be stored locally on each node where the Search application is installed.
- The shared file system used for IBM® Connections must support **advanced locking capabilities**:
  - When using NFS, use NFS v4 because NFS v3 lacks advanced locking capability.
  - Microsoft Distributed File System (DFS) lacks a distributed locking facility for files; thus, it cannot be used for IBM® Connections environments.

For more information see:

<http://www.ibm.com/support/docview.wss?uid=swg21614197>

[http://www-01.ibm.com/support/knowledgecenter/SSYGQH\\_5.5.0/admin/plan/c\\_planning\\_the\\_installation.dita](http://www-01.ibm.com/support/knowledgecenter/SSYGQH_5.5.0/admin/plan/c_planning_the_installation.dita)

A shared file system can be a performance bottleneck therefore pay attention to the network tuning, hardware and other areas to limit this risk. For example, in one performance deployment, the anti-virus software on the application server



was scanning the content on the mapped network drive; this is an extra level of precaution, but one imposing a performance cost.

All of the IBM® Connections Version 5.5 deployments used for the performance benchmarks used a remote file system, mapped over the network, for shared data. Although Small deployments do not strictly require the file content to be available outside the node, we chose for all IBM®'s benchmark systems to emulate this production environment strategy.

**Key point:** The network file share plays a critical role in performance and reliability of a deployment. Size carefully to ensure disk and network I/O is sufficient to meet the demand.

## ***Database Server Hardware***

IBM® Connections makes extensive use of its databases to provide its functionality – so the performance of those databases is crucial to achieve good performance for IBM® Connections. The database server(s) should have a significant amount of memory: at least 8GB on small deployments; as much as 128GB may be helpful on larger deployments. In our benchmark tests we used a minimum of 48GB. In addition, the performance of the disk subsystem where the content is stored is very important. This should be a dedicated storage device, such as an IBM® Storwize® V7000 device, or some other high-performance disk system capable to handle a high disk I/O throughput.

From IBM®'s observations, nearly 1/3 of the disk write operations are due to the databases' transactional log files. It is recommended that you locate these log files on a separate high performance disk subsystem, for example, based on Solid State Disks (SSD).

**Key point:** The database server(s) are a crucial component of an IBM® Connections deployment and the hardware used for them must be adequate for the workload.

## ***Database Server Architecture***

For the database server(s), it is recommended that you use 64-bit architecture exclusively, if possible. A 64-bit database server can access vastly more memory than a 32-bit database server. This provides a performance benefit for the database, as it is able to do greater in-memory caching, which reduces disk I/O and improves overall database performance.

With 64-bit database server(s), it is possible to deploy a single database instance per node, or one database instance per Database. IBM®'s experience with the IBM® Connections databases is that deploying each database in its own instance provides the best performance, at the cost of somewhat higher memory demand. Therefore, if sufficient physical memory is available, we recommend deploying each database in its own instance.

For IBM®'s performance benchmarks we used 64-bit hardware and one DB2 instance for each database.

**Key points:** If possible, deploy the databases used by IBM® Connections on 64-bit hardware, with each database deployed in its own database instance, if sufficient physical memory is available.

## ***Content Caching and Compression Overview***

IBM® Connections includes resources that can be cached and compressed: Implementing a caching solution reduces the impact on the Application Server, therefore increasing capacity. Implementing a compression solution reduces the network bandwidth usage improving users experience. Compression is particularly recommended if the site will be accessed by users with slow network links or who are a long distance from the application server (for example, users on a different continent from the application servers).

## ***Content Caching Details and Implementation Choices***

IBM® Connections includes resources that change infrequently, such as static page elements, JavaScript, style sheets, and images. Also cacheable are some public requests otherwise processed by the application server. For example, the



list of public Wikis can be cached for a short period of time; Caching them reduces the load on the application and database servers for a busy site. We recommend including a caching solution in any large-volume deployment.

There are two choices for caching: using a reverse proxy or enabling caching in the HTTP server. In this section we discuss both options.

**Note:** For our single server benchmark configurations we enabled caching on the HTTP server.

We suggest having the HTTP server reside on a different server than Connections for high workloads. Caching on the HTTP server is a good solution for this setup.

A reverse proxy can be used:

- as an alternative to an HTTP server to cache and compress public content
- to support a high number of concurrent “Long Poll” connections (see [Tuning the PushNotifications server for FileSync and Notifications Center](#)).

Other configurations are viable, but make sure that large cacheable resources are cached and compressed.

The advantage of using a reverse proxy over an HTTP server for caching depends on the topology used. In general it is best to have the caching done on a system other than the application server. If the HTTP server is on the same server as the application server, it is good to use a separate server as a caching reverse proxy.

We used IBM HTTP Server 8.5 in our measurement environments. Our Single-Server configurations have the web server running on the same system as the Connections application server. If, during your monitoring, you notice insufficient processor capacity on the system when running the web server and the Connections server co-located on a single system, consider separating the servers onto different systems.

**Key point:** Caching is recommended for increasing server performance and improving the client experience.

## Content Compression Details and Implementation Choices

IBM® Connections make use of a significant amount of readily-compressible content, such as JavaScript files and style sheets. Therefore we recommend compressing all content except images. This reduces the demand for network bandwidth and provides better performance for clients connecting over slower network.

IBM® Connections content can be compressed at either the HTTP server or the proxy server; it does not make sense to try to compress content at both locations. In IBM®'s lab measurements, we chose to compress the content on the HTTP server tier. See [http://www-01.ibm.com/support/knowledgecenter/SSYGQH\\_5.5.0/admin/install/t\\_ihs\\_config\\_not\\_compressing\\_files.dita](http://www-01.ibm.com/support/knowledgecenter/SSYGQH_5.5.0/admin/install/t_ihs_config_not_compressing_files.dita)

**Key point:** Caching and compressing content is recommended to reduce network bandwidth usage and for an improved client experience, especially over slow networks.

Information about configuring IBM® Connections with a caching proxy can be found in the IBM® Connections Version 5.5 product documentation at:

[http://www-01.ibm.com/support/knowledgecenter/SSYGQH\\_5.5.0/admin/install/t\\_install\\_deploy\\_caching\\_proxy.dita](http://www-01.ibm.com/support/knowledgecenter/SSYGQH_5.5.0/admin/install/t_install_deploy_caching_proxy.dita)

## HTTP Server

The HTTP Server is a mandatory element of IBM® Connections deployments, and is typically configured after IBM® Connections is installed. With IBM® Connections Version 5.5 the Install Manager can configure IHS if pre-installed with WebSphere before installing Connections. All of IBM®'s benchmark environments had the IBM® HTTP Server (IHS) co-located with the application server for ease of management.

IBM HTTP Server on Windows has a Parent process and a single multi-threaded Child process.



On 64-bit Windows OS'es, each instance of IHS is limited to approximately 2000 ThreadsPerChild.

[http://publib.boulder.ibm.com/httperv/ihsdiag/ihs\\_performance.html#tpc\\_win](http://publib.boulder.ibm.com/httperv/ihsdiag/ihs_performance.html#tpc_win)

With PI04922, it's possible to use nearly twice as much memory (nearly twice as many threads). See PI04922

<http://www-01.ibm.com/support/docview.wss?uid=swg1PI04922>

These numbers are not exact and depend on many factors. The upper limits become even more restrictive when loading other modules such as mod\_mem\_cache etc. No specific limits can be provided, but it is suggested that anything approaching the limits indicated above on a Windows operating system could be at risk.

Currently, there is no 64-bit version of IHS on Windows. So, if more threads are needed, a separate non-Windows server will be needed for IHS.

For more information see:

*“IBM WebSphere Application Server Performance Cookbook: Web servers”*

[http://publib.boulder.ibm.com/httperv/cookbook/Web\\_Servers.html](http://publib.boulder.ibm.com/httperv/cookbook/Web_Servers.html)

*“IBM HTTP Server Performance Tuning”*

[http://publib.boulder.ibm.com/httperv/ihsdiag/ihs\\_performance.html#ibm-http-server-performance-tuning](http://publib.boulder.ibm.com/httperv/ihsdiag/ihs_performance.html#ibm-http-server-performance-tuning)

*“Tuning IBM HTTP Server to maximise the number of client connections to WebSphere Application Server”*

<http://www-01.ibm.com/support/docview.wss?uid=swg21167658>

## ***Configuring file downloads through the HTTP Server***

This section is to stress the importance of the following mandatory post configuration steps.

IBM®'s benchmark environments were configured in this way.

**Key point:** You can make downloading files from Files, Wikis, Activities and Library applications more efficient by configuring an IBM® HTTP Server to handle most of the download process instead of the WebSphere® Application Server.

In the default deployment with an IBM® HTTP Server, file download requests are passed from the IBM® HTTP Server to the WebSphere Application Server. The WebSphere Application Server accesses the binary files in a data directory on the file system and returns them to the IBM® HTTP Server, which then passes them to the browser. This means that one thread on the application server is in use while the download is in progress.

When large numbers of users are downloading files, this deployment is inefficient, partly because WebSphere Application Server has a limited thread pool that is tuned for short-lived transactions. In addition, WebSphere Application Server is optimised for Java Platform, Enterprise Edition applications and not for file downloads. In this type of deployment, you might have to create a cluster to handle downloads, especially if you have slow transfer rates.

Configuring IBM HTTP Server to download files makes downloading much more efficient, because IBM HTTP Server is designed specifically for serving files. This configuration leaves WebSphere Application Server to carry out tasks such as security checking and cache validation.

If you choose not to configure the IBM® HTTP Server to download files, you must configure the WebSphere Application Server to transfer data synchronously instead of asynchronously to avoid errors related to using too much memory. For instructions, see the “Excessive native memory use in IBM® WebSphere Application Server” tech note at:

<http://www.ibm.com/support/docview.wss?uid=swg21317658>

Instructions for configuring file downloads through IBM® HTTP Server are provided in the IBM® Connections Version 5.5 Knowledge Center at:

[http://www-01.ibm.com/support/knowledgecenter/SSYGQH\\_5.5.0/admin/install/t\\_install\\_post\\_files\\_downloads.dita](http://www-01.ibm.com/support/knowledgecenter/SSYGQH_5.5.0/admin/install/t_install_post_files_downloads.dita)



## ***Configuring file uploads through the HTTP Server***

This is a new feature in IBM® Connections 5.5 and allows you to Configure the IBM® HTTP Server to manage file uploads from Activities and Files. As we explained for file downloads in the preceding section, this approach is more efficient than using the IBM WebSphere® Application Server to receive files.

Instructions for configuring file uploads through IBM® HTTP Server are provided in the IBM® Connections Version 5.5 Knowledge Center at:

[http://www-01.ibm.com/support/knowledgecenter/SSYGQH\\_5.5.0/admin/install/t\\_install\\_post\\_files\\_uploads.dita](http://www-01.ibm.com/support/knowledgecenter/SSYGQH_5.5.0/admin/install/t_install_post_files_uploads.dita)

## ***Cognos Business Intelligence Server***

IBM® Connections Version 5.0 introduced an enhanced Metrics application. Metrics employs the analytic capabilities of the IBM® Cognos® Business Intelligence server, which is provided as part of the IBM® Connections installation to support the collection of metrics data.

See *Setting up IBM Cognos Metrics to work with IBM Connections* :

[http://www-01.ibm.com/support/knowledgecenter/SSYGQH\\_5.5.0/admin/install/c\\_configure\\_Cognos\\_later.dita](http://www-01.ibm.com/support/knowledgecenter/SSYGQH_5.5.0/admin/install/c_configure_Cognos_later.dita)

Depending on the size of the deployment, the Metrics application can be in same server node as other IBM® Connections application, or consider dedicating an individual node to the Metrics application. For high-volume deployments, we recommend deploying the Cognos Business Intelligence server and the database server on nodes separate from the IBM® Connections application server(s) and database server.

Tuning recommendations for Cognos BI server is not covered in this document. Some references which provide useful information about Cognos BI performance tuning:

- developerWorks: *IBM Cognos 10 BI Community*:  
<https://www.ibm.com/developerworks/mydeveloperworks/groups/service/html/communityview?communityUuid=8d7e4a2b-2364-4719-9f4e-aa9e24db7465>
- developerWorks Technical Library: *IBM Cognos Proven Practices*  
[http://www.ibm.com/developerworks/views/data/libraryview.jsp?site\\_id=1&contentarea\\_by=AllZones&sort\\_by=Date&sort\\_order=2&start=1&end=100&topic\\_by=-1&product\\_by=-1&type\\_by=AllTypes&show\\_abstract=true&search\\_by=IBM Cognos Proven Practices&industry\\_by=-1&series\\_title\\_by=](http://www.ibm.com/developerworks/views/data/libraryview.jsp?site_id=1&contentarea_by=AllZones&sort_by=Date&sort_order=2&start=1&end=100&topic_by=-1&product_by=-1&type_by=AllTypes&show_abstract=true&search_by=IBM+Cognos+Proven+Practices&industry_by=-1&series_title_by=)
- *Performance Tuning for IBM Cognos 10*  
<http://www-01.ibm.com/support/docview.wss?uid=swg21618119>



## Other Considerations

### Search

The Search application is responsible for indexing all the searchable content in IBM® Connections and handling search requests. The search application also powers the social analytics widgets, such as *Recommendations*, *Do You Know*, *Things in Common*, and *Who Connects Us*. Depending on the size of the deployment (*Medium* or *Large* options), consider dedicating an individual node to the Search application.

See “*Making search-related configuration changes to newly added nodes*”

[http://www.ibm.com/support/knowledgecenter/SSYGOH\\_5.5.0/admin/admin/t\\_admin\\_search\\_add\\_search\\_node.dita](http://www.ibm.com/support/knowledgecenter/SSYGOH_5.5.0/admin/admin/t_admin_search_add_search_node.dita)

### Disable Servlet Caching for Files Server

If you expect users to frequently download multiple big files in a single download session, we strongly recommend a *Large* deployment and the following:

- Disable Servlet caching on the Files application server cluster
- Monitor Java Heap usage and size accordingly

**Note:** These instructions are applicable to the “Large Deployment” option where the Files application is deployed to its own application server cluster.

**How to Set:** In the WebSphere Integrated Solutions Console, click:

**Servers** → **Server Types** → **WebSphere application servers** → <**FilesCluster\_server1**> → **Container Settings: Web Container Settings** → **Web container** → **Un-check “Enable servlet and command caching”**

The screenshot shows the 'Web container' configuration page in the WebSphere Integrated Solutions Console. The breadcrumb trail is 'Application servers > FilesCluster\_server1 > Web container'. The page title is 'Web container' and it includes the instruction 'Use this page to configure the web container.' The 'Configuration' tab is active. Under 'General Properties', the 'Default virtual host' is set to 'default\_host'. The 'Enable servlet and command caching' checkbox is unselected and highlighted with a red box. Below it, the 'Disable servlet request and response pooling' checkbox is also unselected. Under 'Asynchronous Servlet Properties', the 'Number of timeout threads' is set to 2. The 'Set timeout' checkbox is checked, with a value of 30000 milliseconds. The 'Use a work manager to start Runnable objects' radio button is selected. The 'Work manager JNDI name' is set to 'wm/default'. At the bottom, there are buttons for 'Apply', 'OK', 'Reset', and 'Cancel'. On the right side, there are links for 'Additional Properties': 'Asynchronous Request Dispatching', 'Custom properties', 'Web container transport chains', and 'Session management'.



## Additional Resources

A comprehensive guide to all possible deployment strategies is beyond the scope of this paper. Some additional information sources, which may be useful are:

- WebSphere Application Server Performance Cookbook  
<http://publib.boulder.ibm.com/htpserv/cookbook/>
- IBM® Connections Version 5.5 Product Documentation:  
[http://www.lotus.com/ldd/lcwiki.nsf/xpViewCategories.xsp?lookupName=Product Documentation](http://www.lotus.com/ldd/lcwiki.nsf/xpViewCategories.xsp?lookupName=Product+Documentation).
- WebSphere Application Server Network Deployment Version 8.5 documentation: [http://www-01.ibm.com/support/knowledgecenter/SSAW57\\_8.5.5/com.ibm.websphere.nd.multiplatform.doc/ae/welcome\\_ndmp.html](http://www-01.ibm.com/support/knowledgecenter/SSAW57_8.5.5/com.ibm.websphere.nd.multiplatform.doc/ae/welcome_ndmp.html).
- DB2 10.5 Information Center: [http://www-01.ibm.com/support/knowledgecenter/SSEPGG\\_10.5.0/com.ibm.db2.luw.kc.doc/welcome.html](http://www-01.ibm.com/support/knowledgecenter/SSEPGG_10.5.0/com.ibm.db2.luw.kc.doc/welcome.html).
- Tuning IBM® System x Servers for Performance: <http://www.redbooks.ibm.com/abstracts/sg245287.html>



## Benchmark System Used for Performance Measurements

This section gives you a context of how we derived IBM®'s tuning recommendations.

### ***Benchmark Workloads***

The teams that evaluate IBM® Connections Version 5.5 performance run two main types of measurement workloads: *Integrated* and *Single* application

The *Integrated* workload is the main reference benchmark measurement, the results of which are published in the IBM® Connections Version 5.5 Benchmark Report. In this workload, all the components are exercised simultaneously as in a real world scenario.

The purpose of the *Single* application measurement is to evaluate the performance of the specific IBM® Connections Version 5.5 application, in relative isolation, outside of interaction with the other suite of offerings (although News and Search are always interacting with all components).

The metrics gathered from these measurements are also used for the IBM® Connections Version 5.5 sizing tool.

### ***Benchmark Deployments***

For all measurements (Integrated and Single components), the performance teams used the *Large* deployment option with each application in its own JVM.

All of the components offered in IBM® Connections Version 5.5 were installed and populated - much in the same manner as a production environment. Each application runs in its own JVM, each application database was installed in its own DB2 instance, and each database was stored in its own “logical drive” to allow us to monitor disk activity in greater detail but that's not needed for production deployments. The logical drives were located on a RAID disk subsystem. The IBM® HTTP was also configured to cache, compress and handle files downloading. For more information, see the [HTTP Server](#) section of this white paper.

### ***Connection Content Manager (CCM)***

CCM was introduced with IBM® Connections Version 4.5. There are two deployment options for CCM, described in the product documentation at [http://www-01.ibm.com/support/knowledgecenter/SSYGQH\\_5.5.0/admin/install/t\\_install\\_cluster.dita](http://www-01.ibm.com/support/knowledgecenter/SSYGQH_5.5.0/admin/install/t_install_cluster.dita) see step #12.

For IBM®'s workloads tests, we chose the following deployment option:

“*Configure Connections Content Manager ..... Install Now (recommended).*”

Here are the key aspect of the CCM deployment we used:

- This install option deploys CCM and FileNet together in their own dedicated application server cluster (e.g. CCMCluster\_server1) in the same cell as the rest of the IBM® Connections components.
- Two new additional databases are created for CCM:
  - FNGCD – FileNet configuration parameters
  - FNOS – Object Store for files metadata
- CCM stores the library files on a shared file system in the `shared_data_directory_root\libraries` directory (e.g. \\fileserv\Connections\data\shared\ccm)



## ***Benchmark Hardware & Operating Systems***

The servers running the measurement use 64-bit hardware, Microsoft Windows Server 2008 R2 Enterprise SP1 64-bit operating system and Red Hat Enterprise Linux Server release 5.11 (Tikanga) 64-bit.

The typical environment used for the benchmark workloads is as follows:

- 1 WebSphere Application Server node – typically an 8 or 12 cores server with 48GB of RAM running the IBM® Connections Version 5.5 product in 19 clusters (JVMs) with the IBM® HTTP Server co-located. Also co-located on this node is the WebSphere Deployment Manager (DM) serving to administer the environment. Note that the DM has minimal impact on the system.
- 1 database server – typically an 8 or 12 cores with 48GB of RAM server running IBM DB2® software.
- 1 LDAP server, running Tivoli® Directory Server – processor load on this tier is minimal.
- A mapped network file-system hosted by the DB2 server. Note that in a large production environment, a dedicated file server is recommended.
- For the DB2 server and the mapped network file-system, we used a high-performing disk subsystem.



## Tuning the Connections System

In general, the following observations about IBM® Connections Version 5.5 can be made:

- Performance across all the components is typically limited by disk I/O, either on the database, the file repository, or both. This is especially true for the News application accessing the Homepage database. Performance can be improved by adding RAM on the DB server to cache more data. From our observations, nearly 1/3 of the disk write operations are due to the databases' transactional log files. We suggest locating these log files on a separate high performance disk subsystem; for example, based on Solid State Disks (SSD).
- End-user experience can be affected by external factors, such as network latency, network bandwidth, and browser rendering times. Configuring a forward proxy server, to cache content near the user, improves client-side performance for public content.
- Compressing content, as explained earlier in this document, provides a significant improvement in terms of network bandwidth usage and end user overall response time.
- Caching content with a reverse caching proxy the HTTP Server near the IBM® Connections server, reduces the workload on the application server and significantly improves throughput.
- With the recommended heap sizes provided in the following section, none of the applications were memory constrained within the JVM heap in IBM®'s integrated benchmarks.

The following sections discuss tuning the various servers in a typical IBM® Connections Version 5.5 deployment.

### Tuning the Application Server Environment

Because IBM® Connections is composed of standard Java web applications, typical tuning practices for WebSphere applications apply. At a minimum, database connection pools, the application server thread pools, and JVM heap should be sized and monitored.

#### *Tuning via the Integrated Solutions Console*

To get to the WebSphere Application Server (WAS) Integrated Solutions Console, start WebSphere and then log in to the WebSphere Integrated Solutions Console via <http://<yourserver>:9060/ibm/console> with the administrator user ID created during the WAS installation.

The port number, 9060, is the default port number in IBM®'s lab deployments, but other deployments may use different ports. To find out the ports in use for your installation, look for 'adminhost' in the following path:

```
<was_profile_root>/config/cells/<cell_name>/nodes/<node_name>/serverindex.xml.
```

For more details on configuring a WebSphere Application Server, see the Tuning section of the information center located at: <http://www.ibm.com/software/webservers/appserv/was/library/>

#### *Tuning the JVM*

**Reference material:** IBM® WebSphere Application Server Information Center “*Network Deployment > Tuning performance > Tuning the full profile > Tuning the application serving environment > Tuning the JVM*” in: [http://www-01.ibm.com/support/knowledgecenter/SSAW57\\_8.5.5/com.ibm.websphere.nd.multiplatform.doc/ae/container\\_tune\\_jvm.html](http://www-01.ibm.com/support/knowledgecenter/SSAW57_8.5.5/com.ibm.websphere.nd.multiplatform.doc/ae/container_tune_jvm.html)



## Heap size

When setting the heap size for an application server, keep the following in mind:

- Size the physical memory of server to fit all of the processes into memory, plus enough for the operating system. When more memory is allocated than the available physical memory on the system, paging occurs and can result in very poor performance.
- The total JVM process size is made up by the Java heap, which you set, and the Java native heap, which grows to its working size after the system has “warmed-up” to its nominal load. For more information, see the developerWorks web site at:  
<http://www.ibm.com/developerworks/linux/library/j-nativememory-linux/index.html>
- We set the minimum and maximum heap sizes to the same values since we are using the generational, concurrent (or ‘gencon’) garbage collection, which helps avoid heap fragmentation. Generational concurrent garbage collection has given the best throughput and response time results in IBM®'s measurements. This is the default setting when installing IBM® Connections.
- Note that running clustered WebSphere nodes may require larger heap sizes compared to stand-alone, non-clustered systems if session replication is being used. This is especially true if using memory-to-memory session replication since session information is also stored in the JVM.
- After doing any heap size tuning, do the following:
  - Monitor the system to make sure that paging is not occurring.
  - Monitor the verbose garbage collection output to determine if the selected size is appropriate. Ideally, the system should spend no more than 3~5% of its time in garbage collection.

By default, when installing IBM® Connections Version 5.5, the IBM® Installation Manager sets the initial heap size to 512MB and the maximum heap size to 2560MB for all deployment options: *Small* or *Medium and Large*

**Key Point:** These values must be reviewed for production systems based on sizing guidelines and then fine-tuned based on usage statistics.

Based on IBM®'s latest integrated (i.e. a mix of all components) and single workload tests, we have revised IBM®'s recommendation for initial Java heap tuning.

We have now grouped the min/max Java Heap size recommendations in three categories:

Small (1.5GB), Medium (2GB), Large (2.5GB).

Small: Bookmarks, Homepage, Common, WidgetContainer, RTE

Medium: Activities, Blogs, Files, Forum, Profiles, Wikis, Mobile

Large: Communities, News, Search, Metrics, CCM,

Special case: PushNotification - see notes below

**Note:** More content will drive a requirement for a bigger heap in *Search* therefore we recommend you monitor the heap usage and adjust accordingly.

On IBM®'s benchmark systems, we saw good performance using the heap sizes given in Table 1.



*Table 1: Base IBM® Connections Version 5.5 Maximum Heap Sizes*

<b>JVM Minimum and Maximum Heap Sizes</b>	
<b>Application</b>	<b>GB</b>
Activities	2
Blogs	2
Bookmarks	1.5
CCM	2.5
Common	1.5
Communities	2.5
Files	2
Forums	2
Homepage	1.5
Metrics	2.5
Mobile	2
Moderation	0.5
News	2.5
Profiles	2
PushNotifications	See Note*
RTE	1.5
Search	2.5~4GB**
WidgetContainer	1.5
Wikis	2

\***Note:** Refer to this chapter [Tuning the PushNotifications server for FileSync and Notifications Center](#)

\*\***Search:** Depending on your data population you may have to increase the heap size to build a full index from scratch. We required a Heap of 4GB to build an index of this size:

Size on disk: 33GB  
No of files: 406  
No of Fields: 526,843  
No of Docs: 7,341,075  
No of Terms: 61,935,171

**How to Set:** In the WebSphere Integrated Solutions Console, click:

**Servers** → **Server Types** → **WebSphere application servers** → **<application\_clusterServer>** → **Server Infrastructure: Java and Process Management** → **Process Definition** → **Java Virtual Machine**

- Initial heap size
- Maximum heap size

Regardless of deployment topology, we recommend enabling verbose garbage collection logging and monitoring heap use for all JVMs. Verbose GC logging has a minimal impact on performance and provides valuable performance diagnostic information.



*Servers* → *Server Types* → *WebSphere application servers* → *<application\_clusterServer>* → *Server Infrastructure: Java and Process Management* → *Process Definition* → *Java Virtual Machine*

- Check the *Verbose garbage collection* check-box.

**Key points:** Monitor verbose garbage collection logs and make sure sufficient heap space is available. When setting the maximum JVM heap size, pay particular attention to ensure sufficient physical memory exists to support the value chosen.

## Nursery Size

Benchmarks show that the generational garbage collector gives the best performance; therefore, it is enabled by default since IBM® Connections, version 3.0. The string `-Xgcpolicy:gencon` is present in the Generic JVM Arguments field in the WAS Admin Console.

The generational garbage collector splits the heap into two sections: the nursery and the tenured region. The collector can be fine-tuned by setting the size of the space where new objects are allocated, the nursery. A large nursery is typically good for throughput, while a small nursery is typically good for low pause times. Moreover, good WebSphere Application Server performance (throughput) usually requires a reasonably large nursery. A good starting point would be to size the nursery to ¼ of the Max Heap size e.g. 512 MB for a 2GB heap size. Next, move up or down to determine optimal value, measure throughput or response times, and analyse GC logs to understand frequency and length of scavenges. In particular watch the “flipped bytes” count and keep that as low as possible.

The nursery size is set with the JVM argument `-Xmn` (for example, `-Xmn512M`)



## *Tuning the Datasources*

### **Connection Pool Size**

By default, when installing IBM® Connections Version 5.5, the IBM® Installation Manager sets the minimum connections pool size to 1 and the maximum connections pool size to 25. These are good values to start with, you should then monitor their usage and adjust to your workload requirements. See [Table 2: Data Sources Tuning](#) used for our tests. Note that each application has its own connection pool.

**How to Set:** In the WebSphere Integrated Solutions Console, click:

**Resources** → **JDBC Providers** → *<provider name>* → **Data Sources** → *<data source name>* → **Connection pool properties**

- Maximum connections
- Minimum connections

### **Prepared Statement Cache Size**

All data sources are configured in a similar manner. By default, IBM® Connections components install with a prepared statement cache value of 100. The default prepared statement cache size is optimum for DB2. However, deployments with a different RDBMS, such as Oracle®, which use larger prepared statements, may want to reduce the size of the cache. In the integrated benchmark using Oracle®, the prepared statement cache size was reduced to 50.

**How to Set:** In the WebSphere Integrated Solutions Console, click:

**Resources** → **JDBC Providers** → *<provider name>* → **Data Sources** → *<data source name>* → **WebSphere Application Server data source properties** → **Statement cache size**.

Be aware that specifying a larger prepared statement cache size can lead to OutOfMemory errors in situations where your application memory is already being highly used by your workload. The prepared statement cache size setting is the maximum allowed cache entries per database connection. So increasing the cache size on a data source that has a large number of connections can quickly increase the heap usage for these cache objects. Any changes should be considered for each individual data source independently instead of across all data sources globally. Before increasing a data source's prepared statement cache size, you should monitor your memory usage under a heavy workload to determine if there is enough JVM heap available to handle an additional increase.



Table 2 summarizes sizes used in IBM®'s benchmarks:

*Table 2: Data Sources Tuning*

Data Source	DB2		Oracle®	
	JDBC Pool (min/max)	Statement Cache Size	JDBC Pool (min/max)	Statement Cache Size
activities	1/50	100	100/100	50
blogs	1/250	100	50/100	50
dogear (Bookmarks)	1/150	100	50/100	50
CCM FNOSDS	75/200	100	75/200	50
CCM FNOSDSXA	25/75	100	25/75	50
CCM FNGCDDS	default	default	default	default
CCM FNGCDDSXA	default	default	default	default
communities	10/200	100	50/100	50
files	10/100	100	50/100	50
forums	50/100	100	50/100	50
homepage	20/100	100	50/100	50
metrics	1/75	100	50/100	50
mobile	1/100	100	50/100	50
news	50/75	100	50/100	50
oauth_provider	1/50	100	01/50	50
profiles	1/100	100	50/100	50
pushnotification	25/75	100	25/75	50
search	50/75	100	50/100	50
urlpreview	50/75	100	50/75	50
widgetcontainer	50/75	100	50/75	50
wikis	1/100	100	50/100	50

It is always recommended to periodically evaluate connection pool sizes and requirements in the actual environment using the IBM® Tivoli® Performance Monitor, integral to WebSphere 8, or other tools.

**Key points:** Increase the JDBC connection pool sizes from the default values and monitor their usage. The prepared statement cache values should not need to be changed if using DB2 as the back-end RDBMS.



## ***Web Container Thread Pools***

By default, when installing IBM® Connections Version 5.5, the IBM® Installation Manager sets the minimum Web Container Thread Pool size to 50 and the maximum to 75. These are good values to start with then you should monitor their usage and adjust to your workload requirements. For a high-volume site, we recommend increasing the WebContainer thread pool to 100 in all application servers to allow more simultaneous requests.

**How to Set:** In the WebSphere Integrated Solutions Console, click:

**Servers → Server Types → WebSphere application servers → <ActivitiesCluster\_server1> → Additional Properties: Thread Pools → Web Container → Thread Pool**

- Minimum size threads
- Maximum size threads

This needs to be repeated for all application servers.

**Key point:** Users should periodically monitor the actual number of threads being used in the pool by enabling the appropriate performance counters in the IBM® Tivoli Performance Monitor.

## ***Service Integration Bus Message Stores***

For high-volume sites, we recommend to tune the Service Integration Bus (SIB) Message File Store. For example, for the Activities application, the following tuning worked well in IBM®'s workloads:

- **Message Store** - High loads on the Activities application generate a high messages traffic, which needs a larger message store for the Activities Messaging Engine (ME).

We noticed good performance using the following values:

- Log size = 500 MB
- Minimum permanent store size = 500 MB
- Maximum permanent store size = 3000 MB
- Minimum temporary store size = 500 MB
- Maximum temporary store size = 1000 MB

Other ME candidates for the same tuning based on traffic: News, Communities, Files

**How to Set:** In the WebSphere Integrated Solutions Console, click:

**Service integration -> Buses > ConnectionsBus > Messaging engines > ActivitiesCluster.000-ConnectionsBus > Message Store**

- **Message Thresholds** - From IBM®'s performance measurements, we observed a significant message traffic, which required us to increase the size of the News Connections bus.

You should monitor the typical queue sizes and growth patterns and tune accordingly.  
We saw good performance using the following values:

- Default messaging engine high message threshold = 300,000
- High message threshold per message point = 300,000



**How to Set:** In the WebSphere Integrated Solutions Console, click:

**Service integration -> Buses -> ConnectionsBus  
Default messaging engine high message threshold = 300,000**

**Service integration -> Buses -> ConnectionsBus -> Messaging engines-> < News cluster ME>,  
High message threshold per message point = 300,000**

**Service integration -> Buses -> ConnectionsBus -> Messaging engines-> < News cluster ME> ->  
Publication points -> connections.events@NewsCluster.000-ConnectionsBus  
High message threshold = 300,000**



## Tuning the Web Server

All performance benchmark measurements were conducted with the IBM® HTTP Server co-located on the node running the WebSphere Application Server. IBM® HTTP Server *fast file downloading* option was configured for: Activities, CCM, Files, Mobile and Wikis. This permits the IHS server to download binary content directly, instead of passing the download requests to the application server, as we found this increased server capacity.

The IBM® HTTP Server provided with IBM® WebSphere Application Server ND requires very little tuning, as most of the default settings are set for excellent performance. The following performance changes were made, in addition to configuring the HTTP Server for SSL:

### Process & Thread Settings

Table 3: IBM® HTTP Server Tuning Directives

Parameter	Windows	Linux
<b>KeepAliveTimeout</b>	5	5
This value is less than the think time defined in IBM®'s scripts to ensure that testing is conservative. Each user is assumed to open a new TCP connection for each page view. However, in a live environment, it can be helpful to increase the KeepAliveTimeout. A higher timeout value can increase contention for HTTP server processes. If you are running out of HTTP processes, decrease this value.		
<b>ThreadLimit</b>	1500 (Same as TreadsPerChild)	25
<b>ThreadsPerChild</b>	1500	25
The higher number of threads per child on Windows is due to a different process model for IHS on Windows.		
<b>MaxRequestsPerChild</b>	0	0
<b>MaxKeepAliveRequests</b>	0	0
Selecting 0 lets an unlimited number of HTTP requests calls to a single TCP connection. This allows maximum reuse of connections between a proxy and the HTTP server.		
<b>StartServers</b>	N/A	2
<b>Access logging</b>	On	On
Access logging was enabled. This setting has little impact on performance and most users choose to enable some logging policy on their HTTP server(s).		
<b>ServerLimit</b>	N/A	80
This value should be set to $\text{MaxClients} \div \text{ThreadsPerChild}$ .		
<b>MinSpareThreads</b>	N/A	25
<b>MaxSpareThreads</b>	N/A	2000
This value should be set to the same values as MaxClients.		
<b>MaxClients</b>	N/A	2000

For Linux, AIX and z/OS, these values assume that IHS is using the default Worker Multi-Processing Module. This module creates multiple processes (servers) to handle requests. Each server, in turn, can create multiple threads. HTTP requests are handled by a thread. So, the number of threads determines how many concurrent connections the server can



service. In the settings in Table 3, notice that the ThreadLimit setting multiplied by the ServerLimit is equal to the MaxClient value. If MaxClients is set lower, each server will not be able to start as many threads as are set in ThreadLimit.

To increase MaxClients, either the ServerLimit, ThreadLimit, or both must also be increased. Note that UNIX-based HTTP servers cache information in each process. So, when increasing settings, raise the number of threads first before raising the number of processes. This reduces the overall memory needed by the HTTP server.

The values used in the performance benchmarks are set to ensure there are sufficient resources to handle hundreds of hits per second from thousands of users. For smaller environments, lower values that use less server resources can be used. For example, to serve several hundred hits per second, ServerLimit could be set to 5, ThreadLimit could be set to 160 and MaxClients could be set to 800.

IBM HTTP Server on Windows has a Parent process and a single multi-threaded Child process. On 64-bit Windows OS'es, each instance of IHS is limited to approximately 2000 ThreadsPerChild.

[http://publib.boulder.ibm.com/httperv/ihsdiag/ihs\\_performance.html#tpc\\_win](http://publib.boulder.ibm.com/httperv/ihsdiag/ihs_performance.html#tpc_win)

With PI04922, it's possible to use nearly twice as much memory (nearly twice as many threads). See PI04922 <http://www-01.ibm.com/support/docview.wss?uid=swg1PI04922>

These numbers are not exact and depend on many factors. The upper limits become even more restrictive when loading other modules such as mod\_mem\_cache etc. No specific limits can be provided, but it is suggested that anything approaching the limits indicated above on a Windows operating system could be at risk.

Currently, there is no 64-bit version of IHS on Windows. So, if more threads are needed, a separate non-Windows server will be needed for IHS.

For more information see:

*"IBM WebSphere Application Server Performance Cookbook: Web servers"*

[http://publib.boulder.ibm.com/httperv/cookbook/Web\\_Servers.html](http://publib.boulder.ibm.com/httperv/cookbook/Web_Servers.html)

*"IBM HTTP Server Performance Tuning"*

[http://publib.boulder.ibm.com/httperv/ihsdiag/ihs\\_performance.html#ibm-http-server-performance-tuning](http://publib.boulder.ibm.com/httperv/ihsdiag/ihs_performance.html#ibm-http-server-performance-tuning)

*"Tuning IBM HTTP Server to maximize the number of client connections to WebSphere Application Server"*

<http://www-01.ibm.com/support/docview.wss?uid=swg21167658>

**Key point:** The IBM® HTTP Server should need very little tuning, but make sure to configure File Serving if appropriate for the deployment, and enable compression for file types if not doing so on the proxy server.

## Enabling Caching in IHS

For the IBM® Connections Version 5.5 performance evaluation, disk caching was used as in-memory caching is deprecated. With proper setup of the operating system's file system buffering, disk caching achieves results similar to in-memory caching. Disk caching implicitly uses system memory as file system cache which can be swapped out if needed.

### How to Set Disk Caching

These values are set in the HTTP server's httpd.conf file:

```
LoadModule cache_module modules/mod_cache.so
<IfModule mod_cache.c>
    # Used in conjunction with %{CACHE_MISS}e in LogFormat we can track cached
and not cached requests
    SetEnv CACHE_MISS 1
    # Disable IHS caching of server status
    CacheDisable /server-status
```



```
LoadModule disk_cache_module modules/mod_disk_cache.so
<IfModule mod_disk_cache.c>
    CacheRoot E:/ihscache
    CacheEnable disk /
    CacheDirLevels 2
    CacheDirLength 1
    CacheIgnoreHeaders Set-Cookie
    CacheMaxFileSize 10000000
    CacheIgnoreNoLastMod On
    CacheDefaultExpire 86400
</IfModule>
</IfModule>
```

On Linux make sure that the the HTTP server has permission to write to the location specified by CacheRoot by running the command: `chown nobody /ihscache`

Make sure to select either disk or in memory caching, but not both. If a caching reverse proxy is used, in most cases there is no need for caching on the HTTP server as well.

## Performance Considerations

For best performance, it may be necessary to place the disk cache on a separate disk. This not only helps IO throughput, but also allows changing mount options on the cache disk. For best performance on Linux, set the `noatime` option on the filesystem to prevent writing access times on each file read.

More information on configuring disk caching can be found in this blog post on IBM developerWorks® [https://www.ibm.com/developerworks/mydeveloperworks/blogs/portalops/entry/apache\\_mod\\_mem\\_cache\\_versus\\_mod\\_disk\\_cache\\_choose\\_mod\\_disk\\_cache20?lang=en](https://www.ibm.com/developerworks/mydeveloperworks/blogs/portalops/entry/apache_mod_mem_cache_versus_mod_disk_cache_choose_mod_disk_cache20?lang=en).

## Content Compression Directives

To reduce network traffic, GZIP compression was enabled on the HTTP server for all benchmarks.

Follow the steps documented in the IBM® Knowledge Center:

[http://www-01.ibm.com/support/knowledgecenter/SSYGQH\\_5.5.0/admin/install/t\\_ihs\\_config\\_not\\_compressing\\_files.dita](http://www-01.ibm.com/support/knowledgecenter/SSYGQH_5.5.0/admin/install/t_ihs_config_not_compressing_files.dita)

If a proxy server is used then we recommend to enable compression on the Proxy server instead.

**Note:** Do not compress content at both the HTTP server and the proxy server.

## Monitoring HTTP Server Activity

Users are advised to enable monitoring and adjust the number of workers, if needed.

We enabled the server-status module (`mod_status.so`) so that the number of running and available Web server processes could be monitored under load. This enables appropriate tuning of the parameters shown in [Table 3](#).

## HTTP Logging / Problem Determination

Access logs are crucial to problem determination. We recommend you enable logging of requests at one tier within the Connections deployment, such as at the proxy or HTTP server. In our performance measurements, we enabled this on the HTTP server.

We found it useful to add the response times to the access logs. This can be done by adding the `%D` token (for response



times) to the combined log format, 'combined' log format, as in this example:

```
LogFormat "%h %l %u %t \"%r\" %>s %b RESPONSE_TIME: [%D microseconds] \"%{Referer}i\" \"%{User-Agent}i\" combined
```

```
CustomLog logs/access.log combined
```

For further debugging, we found it useful to add some additional details, as in this example:

```
LogFormat "%h:%p %l %u %t RESPONSE_CODE:[%>s] AGE: [%{Age}o] CACHE_MISS: [%{CACHE_MISS}e] CACHE-STATUS: [%{cache-status}e] SIZE: [%b] RESPONSE_TIME: [%T/%D] (microseconds) WAS: [%{WAS}e] \"%r\" %>s REFERER: \"%{Referer}i\" USER_AGENT: \"%{User-Agent}i\" combined
```

References:

<http://www.ibm.com/developerworks/library/co-websphere-access-feature/>

<http://www-01.ibm.com/support/docview.wss?uid=swg21570830>

**Note:** Make sure you manage logs size and logs rotation to avoid running out of space on the disk where you are writing the logs too.

Example:

```
CustomLog "|C:/IBM/HTTPServer/bin/rotatelog.exe -l logs/access_%Y-%m-%d_%H_%M_%S.log 500M" combined
```

References:

<http://publib.boulder.ibm.com/htpserv/manual70/programs/rotatelog.html>



# Tuning the Database – DB2®

## General Tuning

DB2 database tuning in general revolves around reorganisation and runstats. *Reorganisation* changes the physical arrangement of data on disk, and should be done when a significant amount of data is added to the database. *Runstats* provides DB2 with statistics about table contents, allowing it to pick efficient paths to evaluate queries. *Runstats* should be run regularly to ensure that queries are being executed optimally.

**Note:** When applying DB2 fix packs several steps are needed to ensure a smooth upgrade and good operational performance when the upgrade is all done. Consult the IBM DB2 Knowledge Center for the "Applying fix packs" sections:

[http://www.ibm.com/support/knowledgecenter/SSEPGG\\_10.5.0/com.ibm.db2.luw.qb.server.doc/com.ibm.db2.luw.qb.server.doc-gentopic2.html?lang=en](http://www.ibm.com/support/knowledgecenter/SSEPGG_10.5.0/com.ibm.db2.luw.qb.server.doc/com.ibm.db2.luw.qb.server.doc-gentopic2.html?lang=en)

IBM® Connections bundles scripts in the product's connections.sql folder to reorganise and update statistics on each of the application databases. We recommend using the supplied scripts rather than doing a subset of tables manually, as the supplied scripts are optimised for IBM® Connections.

We also recommend a frequent review of the DB2 general error log – db2diag.log – looking for general and critical warnings.

We applied the following setting to all database instances before doing database migrations:

```
db2set DB2CODEPAGE=1208
```

## Autonomic Features

IBM® Connections takes advantage of the self-configuring and self-maintaining features of DB2 V10. These should be left on so that DB2 can optimise itself based on actual workload.

**Key point:** Run the reorg and statistics update scripts for each installed application, bundled with the product. These are especially critical after an initial data load or significant shifts in population.

## DB2\_MINIMIZE\_LISTPREFETCH

We recommend applying this setting otherwise you may experience a high CPU utilisation with a relative small load.

```
db2set DB2_MINIMIZE_LISTPREFETCH=YES
```

*List prefetch* is a special table access method that involves retrieving the qualifying row IDs from an index, sorting them by page number and then prefetching the data pages. Sometimes the optimizer does not have accurate information to determine if list prefetch is a good access method. This might occur when predicate selectivities contain parameter markers or host variables that prevent the optimizer from using catalog statistics to determine the selectivity.

This occurs on a number of the query plans in the IBM® Connections databases. Applying this setting reverts the access



plan back to simple IXSCAN (index scan), thus reducing database CPU utilisation

**Note:** Apply this setting to each DB2 server instance used for the the Connections databases and restart the instance for it to take effect.

For more information on this option see the IBM® DB2 10.5 Knowledge Center

[http://www.ibm.com/support/knowledgecenter/SSEPGG\\_10.5.0/com.ibm.db2.luw.admin.regvars.doc/doc/r0005664.html](http://www.ibm.com/support/knowledgecenter/SSEPGG_10.5.0/com.ibm.db2.luw.admin.regvars.doc/doc/r0005664.html)

## DB2\_SELECTIVITY

We recommend applying this setting to improve the performance of Connections queries that include selectivity hints.

```
db2set DB2_SELECTIVITY=YES
```

**Note:** Apply this setting to each DB2 server instance used for the the Connections databases and restart the instance for it to take effect.

For more information on this option see the IBM® DB2 10.5 Knowledge Center

[http://www.ibm.com/support/knowledgecenter/SSEPGG\\_10.5.0/com.ibm.db2.luw.admin.regvars.doc/doc/r0005664.html](http://www.ibm.com/support/knowledgecenter/SSEPGG_10.5.0/com.ibm.db2.luw.admin.regvars.doc/doc/r0005664.html)

## DB2 Data Compression

DB2 provides a compression technology option that reduces storage requirements, improves I/O efficiency and performance and provides quicker access to data from disk.

**Note:** To use compression, you **must have a license** for the DB2 Storage Optimisation Feature.

IBM®'s measurements with table and index compression applied to all the Homepage database tables yielded the following overall performance improvements:

- 25% reduction in database size on disk
- A 16% reduction in IOPS (Homepage generates by far the single highest number of IOPS at 64% of the total test workload)
- The DB2 and WebSphere CPU utilisation values were quite similar between the un-compressed and compressed DB workloads at the same throughput, i.e. we did not observe a significant CPU penalty when applying DB compression.

**Note:** In addition to the disk space and disk I/O improvements listed above, DB2 Data Compression typically reduces the DB2 buffer pool memory requirements as well. This means that for memory limited DB2 servers and DB2 servers having no farther capacity to expand memory, DB2 Data Compression can help reduce host memory stresses.

For more information on this compression option, see the following documentation:

- IBM® DB2 10.5 Information Center - Data Compression [http://www-01.ibm.com/support/knowledgecenter/SSEPGG\\_10.5.0/com.ibm.db2.luw.admin.dboj.doc/doc/c0055401.html](http://www-01.ibm.com/support/knowledgecenter/SSEPGG_10.5.0/com.ibm.db2.luw.admin.dboj.doc/doc/c0055401.html)
- IBM® developerWorks – Technical article on data compression in DB2 <http://www.ibm.com/developerworks/data/library/techarticle/dm-0605ahuja/>



## ***Enable compression and compress***

As a companion to this tuning guide we have provided in [Appendix C – DB2 Compress Homepage](#) a compression SQL script for the Homepage database.

### **Step 1. Enable Compression on the database**

From a DB2 command window on the DB2 server, issue the following command to run the provided compression script:

```
set DB2INSTANCE=<DB2_instance_name>  
db2 -td@ -vf Compress_homepage_DB2.sql
```

**Note:** <DB2\_instance\_name> is the name of the DB2 instance serving the database. The default DB2 instance is called “DB2”. If you use multiple instances e.g. one for each database, you need to identify the instance serving the Homepage database.

Example:

```
set DB2INSTANCE=DBI_HP  
db2 -td@ -vf Compress_homepage_DB2.sql
```

### **Step 2. Compress the databases**

The compression takes effect only once the table/index **dictionaries** are built, which is usually during the REORG phase.

The “reorg.sql” scripts are provided for all databases with the IBM® Connections DBWizard. Run the one for the Homepage database. Follow the detailed steps documented in the IBM® Knowledge Center “*Improving access performance and defragmenting DB2 database data*”:

[http://www-01.ibm.com/support/knowledgecenter/SSYGQH\\_5.5.0/admin/admin/t\\_admin\\_db\\_maintain\\_compact.dita](http://www-01.ibm.com/support/knowledgecenter/SSYGQH_5.5.0/admin/admin/t_admin_db_maintain_compact.dita)

For more information about building and updating/maintaining the compression dictionaries read:

- “*Table-level compression dictionary creation*”  
[http://www-01.ibm.com/support/knowledgecenter/SSEPGG\\_10.5.0/com.ibm.db2.luw.admin.dboj.doc/doc/c0052331.html?cp=SSEPGG\\_10.1.0](http://www-01.ibm.com/support/knowledgecenter/SSEPGG_10.5.0/com.ibm.db2.luw.admin.dboj.doc/doc/c0052331.html?cp=SSEPGG_10.1.0)
- “*Index compression*”  
[http://www-01.ibm.com/support/knowledgecenter/SSEPGG\\_10.5.0/com.ibm.db2.luw.admin.dboj.doc/doc/c0054539.html?cp=SSEPGG\\_10.1.0/3-0-0-2-5-1](http://www-01.ibm.com/support/knowledgecenter/SSEPGG_10.5.0/com.ibm.db2.luw.admin.dboj.doc/doc/c0054539.html?cp=SSEPGG_10.1.0/3-0-0-2-5-1)



## Tuning the Database – Oracle®

This section presents key performance recommendations for the configuration and set-up of Oracle® on Red Hat Linux for the IBM® Connections product. The recommendations here are based on best practices as well as our empirical results.

The recommendations are mainly divided into five main categories. These are as follows:

1. Oracle® Buffer Cache and Shared Memory Management
2. Use of SQL Profiles
3. Securefile LOBs
4. Oracle® Initialisation Parameters
5. 64-bit Linux Specific Tuning

### Notes:

- IBM Connections 5.5 moved support to Oracle 12c
- These performance recommendations were first evaluated on the Oracle® 11g platform with Connections 5.0 and verified on the Oracle® 12c platform, specifically Oracle® 12.1.0.0.0, with Connections 5.5.
- These tuning settings are based on a system with 12 physical cores, hyper threading enabled, 56GB RAM

## Buffer Cache and Shared Memory Management

Two approaches are presented for the configuration of the Shared Global Area (SGA) and Program Global Area (PGA); these are the Automatic Shared Memory Management (ASMM) and Automatic Memory Management (AMM). For large Connections deployments IBM supports only Automatic Memory Management (AMM), the implementation details of each is presented in the following sections.

### Notes:

- The workload characteristics of IBM® Connections are suited to using AMM. This is what we used for our benchmarks.
- Regardless of what approach you choose you should:
  - Keep paging (memory swap to disk) to a minimum by ensuring sufficient physical memory is available to fit all processes running on the system: the OS itself, Oracle's individual server processes, any other programs running on the system;
  - Oracle recommends that the entire SGA should fit in real memory (i.e.no paging to disk);
  - To attain an optimal memory configuration might require an iterative process of monitoring and further adjustments.

For more information on configuring and tuning memory on Oracle see the following documentation:

- Oracle Database Administrators Guide:  
[https://docs.oracle.com/cd/E11882\\_01/server.112/e25494/memory.htm#ADMIN00207](https://docs.oracle.com/cd/E11882_01/server.112/e25494/memory.htm#ADMIN00207)
- Oracle Database Performance Tuning Guide:  
[https://docs.oracle.com/cd/E11882\\_01/server.112/e41573/memory.htm#PFGRF014](https://docs.oracle.com/cd/E11882_01/server.112/e41573/memory.htm#PFGRF014)

### *Automatic Shared Memory Management (ASMM)*

Automatic Shared Memory Management was introduced in Oracle® 10g and allows the majority of the components of Oracle® shared memory (SGA) to be automatically sized. These components include the buffer pools, the shared pools, the large pool and other areas of memory shared between server processes. The exception is the log buffer which is



sized separately.

ASMM is implemented as follows (using suggested initial values) :

1. Turn off AMM: MEMORY\_TARGET=0
2. Set the maximum amount of memory for the SGA: e.g. SGA\_MAX\_TARGET=28G
3. Set the value for SGA\_TARGET: e.g. SGA\_TARGET=26G
4. Set the value for PGA\_AGGREGATE\_TARGET: e.g. PGA\_AGGREGATE\_TARGET=2G

#### Notes on SGA

- The host must have at least the amount of shared memory specified by SGA\_MAX\_TARGET available to Oracle® (see section on Linux Shared Memory for additional information on how to size it).
- The view V\$SGA\_TARGET\_ADVICE can provide guidance as to the effectiveness of the value set.

In the above recommendation a value is also included for PGA Memory Management (i.e. the aggregate memory specifically used by dedicated server processes). Note that the ratio of SGA\_TARGET to PGA\_AGGREGATE\_TARGET is greater than 10 to 1. This is an indication that the workload characteristics of IBM® Connections leans heavily towards an OLTP workload categorisation (as opposed to a DSS categorisation) and that the SGA/PGA ratio can be aggressively set in favour of the SGA.

#### Notes on PGA

- It should be noted that for IBM®'s suggested initial value, the PGA\_AGGREGATE\_TARGET is set above 1 GB which is a lower limit that should be observed. When the PGA aggregate target is less than 1GB the amount of memory available to individual processes is normally restricted to 200 MB which can present a limitation in performance.
- The view V\$PGA\_AGGREGATE\_ADVICE can provide guidance as to the effectiveness of the value set.

## *Automatic Memory Management (AMM)*

As we have seen previously when using ASMM we can set specific target values for the SGA and PGA. In Oracle® 11g the responsibility for computing the relative values for PGA and SGA can be assigned to Oracle®. As opposed to having a specific hard ratio for SGA and PGA, Oracle® will determine the optimum trade-off between the various areas and periodically adjust memory allocations if required. This is performed by Automatic Memory Management (AMM).

AMM is implemented as follows (using suggested initial values):

1. Set a value for MEMORY\_MAX\_TARGET. Sufficient OS memory is required to support the value set. MEMORY\_MAX\_TARGET=30G.
2. The SGA\_TARGET and PGA\_AGGREGATE\_TARGET can be set to 0. If these values are non-zero then it defines the minimum size for the specified region. We set a minimum for the SGA, e.g SGA\_TARGET=24G and PGA\_AGGREGATE\_TARGET=0
3. Set MEMORY\_TARGET to the total amount of memory you want to share between SGA and PGA. e.g. MEMORY\_TARGET=30G.

#### Notes on AMM

- The view V\$MEMORY\_TARGET\_ADVICE can provide guidance as to the effectiveness of the value set.
- The workload characteristics of IBM® Connections are suited to using AMM.

#### Additional Notes

- Starting with Oracle® Database 11g, the Automatic Memory Management feature requires more shared memory (/dev/shm) and file descriptors. The shared memory should be sized to be at least the greater of MEMORY\_MAX\_TARGET and MEMORY\_TARGET for each Oracle® instance on that computer.
- To determine the amount of shared memory available, enter the following command:
- # df -h /dev/shm/



- MEMORY\_MAX\_TARGET and MEMORY\_TARGET cannot be used when LOCK\_SGA is enabled or with HugePages on Linux.



## Oracle® SQL Profiles

### *Optimizing the Execution Plan Using SQL Profiles*

When a SQL statement is processed an execution plan is generated which details the method by which Oracle® will execute that statement. Oracle® uses a “cost-based” optimizer. The cost optimizer determines the most efficient execution plan for a SQL statement. The optimizer depends heavily on the statistics that you (or the database) gather. In most scenarios, this results in a well-performing execution plan. However, there are certain situations in which an adjustment to the optimizer's standard choice of execution plan can be performed which results in an overall performance benefit. One of the features that can be used to influence the decision path the optimizer uses when selecting a plan is the use of a SQL Profile. A SQL Profile is generated by the SQL Tuning Advisor and effectively provides the optimizer with more information so it can generate a potentially better plan.

### *High CPU and Disk utilisation for Homepage queries due to sub-optimal Oracle® Execution Plans*

When analysing the performance of the IBM® Connections database in our benchmark workload it was observed that certain SQL statements which queried tables related to the HOMEPAGE schema had sub-optimal execution plans associated with them. These queries displayed symptoms that included in certain cases high values for CPU utilisation and high values for “User I/O Wait time” (which is a symptom of high disk utilisation). Together these increased CPU load and disk activity on the database server.

Typically the problem queries were of the following format:

```
SELECT STORY_ID,  
       SOURCE_TYPE ,  
       ROLLUP_ENTRY_ID,  
       CREATION_DATE ,  
       IS_BROADCAST,  
       SAVED_STORY_ID ,  
       SAVED_READER_ID,  
       IS_VISIBLE  
FROM ...
```

These queries can be divided into two categories (High Physical Reads and High Logical Reads). They access the following table segments in the HOMEPAGE NEWSREGTABSSPACE tablespace:

#### High Physical Reads

- NR\_COMMUNITIES\_READERS
- NR\_PROFILES\_READERS
- NR\_FILES\_READERS
- NR\_STATUS\_UPDATE\_READERS
- NR\_AGGREGATED\_READERS

#### High Logical Reads

- NR\_PROFILES\_VIEW
- NR\_AGGREGATED\_READERS
- NR\_COMMUNITIES\_VIEW
- NR\_STATUS\_UPDATE\_READERS

A high number of physical reads was observed against the first set of tables above and this resulted in a high elapsed



time per SQL execution combined with a high value for “User I/O wait time”. A high number of logical reads was observed against the second set of tables and this resulted in a high elapsed time per SQL execution combined with high CPU utilisation.

## Effects of SQL Profile

A SQL Profile was applied to the problem queries, which resulted in substantial beneficial performance effects, including:

- An increase in the overall DB throughput in terms of transactions per second
- A reduction in the end-user transaction response time
- A reduction in the disk I/O subsystem activity (both in terms of IOPS and disk throughput KB/s)
- A reduction in the CPU utilisation of the Database system
- An more efficient use of Database and system memory by accommodating the workload with a reduction in the sizing requirement for the SGA (in particular the Buffer Cache).

The application of SQL Profiles to the HOMEPAGE queries is a very effective solution to this performance issue and in particular addresses the root cause of the issue rather than treating the symptoms.

## Creating and Accepting a SQL Profile

In order to manually create a SQL Profile the SQL Tuning Advisor is run against the problem query or queries.

A problem query can be identified in a number of ways, such as;

- in real-time via the Performance tab of the Oracle® Enterprise Manager (OEM) Console;
- by analysing a snapshot;
- by examining the results of the daily Automatic SQL Tuning job.

When the SQL Tuning Advisor is run it may recommend a SQL Profile. It will show the following information:

- the problem SQL text and SQL ID;
- the expected performance benefit as a percentage;
- a comparison between the existing execution plan and the execution plan that includes the SQL Profile.

An example is shown below in which the expected benefit is 99.98%. The SQL Profile is created by simply clicking the “implement” button to accept the SQL Tuning Advisor recommendation.

Recommendations for SQL ID:7ssq12rx3zdrs Return

---

**Confirmation**  
The recommended SQL Profile has been created successfully.

Only one recommendation should be implemented.

**SQL Text**  
select STORY\_ID, SOURCE\_TYPE, ROLLUP\_ENTRY\_ID, CREATION\_DATE, IS\_BROADCAST, SAVED\_STORY\_ID, SAVED\_READER\_ID, IS\_VISIBLE from( select NEWS\_FEED.STORY\_ID, NEWS\_FEED.SOURCE\_TYPE, NEWS\_FEED.ROLLUP\_E...

**Select Recommendation**  
Original Explain Plan (Annotated)

Implement

Select Type	Findings	Recommendations	Rationale	Benefit Other (%) Statistics	New Explain Plan	Compare Explain Plans
SQL Profile	A potentially better execution plan was found for this statement.	SQL profile "SYS_SQLPROF_01468atbd20b0001" was created by the user for this recommendation and currently has status "ENABLED".		99.98		

Figure 1: Manually Accepting a SQL Profile



## *Automatically Accepting SQL Profiles*

In Oracle® Database 11g or higher the Automatic SQL Tuning job runs on a daily basis (generally during the maintenance window). The Automatic SQL Tuning job identifies problem queries and automatically runs the SQL Tuning Advisor which may generate (but not implement) a SQL Profile.

To view the Automatic SQL Tuning Results using OEM select the following:

- Home->Advisor Central->Automatic SQL Tuning Results
- Select the appropriate time period and select “SQL Profile”

A list of SQL statements that would benefit from a SQL Profile are shown. These can be individually reviewed or all of them can be implemented directly via the “Implement All SQL Profiles” button.

The above operation is a manual operation; however, Oracle® presents the option to automatically implement all SQL Profiles that the Automatic SQL Tuning job identifies. Automatic implementation of SQL Profiles is enabled as follows:

- Home->Advisor Central->Automatic SQL Tuning Results
- Select “Configure” Automatic Implementation of SQL Profiles

### **Automatic SQL Tuning Result Summary**

The Automatic SQL Tuning runs during system maintenance windows as an automated maintenance task, searching for ways to improve the execution plans of high-load SQL statements.

#### **Task Status**

Automatic SQL Tuning (SYS\_AUTO\_SQL\_TUNING\_TASK) is currently **Enabled** [Configure](#)

**Figure 2: Enable Automatic Acceptance of SQL Profiles**

The decision to enable the automatic implementation of SQL Profiles or to selectively implement SQL Profiles manually is one best left to the judgement of the custodians of the specific database system and the policies that govern the administration of the database. The main recommendation is that a review of specific problem queries will benefit from the application of SQL Profiles.

## *Additional Notes*

Many of the tasks highlighted in this section are accomplished via the OEM Console however they can equally be performed via the appropriate DBMS Utility. For example, to automatically enable the acceptance of SQL Profiles the following statement can be executed.

```
BEGIN  
dbms_sqltune.set_auto_tuning_task_parameter('ACCEPT_SQL_PROFILES', 'TRUE');  
END;
```

## *Licensing Requirements for SQL Profiles*

SQL Profiles (because they are part of the SQL Tuning Advisor) requires a license for the Oracle® Tuning Pack (which requires a license for the Diagnostics Pack).

You can determine whether your system is licensed by executing the following SQL statement:



```
SQL> sho parameter control_management_pack_access  
NAME                                TYPE                                VALUE  
-----                                -
```

**Figure 3: List Management Pack Access**



## Oracle® Securefile LOBS

**Note:** “Beginning with Oracle Database 12c, SecureFiles LOB storage became the default in the `CREATE TABLE` statement. If no storage type is explicitly specified, new LOB columns use SecureFiles LOB storage.” see “Using Oracle LOB Storage” [https://docs.oracle.com/database/121/ADLOB/adlob\\_smart.htm#ADLOB4444](https://docs.oracle.com/database/121/ADLOB/adlob_smart.htm#ADLOB4444)

The IBM® Connections tables contain a significant number of LOBs or large objects that are used to store unstructured data. In 11g/12c a LOB can now be stored as a *securefile* or as the standard *basicfile*. By default Oracle® 11g/12c allows for the use of both *securefile* and *basicfile* LOBs in the same database. IBM® Connections tablespaces are created using automatic segment space management (ASSM) and this allows the use of *securefile* LOBs within those tablespaces.

The *securefile* LOB offers a number of performance improvements over *basicfile* LOBs such as:

- Better caching; which improves the speed of LOB writes.
- Less Fragmentation: Oracle® *securefile* allocates contiguous chunks of disk to reduce LOB fragmentation and reduce physical read I/O.
- Sequential prefetch: *securefile* prefetch allows Oracle® to read-ahead and grab the remaining data blocks for a chunk, thereby reducing disk I/O.
- Compression: tables containing *securefile* LOBs can be easily compressed using Oracle® Advanced Compression if required (this can not be performed with a table containing a *basicfile* LOB).

The dynamic parameter `DB_SECUREFILE` is used to specify how to store large objects. The default value in 11g is “Permitted” which allows for *securefile* LOBs to be created. We recommend setting this to “Always” which directs Oracle® to create LOBs as a *securefile* and revert back to a *basicfile* if it can not do so.

The parameter can be changed with an `ALTER SYSTEM` statement before creating the databases, e.g:

```
alter system set db_securefile='ALWAYS' scope=both;
```



## Oracle® Initialisation Parameters

In many cases the default initialisation parameters provided by the Oracle® installation are sufficient, however there are some that may need to be adjusted dependant on the volume of the workload or the number of connected sessions. These are highlighted below.

Dedicated server mode (the default instance configuration) was used in IBM®'s environment. Dedicated server mode ensures that one server is started for every client connection and destroyed when the client disconnects. The values of some initialisation parameters may need to be evaluated based on the expected number of client connections and the physical characteristics of the server machine.

### *Processes*

The processes parameter is effectively equivalent to the maximum number of concurrent users, plus the number of background processes (~22). The default value for certain environments may be low so this should be adjusted to an estimate that takes into account the expected number of concurrent users. The general rule is to set the number of processes to be ~1.5 times the total number of connections across all Application Server nodes. E.g:

```
Total JDBC Pool per AppServer node = 1,850
processes 1 node = 1,850 * 1.5 = 2,775
processes 3 nodes = 2,775 * 3 = 8,325
```

### *Sessions*

The sessions parameter is derived accordingly from the processes parameter as follows:

```
sessions = (1.5 * PROCESSES) + 22
```

This accounts for recursive sessions which may share a server process or connection.

```
Sessions 1 node = (1.5 * 2,775) + 22 = 4,185 (rounded up)
```

### *Open Cursors*

The value of the OPEN\_CURSORS parameter should be set high enough to prevent the application from running out of open cursors (handles to private SQL areas). In IBM®'s environment this was set to 1000.

### **DB Writer Processes**

The DB\_WRITER\_PROCESSES parameter can be set to take advantage of a multi-cpu system that modifies data heavily by enabling multiple DB writer processes. This should be set using the formula:

```
DB_WRITER_PROCESSES = CPU_COUNT / 8
```

e.g. in a 24 cpu\_count environment, DB\_WRITER\_PROCESSES = 3

## 64-bit Linux Specific

### *Virtual Memory*

This section considers parameters that can be modified to tune the System's Virtual Memory. For certain workloads this tuning advice may prevent performance from degrading in a database environment that has a high demand on virtual memory or has a large amount of RAM and is subject to write-heavy loads.



To understand the rationale behind the tuning advice a brief overview is provided of the Linux Swapping and the Page Cache.

## Swapping

Swapping occurs when the kernel moves data between memory and swap space on disk. The kernel swaps out when it needs to free up physical memory (RAM) by pushing data out to disk. It swaps in when an application needs to access data that was swapped out. Swapping is detrimental to application or database performance. Swapping can be prevented by ensuring that there is enough physical memory to accommodate the database instance.

## Tuning Swap Activity

The `vm.swappiness` parameter (value range 0 to 100) defines the kernel's attitude to swapping. The larger the value the more willing the kernel is to swap an application to disk rather than free some memory from its own buffers. By default Linux uses a value of 60 which is tuned toward moderately aggressive swapping to increase disk cache. With this setting the kernel may be encouraged to use the available swap space. As Oracle® is commonly the biggest memory consumer on the system, it may end up being affected by swapping. In such a scenario, the kernel may decide that expanding its own caches is more important than preserving some of Oracle®'s buffers in memory. When running with a dedicated Oracle® server, `vm.swappiness` can be set to 0 to discourage the kernel from swapping out application memory.

## Page Cache and Pdflush

During write-heavy loads, as data is written out, ultimately intended for disk, Linux caches this information in an area of memory called the **Page Cache**. The full set of metrics associated with the page cache can be seen up by executing `cat /proc/meminfo`

As pages are modified, the size of the "Dirty" section will increase. Once writes to disk have begun the "Writeback" figure will go up until the write is finished. Linux writes data out of the page cache using the **pdflush** process. At any moment, between 2 and 8 pdflush threads are running on the system. You can monitor how many are active by executing:

```
cat /proc/sys/vm/nr_pdflush_threads
```

The manner in which pdflush allocates threads is as follows:

- whenever all existing pdflush threads are busy for at least one second, an additional pdflush thread is spawned.
- The new ones try to write back data to device queues that are not congested, aiming to have each device that's active get its own thread flushing data to that device.
- each time a second has passed without any pdflush activity, one of the threads is removed.

The minimum and maximum number of pdflush threads can be modified, but it's very rare they need to be adjusted.

## Tuning Pdflush

The pdflush thread is controlled by a series of parameters in the directory `/proc/sys/vm`.

The page cache on a Linux system is the area where file system based I/O is cached and the designated parameters affect how the cache is utilised by the kernel, more specifically, tuning pdflush for how much RAM based cache to use for dirty pages (pages to be written to disk) and for how frequently to flush those pages by writing the pages to disk. Dirty pages are the pages in memory (page cache) that have been updated and therefore have changed from what is currently stored on disk.

`dirty_writeback_centisecs` (default 500)

In hundredths of a second, this is how often pdflush wakes up to write data to disk. The default wakes up the two (or more) active threads every five seconds.



`dirty_expire_centiseconds` (default 3000)

This is first thing `pdflush` works on is writing pages that have been dirty for longer than it deems acceptable. In hundredths of a second, how long data can be in the page cache before it's considered expired and must be written at the next opportunity. Note that this default is very long: a full 30 seconds. and it means that under normal circumstances, unless you write enough to trigger the other `pdflush` method, Linux won't actually commit anything you write until 30 seconds later.

`dirty_background_ratio` (default 10)

The second thing `pdflush` will work on is writing pages if memory is low. Maximum percentage of memory (where  $\text{memory} = \text{MemFree} + \text{Cached} - \text{Mapped}$ ) that can be filled with dirty pages before `pdflush` begins to write them out.

Based on the above default configuration, data written to disk will sit in memory until either a) they're more than 30 seconds old, or b) the dirty pages have consumed more than 10% of the active, working memory. If you are writing heavily, once you reach the `dirty_background_ratio` driven figure worth of dirty memory, you may find that all your writes are driven by that limit. It's easy to get in a situation where pages are always being written out by that mechanism well before they are considered expired by the `dirty_expire_centiseconds` mechanism.

`dirty_ratio` (default 40)

Maximum percentage of total memory that can be filled with dirty pages before processes are forced to write dirty buffers themselves during their time slice instead of being allowed to do more writes. Note that all processes are blocked for writes when this happens, not just the one that filled the write buffers.

A potential issue for write heavy workloads is that because so much must be written before any physical writes start, the I/O appears more bursty than would seem optimal. This can result in long periods where no physical writes happen at all, as the large page cache is filled, followed by writes at the highest speed the device can achieve once one of the `pdflush` triggers is tripped.

The rationale behind the recommended values is that for systems with large amounts of RAM is that the I/O associated with flushing dirty pages happens more often but with small amounts of I/O as opposed to less often with larger and possibly debilitating amounts of I/O.

## Virtual Memory Tuning Advice

We recommend that you first benchmark your system with the default values before applying the changes listed below.

`dirty_background_ratio` (default=10, recommended=3)

This is the primary value to adjust downward. The goal is to reduce the amount of data Linux keeps cached in memory, so that it writes it more consistently to the disk rather than in a batch, lowering `dirty_background_ratio` is the most effective way to do this. It is more likely the default is too large in situations where the system has large amounts of memory and/or slow physical I/O.

`dirty_ratio` (default=40, recommended=15)

Secondary tunable to adjust.

`dirty_expire_centiseconds` (default=3000, recommended=500)

Test lowering, but not to extremely low levels. Attempting to speed how long pages sit dirty in memory can be accomplished here, but this will considerably slow average I/O speed because of how much less efficient this is. This is particularly true on systems with slow physical I/O to disk. Because of the way the dirty page writing mechanism works, trying to lower this value to be very quick (less than a few seconds) is unlikely to work well. Constantly trying to write dirty pages out will just trigger the I/O congestion code more frequently.



`dirty_writeback_centisecs` (default=500, recommended=100)  
This parameter is the least likely to have any effect and can be potentially be left alone.

To enable the above tuning advice add the following to `/etc/sysctl.conf`.

```
vm.swappiness=0
vm.dirty_background_ratio=3
vm.dirty_ratio=15
vm.dirty_expire_centisecs=500
vm.dirty_writeback_centisecs=100
```

Run `sysctl -p` to load the new values.

## Shared Memory

Shared memory tuning depends upon the size of Oracle®'s system global area (SGA). The amount of server RAM limits the SGA size. To tune shared memory for the SGA, you must calculate:

- the maximum amount of shared memory for the kernel to allocate i.e. the sum of all shared memory segments (shmall);
- the maximum size of any one memory segment (shmmax);
- the maximum number of segments (shmmni).

### Maximum amount of shared memory (shmall)

The maximum amount of shared memory (measured in units of page size) should be configured to support the amount of physical memory in the system.

Obtain page size as follows:  
# getconf PAGE\_SIZE ## page size in bytes  
4096

Example  
If the amount of RAM is 36 GB,  
 $shmall = (36 \text{ GB} / \text{page size}) = 36 * 1024^3 / 4096 = 9,437,184$  [pages]

### Maximum size of a shared segment (shmmax)

The maximum size of a shared segment (shmmax) should be configured to be at least one-half the size of physical memory.

Example  
For our 36 GB example, this parameter should be set to  $36 * 1024^3 / 2$  or 19,327,352,830 [bytes]

### Maximum number of shared-memory segments (shmmni)

Oracle® recommends setting the maximum number of shared-memory segments to 4,096 i.e. it does not scale with RAM unlike shmall and shmmni which are proportional to the amount of server RAM.

### Example

Add the following to `/etc/sysctl.conf`, remembering to adjust shmall, and shmmax to your server's RAM:

```
# The following values are for 36 GB of RAM
kernel.shmall=9437184
kernel.shmmax=19327352830
kernel.shmmni=4096
```



Run `sysctl -p` to load the new values.



## Tuning the Operating System

In any high-load environment, the network must be closely monitored to ensure that its performance is acceptable and consistent. The settings mentioned in the following sections are not necessarily optimal for all environments. Rather, they are mentioned to highlight that network tuning needs to be managed in the performance environment as part of any bottleneck resolution process.

See also Tuning IBM® System x Servers for Performance at: <http://www.redbooks.ibm.com/abstracts/sg245287.html>

### Windows

#### *Network Tuning*

The following registry settings were used on all of the Windows-based systems in IBM®'s environment to improve network performance:

[http://www-01.ibm.com/support/knowledgecenter/SSAW57\\_8.5.5/com.ibm.websphere.nd.doc/ae/welc6toptuning.html](http://www-01.ibm.com/support/knowledgecenter/SSAW57_8.5.5/com.ibm.websphere.nd.doc/ae/welc6toptuning.html)

Windows registry settings can be modified in the Registry Editor tool in the following section:

```
[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters]
```

This can be started from a command line with the command `regedit`. See Table 8 for a discussion of each of these parameters.

*Table 8: Network Tuning Windows*

Parameter	Value
<code>TcpTimedWaitDelay</code>	dword:0000001e

The parameter *TCPTimedWaitDelay* adjusts the amount of time that TCP waits before completely releasing a socket connection for re-use.

To increase the throughput of the server, we decreased *TCPTimedWaitDelay* from default 120 seconds to 30 (0x1E) seconds.



## Linux

### Swapping

For the Application Server we recommend `vm.swappiness` to be set to 0 to discourage the kernel from swapping out application memory. See [Tuning Swap Activity](#) in this guide.

### Network Tuning on Linux

We added the settings listed in Table 9 to file `/etc/sysctl.conf` and then ran the `sysctl -p` command.

Table 9: Network Tuning Linux

Parameter	Value
<code>net.ipv4.ip_forward</code>	0
<code>net.ipv4.conf.default.rp_filter</code>	1
<code>net.ipv4.conf.default.accept_source_route</code>	0
<code>net.core.rmem_max</code>	16777216
<code>net.core.wmem_max</code>	16777216
<code>net.ipv4.tcp_rmem</code>	4096 87380 16777216
<code>net.ipv4.tcp_wmem</code>	4096 65536 16777216
<code>net.ipv4.tcp_fin_timeout</code>	30
<code>net.core.netdev_max_backlog</code>	3000
<code>net.core.somaxconn</code>	10000
<code>net.ipv4.tcp_keepalive_intvl</code>	15
<code>net.ipv4.tcp_keepalive_probes</code>	5

To inspect current TCP parameters, run the following command: `sysctl -a | fgrep tcp`



## Setting Shell Limits for the Oracle® User on the Database System

To improve the performance of the Oracle® software, increase the shell limits for the Oracle® user as shown in Table 10.

Table 10: Setting Shell Limits for the Oracle® User

Shell Limit	Limits.conf Parameter	Hard Limit
Maximum number of open file descriptors	nofile	65536
Maximum number of processes available to a single user	nproc	16384
Maximum size of the stack segment of the process	stack	10240

To increase the shell limits:

1. Add the following lines in the `/etc/security/limits.conf` file:

```
oracle soft nproc 2047
oracle hard nproc 16384
oracle soft nofile 1024
oracle hard nofile 65536
oracle soft stack 10240
```

2. Add the following line to the `/etc/pam.d/login` file, if it does not already exist:

```
session required pam_limits.so
```

3. Depending on the Oracle® user's default shell, make the following changes to the default shell startup file:

- For the Bourne, Bash, or Korn shell, add the following lines in the `/etc/profile` file (or the `/etc/profile.local` file on SUSE Linux Enterprise Server systems):

```
if [ $USER = "oracle" ]; then
    if [ $SHELL = "/bin/ksh" ]; then
        ulimit -p 16384
        ulimit -n 65536
    else
        ulimit -u 16384 -n 65536
    fi
fi
```

- For the C shell, add the following lines in the `/etc/csh.login` file (or the `/etc/csh.login.local` file on SUSE Linux Enterprise Server systems):

```
if ( $USER == "oracle" ) then
    limit maxproc 16384
    limit descriptors 65536
endif
```



## Setting Shell Limits for the Connection User on the Application Server Tier

To improve the performance of the WAS software, increase the following shell limits for the user running Connections.

Table 11: Setting Shell Limits for the Connections User

Shell Limit	Limits.conf Parameter	Hard Limit
Maximum number of open file descriptors	nofile	65536
Maximum number of processes available to a single user	nproc	16384
Maximum size of the stack segment of the process	stack	10240

To increase the shell limits:

1. Add the following lines in the `/etc/security/limits.conf` file:

```
root soft nproc 2047
root hard nproc 16384
root soft nofile 1024
root hard nofile 65536
root soft stack 10240
```

2. Add the following line to the `/etc/pam.d/login` file, if it does not already exist:

```
session required pam_limits.so
```

3. Depending on the root user's default shell, make the following changes to the default shell startup file:

- For the Bourne, Bash, or Korn shell, add the following lines in the `/etc/profile` file (or the `/etc/profile.local` file on SUSE Linux Enterprise Server systems):

```
if [ $USER = "connections" ]; then
    if [ $SHELL = "/bin/ksh" ]; then
        ulimit -p 16384
        ulimit -n 65536
    else
        ulimit -u 16384 -n 65536
    fi
fi
```

- For the C shell, add the following lines in the `/etc/csh.login` file (or the `/etc/csh.login.local` file on SUSE Linux Enterprise Server systems):

```
if ( $USER == "connections" ) then
    limit maxproc 16384
    limit descriptors 65536
endif
```



## ***Other Environments***

We did not run benchmark measurements on other operating systems, but encourage IBM® Connections users to check their platform documentation for any suggested network tuning.

***Key point:*** In all cases, monitor the deployment for good network health and expected speeds and throughput.



## Tuning IBM® Connections Applications

### *Profiles Login Attributes*

A user can log in with a number of attributes, depending on their company policy, (User ID, eMail ID, and so on). These attributes can be specified in `profiles-config.xml`. The default out of the box installation of IBM® Connections specifies three of the most common attributes that customers may use to login, as in the following example. A customer can also specify other attributes:

```
<directory>
  <!--
    Specifies the profiles field that is used to resolve person
records via WALTZ / Javelin
  -->
  <lconnUserIdField>guid</lconnUserIdField>

  <!--
    Lists fields that will be used to resolve user at login time
  -->
  <loginAttributes>
    <loginAttribute>uid</loginAttribute>
    <loginAttribute>email</loginAttribute>
    <loginAttribute>loginId</loginAttribute>
  </loginAttributes>
</directory>
```

Each attribute listed in the `profiles-config.xml` incurs a *look-up* performance cost. For the majority of customers using a single attribute login option, the system can be configured to use a single value to avoid unnecessary multiple look-ups and therefore improve performance.

Here is an example for a customer using the UID as the login attribute:

```
<loginAttributes>
  <loginAttribute>uid</loginAttribute>
</loginAttributes>
```

For more information see:

- Choosing login values:  
[http://www.ibm.com/support/knowledgecenter/SSYGQH\\_5.5.0/admin/install/c\\_install\\_choose\\_login.dita](http://www.ibm.com/support/knowledgecenter/SSYGQH_5.5.0/admin/install/c_install_choose_login.dita)

### *Tuning the Social Analytics Indexer*

To build the Social Analytics (SaND) index there is a trade-off between the time it takes and the memory it uses for this task. There is a configuration option to control the number of iterations of relationship data collection in the graph. This affects the trade-off as follows:

- The higher the number of iterations, the longer it takes to build the graph, but its memory consumption is much lower.
- The lower the number of iterations, the less time it takes to build the graph, but its memory consumption is higher



The default value is set to 100. On large data populations, we saw good performance improvement by lowering this value to 20.

- `SearchCellConfig.setSandIndexerTuning("graph",20)`

For more information, see: [http://www-01.ibm.com/support/knowledgecenter/SSYGQH\\_5.5.0/admin/admin/t\\_admin\\_search\\_tune\\_sand\\_indexing.dita](http://www-01.ibm.com/support/knowledgecenter/SSYGQH_5.5.0/admin/admin/t_admin_search_tune_sand_indexing.dita)

### ***Tuning CCM via the Administration Console for Content Platform Engine (ACCE)***

Start the Administration console:

```
http://<connections_server>/acce
```

**Note:** For detailed instruction, see: <http://pic.dhe.ibm.com/infocenter/p8docs/v5r2m0/index.jsp>  
Then navigate to: *Administering FileNet P8 > Administering Content Platform Engine > Getting Started*

### ***Principal Cache Max Entries:***

The Principal Cache caches basic information (such as IDs) about individual users, groups and communities. The optimal number may differ depending on your system and how many users, groups and communities are typically in active use over the course of several hours.

See also:

[http://www-01.ibm.com/support/knowledgecenter/SSNW2F\\_5.2.0/com.ibm.p8.ce.dev.ce.doc/sec\\_concepts.htm#sec\\_concepts\\_security\\_caching?cp=SSNW2F\\_5.2.0/10-2-1-2&lang=en](http://www-01.ibm.com/support/knowledgecenter/SSNW2F_5.2.0/com.ibm.p8.ce.dev.ce.doc/sec_concepts.htm#sec_concepts_security_caching?cp=SSNW2F_5.2.0/10-2-1-2&lang=en)

The value suggested below is what worked well with our workload and population.

#### **How to set:**

In the Administration Console for Content Platform Engine (ACCE), click:

***ICDomain → Properties → Subsystem Configurations → From drop-down select: “6) Server Cache Configuration” → Principal Cache Max Entries***

Set the value to: 98304

Then save.

### ***Verify the User Token Cache Entry TTL***

In the Administration Console for Content Platform Engine (ACCE), click:

***ICDomain → Server Cache Subsystem → User Token Cache → Entry TTL***

The value should be already set to 3600 seconds.

### ***Scheduled Jobs***

In a deployment of IBM® Connections, a number of maintenance tasks are performed regularly using the IBM® WebSphere Application Server scheduling service. In some cases, administrators might want to change the default



schedule or disable tasks not necessary to the environment.

Please consult the product documentation for instructions on administering managed tasks for IBM® Connections.

## ***Social Analytics Widgets***

The social analytics (SaND) widgets in IBM® Connections have a performance impact in deployments with rich social networks, and may reduce system capacity if left in the default configuration.

A tuning strategy to mitigate this performance cost is relocating the widgets so they are displayed less frequently. For example, in the Profiles application, these widgets were moved from the default view of a user's profile, to separate tabs within that view; thus, they are only executed when explicitly requested. In a performance benchmark, so configured, capacity increased 10%.

Full instructions for managing these widgets are documented in the IBM® Connections wiki at:

[http://www-01.ibm.com/support/knowledgecenter/SSYGOH\\_5.5.0/admin/admin/t\\_admin\\_profiles\\_edit\\_widgets.dita](http://www-01.ibm.com/support/knowledgecenter/SSYGOH_5.5.0/admin/admin/t_admin_profiles_edit_widgets.dita)

## ***Public Blogs Page Cache-Control***

To reduce the CPU usage of both WAS and the Database servers, the Public Blogs page has a public cache-control which makes it cachable by a caching proxy and the browser. The default age is 5 seconds.

```
Cache-Control:"public, max-age=5, s-maxage=5"
```

For web sites that don't have frequent blog updates or where a delay of the new items being displayed on the public blog page can be tolerated, we can tune the cache control age to a longer time. For example 1 minute or even longer depending on the usage.

To tune this value edit the following file:

```
<was_profile_root>\installedApps\<cell_name>\Blogs.ear\blogs.war\WEB-INF\classes\roller.properties
```

Then change the value of the following property: (5 indicates 5 seconds)

```
http_cache.maxAge.html=5
```

## ***Tuning Activities Subject Cache***

In Connections 5.5 we have increased the subject cache to support up to 2000 concurrent users. For *Activities* workloads with more than 2000 concurrent users per cluster member we recommend to tune the Activities Subject Cache according the number of concurrent users accessing the system. NOTE: *Communities* concurrent users are also implicitly *Activities* concurrent users as well.

### **Main Steps**

1. Tune subject cache size
2. Tune Activities Application Server Java Heap
3. Monitor the cache size occupancy and hit rate

### **Step 1: Tune subject cache size**



Calculate the Subject Cache size with the following formula:

$$(\text{number of active users} / \text{number of nodes in cluster}) * 3$$

For example for a population of 5,000 active users, we would calculate a cache size of  $(5000 / 2) * 3 = 7500$ .

Edit

```
<was_profile_root>\installedApps\<cell_name>\Activities.ear\oawebui.war\cacheinstances.properties
```

change the cacheSize line:

```
cache.instance.0=/services/cache/activitiesSubject  
cache.instance.0.cacheSize=7500
```

## Step 2: Tune Activities Application Server Java Heap

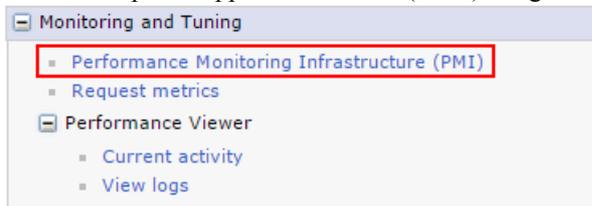
Increase the heap size of the Activities server by 10MB for each 5,000 entries allowed in the cache (to avoid running low on heap space due to the larger cache).

## Step 3. Monitor the cache size occupancy and hit rate.

Ensure that the cache has sufficient space (i.e. not full) and verify hitrate %

### Setup Monitoring

From the WebSphere Application Server (WAS) Integrated Solutions Console



Performance Monitoring Infrastructure (PMI) ?

**Performance Monitoring Infrastructure (PMI) > ActivitiesCluster\_server1**

Use this page to configure Performance Monitoring Infrastructure (PMI)

**Runtime** Configuration

---

**General Properties**

- Use sequential counter updates
- Persist my changes

**Currently monitored statistic set**

- None  
No statistics are enabled.
- Basic  
 Provides basic monitoring, including Java EE and the top 38 statistics.
- Extended  
 Provides extended monitoring, including the basic level of monitoring plus workload monitor, performance advisor, and Tivoli resource models.
- All  
 All statistics are enabled.
- Custom**  
Provides fine-grained control to selectively enable statistics.

Apply OK Reset Cancel



IBM® Connections Version 5.5 – Social Software for Business IBM® Connections  
Performance Tuning Guide

Performance Monitoring Infrastructure (PMI)

Performance Monitoring Infrastructure (PMI) > ActivitiesCluster\_server1 > Custom monitoring level

Use this page to configure Performance Monitoring Infrastructure (PMI)

Runtime Configuration

Dynamic Caching

Object: com/ibm/connections/

Object: ic/services/cache/OA

Object: ic/services/cache/OA

Object: services/cache/activit

Object Cache

Counters

Object: services/cache/organ

Object: ws/WIMInitCache

Object: ws/WSSecureMap

Object: ws/WSSecureMapNo

Enable Disable

Select	Counter	Type	Description	Status
<input checked="" type="checkbox"/>	InMemoryCacheEntryCount	CountStatistic	The current number of in-memory cache entries.	Enabled
<input type="checkbox"/>	MaxInMemoryCacheEntryCount	CountStatistic	The maximum number of in-memory cache entries.	Disabled
Total 2				

Performance Monitoring Infrastructure (PMI)

Performance Monitoring Infrastructure (PMI) > ActivitiesCluster\_server1 > Custom monitoring level

Use this page to configure Performance Monitoring Infrastructure (PMI)

Runtime Configuration

Dynamic Caching

Object: com/ibm/connections/

Object: ic/services/cache/OA

Object: ic/services/cache/OA

Object: services/cache/activit

Object Cache

Counters

Object: services/cache/organ

Object: ws/WIMInitCache

Object: ws/WSSecureMap

Object: ws/WSSecureMapNo

Object: ws/com.ibm.workplac

Object: ws/com.ibm.ws.wsse

Object: ws/com.ibm.ws.wsse

Object: ws/ldap/AttributesCac

Object: ws/ldap/SearchResult

Servlet: baseCache

JDBC Connection Pools

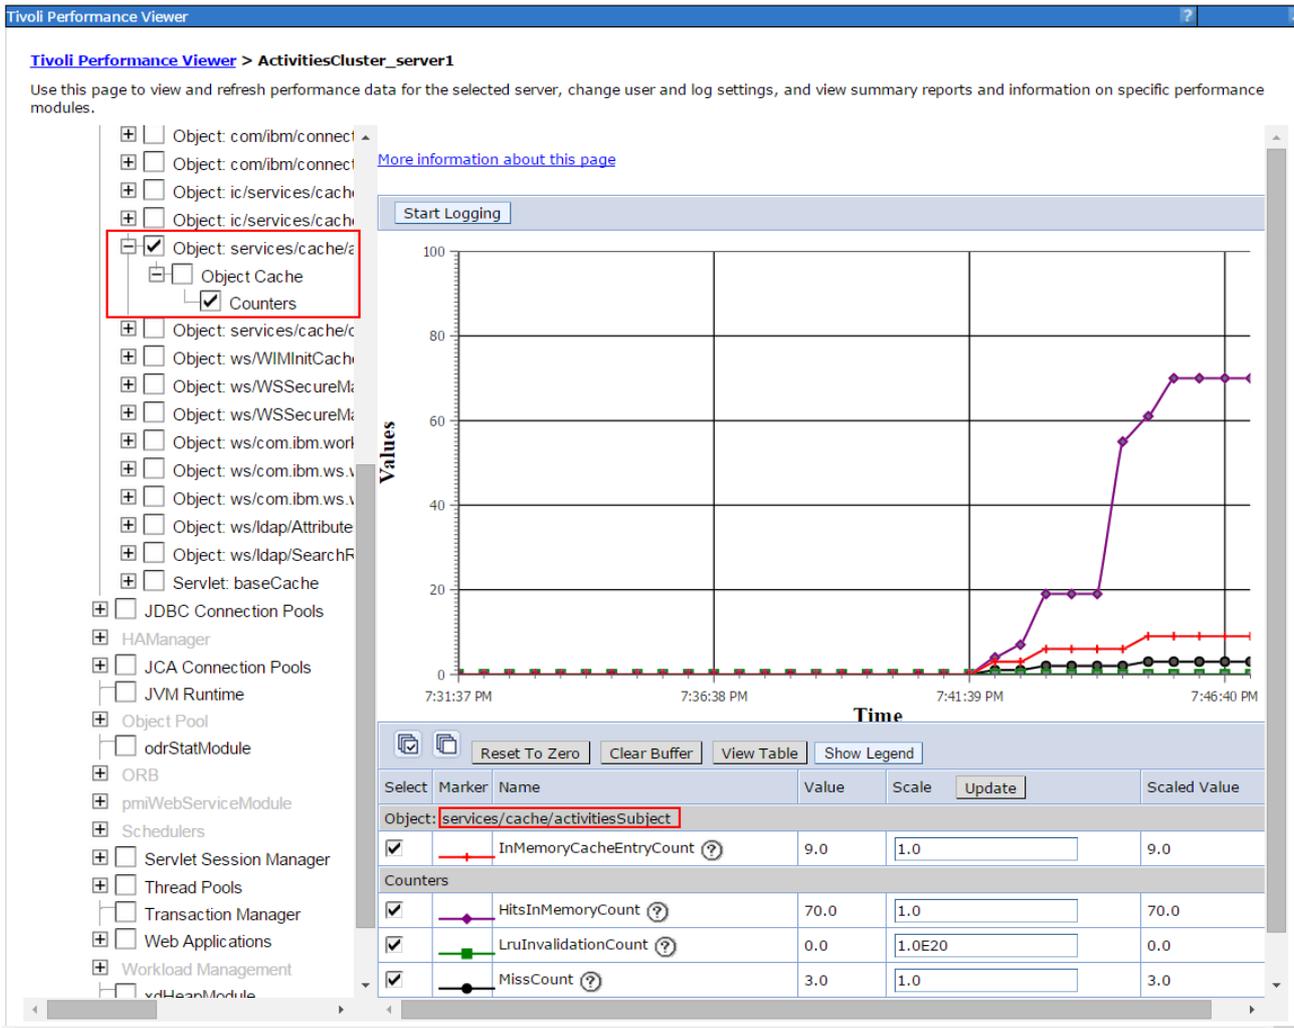
HAManager

Enable Disable

Select	Counter	Type	Description	Status
<input type="checkbox"/>	ClientRequestCount	CountStatistic	The number of requests for cacheable objects that are generated by applications running on this application server.	Disabled
<input type="checkbox"/>	DistributedRequestCount	CountStatistic	The number of requests for cacheable objects that are generated by cooperating caches in this replication domain.	Disabled
<input type="checkbox"/>	ExplicitDiskInvalidationCount	CountStatistic	The number of explicit invalidations resulting in the removal of an entry from disk.	Disabled
<input type="checkbox"/>	ExplicitInvalidationCount	CountStatistic	The number of explicit invalidations.	Disabled
<input type="checkbox"/>	ExplicitMemoryInvalidationCount	CountStatistic	The number of explicit invalidations resulting in the removal of an entry from memory.	Disabled
<input checked="" type="checkbox"/>	HitsInMemoryCount	CountStatistic	The number of requests for cacheable objects that are served from memory.	Enabled
<input type="checkbox"/>	HitsOnDiskCount	CountStatistic	The number of requests for cacheable objects that are served from disk.	Disabled
<input type="checkbox"/>	InMemoryAndDiskCacheEntryCount	CountStatistic	The current number of used cache entries in memory and disk.	Disabled
<input type="checkbox"/>	LocalExplicitInvalidationCount	CountStatistic	The number of explicit invalidations generated locally, either programmatically or by a cache policy.	Disabled
<input type="checkbox"/>	LruInvalidationCount	CountStatistic	The number of cache entries that are removed from memory by a Least Recently Used (LRU) algorithm.	Enabled
<input checked="" type="checkbox"/>	MissCount	CountStatistic	The number of requests for cacheable objects that were not found in the cache.	Enabled
<input type="checkbox"/>	RemoteCreationCount	CountStatistic	The number of cache entries that are received from cooperating dynamic caches.	Disabled
<input type="checkbox"/>			The number of explicit invalidations received	



### Monitor cache statistics under load



### Calculate Hit rate

$$\text{hitrate}\% = (\text{HitsInMemoryCount}) / (\text{HitsInMemoryCount} + \text{MissCount}) * 100$$

e.g.  $70 / (70+3) * 100 = 95.89\%$

### Cache Size Occupancy

Ensure the InMemoryCacheEntry count is below the `cache.instance.0.cacheSize` you have set in *Step 1*

**Note:** Poor hit rate would also suggest the cache size was not sufficient.



## ***Tuning the PushNotifications server for FileSync and Notifications Center***

The *PushNotifications* server is used to push updates to the Notifications Center and to the File Sync clients using a technique called “Long Polling”; a connection is held open between the web client and the web server until:

- a) the server has new information and it pushes it to the client. That connection is then closed;
- b) a default time-out of 90 seconds, the connection is then closed. This is to prevent problems with infinite connections which are not expected by most network topologies.

The connection is then re-opened anew until a) or b) occur again in an infinite loop.

**Note:** For more than 5,000 concurrent connections a Proxy server that supports “non-blocking I/O” such as WebSphere Application Server (WAS) Proxy or an F5 server is needed to route the *Long-Poll* connections to the *PushNotifications* server cluster.

For instructions on how to configure and tune the WebSphere Application Server (WAS) Proxy please refer to the IBM Connections 5.5 product documentation in Knowledge Center.

*PushNotifications* is very different from the other Connections components in terms of architecture, design and tuning so it runs in its own separate server cluster to achieve better performance and avoid affecting other components' performance.

### **Notes:**

- The tuning settings listed below are for one PushNotifications node supporting up to 30,000 concurrent connections. For a higher number of concurrent connections, we recommend scaling horizontally on multiple nodes.
- The “WebSphere authentication cache” is a global cache and needs to be multiplied by the number of nodes. See details below in step 7.

### **Step 1. JVM Tuning for PushNotifications**

**How to Set:** In the WebSphere Integrated Solutions Console, click:

***Servers → Server Types → WebSphere application servers → <PushNotificationCluster\_server1> → Server Infrastructure: Java and Process Management → Process Definition → Java Virtual Machine***

Less than 15,000 concurrent connections per node:

Initial and maximum heap size: **4096**

Generic JVM arguments: -Xgcpolicy:**gencon** -Dsun.nio.MaxDirectMemorySize=**1048576000**

Up to 20,000 concurrent connections per node:

Initial and maximum heap size: **6144** MB

Generic JVM arguments: -Xgcpolicy:**balanced** -Dsun.nio.MaxDirectMemorySize=**1572864000**

Up to 30,000 concurrent connections per node:

Initial and maximum heap size: **8192** MB

Generic JVM arguments: -Xgcpolicy:**balanced** -Dsun.nio.MaxDirectMemorySize=**2097152000**

### **Step 2. PushNotifications WorkManager**

Tune this work manager thread pool larger than default to avoid thread hang during peak hours:

**How to Set:** In the WebSphere Integrated Solutions Console, click:



**Resources** → **Asynchronous beans** → **Work managers** → **PushNotificationWorkManager**

- Work request queue size: 400
- Work request queue full action: fail
- Number of alarm threads: 15
- Minimum number of threads: 1
- Maximum number of threads: 40
- Growable: false / unchecked

**Step 3.** DB2 data source: jdbc/pushnotification

**How to Set:** In the WebSphere Integrated Solutions Console, click:

**Resources** → **JDBC** → **Data sources** → **<pushnotification>** → **Connection pool properties**

Maximum connection for the pool: 100

**Step 4.** TCP transport channel: Increase the number of Maximum Open Connections

**How to Set:** In the WebSphere Integrated Solutions Console, click:

**Servers** → **Server Types** → **WebSphere application servers** → **<PushNotificationCluster\_server1>** → **Web Container Settings** → **Web container transport chains**

**Edit the following transport chains:**

- WCInboundDefault
- WCInboundDefaultSecure
- HttpQueueInboundDefault
- HttpQueueInboundDefaultSecure

Click on “TCP Inbound channel” and Set Maximum open connections to: 50,000

This is a good value to support 30,000+ concurrent users. Also leaves some buffer for TIME\_WAIT, CLOSE\_WAIT etc..

**Servers :** Apply all the above to the following Application Server cluster nodes:

- PushNotifications
- Files
- News

**OS TCP Tuning:** When increasing the WAS limits for TCP connections please also check the OS limits.

E.g. on Linux run:

```
cat /proc/sys/net/ipv4/ip_local_port_range
```

and if the reported range is too small then adjust the limits accordingly. E.g. on Linux do

```
echo "1024 64000" >/proc/sys/net/ipv4/ip_local_port_range
```

**Step 6.** Web-Container thread pool

**How to Set:** In the WebSphere Integrated Solutions Console, click:

**Servers** → **Server Types** → **WebSphere application servers** → **<PushNotificationCluster\_server1>** →



***Additional Properties: Thread Pools → Web Container → Thread Pool***

Maximum connections: 100

**Step 7.** WebSphere authentication cache

**How to set:** In the WebSphere Integrated Solutions Console, click:

**Global Security → Authentication cache setting → Max cache size = 40,000**

**Notes:**

- The “WebSphere authentication cache” is a global cache and needs to be multiplied by the number of nodes e.g:
  - 30K connections == 1 node == Auth Cache: 40,000
  - 90K connections == 3 nodes == Auth Cache: 120,000
- You may have to increase these values if other workloads (concurrent users) are hitting the Connections system.

**Step 8.** LDAP tuning

This depends on the brand of LDAP you use. The goal for this step is to tune the LDAP server to support more concurrent connections, to support possible authentication floods caused by timed LTPA token expiries.

**How to set :** For IBM TDI,  
locate the ibmsldapd.conf file e.g. `c:\idslapd-dsrdbm01\etc\ibmsldapd.conf`

in the section `dn: cn=Front End, cn=Configuration`  
add the following directive: `ibm-slapdSetenv: SLAPD_OCHANDLERS=30`

Locate the existing directive: `ibm-slapdDbConnections: 20`  
increase the value to: `ibm-slapdDbConnections: 50`

Reference: <http://www-01.ibm.com/support/docview.wss?uid=swg21165974>

**Step 9:** Increase LTPA time-out to reduce re-authentications.

The WebSphere LTPA time-out has a default of 120 minutes. This value was considered too small for dedicated deployments of the *PushNotifications* service. Increasing the time-out reduces the number of re-authentications.

In our load tests at 40,000 long-poll concurrent connections per server node with the default LTPA time-out, the load caused by re-authentications became significant: ~ 6 re-authentications/second. We experimented with a LTPA time-out of 18.5 hours; the re-authentication load was reduced to 0.6 re-authentications/second

**How to Set:** In the WebSphere Integrated Solutions Console, click:

**Security → Global security → LTPA**

LTPA timeout: 1120

**Step 10:** Hystrix Thread Pool



**How to set :** Edit

```
C:\IBM\WebSphere\AppServer\profiles\Dmgr01\config\cells\appserver1Cell101\LotusConnections-config\hystrix-config.properties
```

Core thread-pool size. This is the maximum number of concurrent HystrixCommands that can be executed. Increase the value to 75

```
hystrix.threadpool.default.coreSize=75
```

Monitor the SystemOut.log of *PushNotifications* server nodes. If you see any of these messages you may need to increase that value.

```
com.netflix.hystrix.exception.HystrixRuntimeException:  
DS_SearchAccountByExactLoginUserNameMatch_0 could not be queued for execution and no  
fallback available.  
java.util.concurrent.RejectedExecutionException: Rejected command because thread-pool  
queueSize is at rejection threshold.  
at com.netflix.hystrix.HystrixCommand.subscribeWithThreadIsolation(HystrixCommand.java:1158)
```

**Step 11:** Long-Poll time-out

Ideally the Long-Poll connection time-out would be infinite but this is limited by various tiers in the network infrastructure such as proxies, routers, etc. The recommendation is to set it as long as possible as allowed by the network delivery chain i.e. to the network device with the lowest connection time-out. Longer time-outs result in less-frequent requests to the PushNotifications server and thus lower CPU load on the application server.

The default time-out is 90 seconds and fits most networks but you can change it to suit your network delivery infrastructure.

**How to set :** Edit

```
C:\IBM\WebSphere\AppServer\profiles\Dmgr01\config\cells\appserver1Cell101\LotusConnections-config\pushnotification-config.xml
```

```
<server>  
  <transports interval="0" maxQueue="-1" metaConnectDeliverOnly="true" timeout="90000" />  
</server>
```

**Note:** Time-out is expressed in milliseconds.



# System Monitoring

## Overview

In addition to performance tuning, it is important to monitor the system for possible performance bottlenecks. Performance monitoring is a very broad topic, with a great number of tools available to make the task more efficient. This document is not going to attempt to cover the entire field of performance monitoring. Rather, this section will discuss important metrics which we have found useful when monitoring systems in IBM®'s performance measurements.

## Areas to Monitor

### *Metrics*

The following metrics are important to monitor in a IBM® Connections deployment:

- **Processor usage:** Typically this is shown as a percentage of the available processor power available, so 50% utilisation indicates that half of the available processing power is in use. Hardware multi-threading can make a single physical processor core appear as two or more processors. Keep this in mind when looking a processor usage, since operating systems typically report utilisation as a percentage of available processors.
- **Physical memory:** Monitor the amount of free memory available within the machine. If possible, also monitor paging activity (pages swapped in/out). Together, these provide a good picture of whether enough physical memory is available.
- **JVM memory:** In addition to the total memory on the system, monitor the memory usage within the Java heap of all IBM® Connections application servers. Enable verbose garbage collection on those application servers so that it is possible to view memory usage statistics within the Java heap.
- **Disk I/O:** The metrics available for monitoring disk I/O will vary by platform. The key is to be able to determine if the disk system is overloaded. Look at the portion of time the disk is busy (or on Windows, the portion of the time the disk is idle). Another useful metric where it's available is the disk queue length.
- **Network I/O:** Look at both read and write rates in bytes per second (or megabytes per second).
- **Logs:** Periodically check the logs from the servers in the environment. This includes `SystemOut.log` and `SystemErr.log` on the application servers, `db2diag.log` on DB2 databases (or corresponding logs on other databases), and any logs on the HTTP servers or proxy servers.

### *Systems*

Remember to monitor all of the systems which make up the IBM® Connections deployment. Start with the entry point to the system, such as a load balancer or reverse proxy, go through the HTTP servers and application servers, to the database servers and directory (LDAP) servers. Also make sure to monitor the network link to the system entry point as well as network links between different tiers. Network bandwidth and latency have significant impact on the overall response time of the client (the browser).



## *Virtualised Environments*

Virtualised environments provide additional deployment flexibility and may reduce overall hardware costs. However, they can make performance monitoring more difficult. It's important to monitor utilisation within the virtual machines to see if a specific virtual machine is seeing a performance problem. It's equally important to monitor the physical hosts where the virtual machines are located. This is especially important if resources on the physical host are over committed, where more processors, memory, or other resources are allocated to virtual machines than exist in the physical machine. This is especially true for memory which should never swap to disk.

## What to look for

It is hard to make broad, accurate generalisations about what performance metrics indicate a problem. Therefore consider these as guidelines rather than hard-and-fast rules of what does or doesn't indicate a problem. With that in mind, consider the following:

- **Processor usage:** On the reverse proxy and application server tiers, a processor load over 70% that is sustained for at least 5 minutes indicates that the processor may be heavily loaded. On database servers, a processor load over 50% for at least 5 minutes may indicate heavy load. If the processor is heavily loaded, more nodes or nodes with more processors may be needed to handle additional load.
- **Physical memory:** The system should always have enough physical memory to fit all running processes in memory; only minimal paging activity should be occurring. Sustained paging activity will cause significant performance problems, particularly on the application server tier.

When attempting to estimate the memory requirement for the application server tier, remember that the process size for an application server is larger than just its heap size. For example, a JVM which has a heap size of 1 GB might consume a total of 1.3 GB or more.

- **JVM memory:** IBM® Connections uses the generational (or 'gencon') garbage collector by default. This setting means that two types of garbage collections will be run: short, frequent collections of the 'nursery' region of the heap, and longer, infrequent collections of the 'tenured' region of the heap.

Nursery garbage collections typically should account for no more than 10% of elapsed time. If nursery collections are consuming a higher portion of time, consider increasing the nursery size with the `-Xmn` JVM parameter.

Garbage collections of the tenured region should come at least 10 minutes apart under high load. If they are coming more frequently, consider increasing the total heap size to make more memory available in the tenured region. When increasing the heap size, make sure to not exceed the available physical memory.

- **Disk I/O:** The performance of disk systems will vary based on the workload it is handling, so it is not usually possible to compare the current workload against the maximum capacity of the disk system.

We have found it useful to monitor the portion of the time the disk system is idle. Our observation was that when the amount of idle time on a disk system drops below 50% for a sustained period (at least 5 minutes), that disk system may start to impact the performance of the IBM® Connections system.

The other indicator we found useful was to look at the disk queue length (on Windows systems). When the disk queue length was higher than the number of physical disks in the disk array for a sustained period, this indicated that the disk system was becoming overloaded and impacting system performance.

- **Network I/O:** With modern network switches, it is usually possible to run a network interface at 60% of its stated capacity for a sustained period of time, with higher bursts possible. For example, a 100 Mbps network is usually able to sustain 7.5 megabytes (60 megabits) per second of traffic. Sustained use above this level can



indicate a network bottleneck.

Make sure to check the actual network interface speed and duplex setting, not the adapter's rated speed. A 1Gbps full-duplex network adapter which is running at 100Mbps half-duplex will not be able to handle 40 megabytes/second of network traffic.

- **Logs:** Logs can contain a wealth of information. The following are messages we have seen, which can indicate performance problems:

- WebSphere SystemOut - HMGR0152W: CPU Starvation detected. Current thread scheduling delay is 9 seconds.

This can indicate excessively high processor load on the application server, or possibly other performance problems. More discussion of this message is available at <http://www.ibm.com/support/docview.wss?uid=swg21236327>.

- WebSphere SystemOut - WSVR0605W: Thread *threadname* has been active for *hangtime* and may be hung. There are *totalthreads* threads in total in the server that may be hung.

If *threadname* is a WebContainer thread, such as WebContainer : 3, this can indicate a problem, as the WebContainer threads should typically handle requests quickly. However, IBM® Connections has other threads which perform background work and which may execute for long periods of time, so this message is of less concern if it references one of those background threads.



## Appendix A - Reference Material

- IBM® Connections Version 5.5 Product Documentation:  
[http://www.lotus.com/ldd/lcwiki.nsf/xpViewCategories.xsp?lookupName=Product Documentation](http://www.lotus.com/ldd/lcwiki.nsf/xpViewCategories.xsp?lookupName=Product+Documentation).
- WebSphere Application Server Performance Cookbook  
<http://publib.boulder.ibm.com/httserv/cookbook/>
- WebSphere Application Server Network Deployment Version 8.5 documentation:  
[http://www-01.ibm.com/support/knowledgecenter/SSAW57\\_8.5.5/com.ibm.websphere.nd.multiplatform.doc/ae/welcome\\_ndm\\_p.html](http://www-01.ibm.com/support/knowledgecenter/SSAW57_8.5.5/com.ibm.websphere.nd.multiplatform.doc/ae/welcome_ndm_p.html).
- DB2 10.5 Information Center:  
[http://www-01.ibm.com/support/knowledgecenter/SSEPGG\\_10.5.0/com.ibm.db2.luw.kc.doc/welcome.html](http://www-01.ibm.com/support/knowledgecenter/SSEPGG_10.5.0/com.ibm.db2.luw.kc.doc/welcome.html).
- Tuning IBM® System x Servers for Performance:  
<http://www.redbooks.ibm.com/abstracts/sg245287.html>
- Understanding how the JVM uses native memory on Windows and Linux  
<http://www.ibm.com/developerworks/linux/library/j-nativememory-linux/index.html>
- DB2 best practices: Storage optimisation with deep compression  
<http://www.ibm.com/developerworks/data/bestpractices/deepcompression/index.html>
- DB2 best practices: Managing data growth  
<http://www.ibm.com/developerworks/data/bestpractices/managingdatagrowth/index.html>
- Best Practices: Tuning and Monitoring Database System Performance  
<http://www.ibm.com/developerworks/data/bestpractices/systemperformance/>
- Troubleshooting:  
[http://www-01.ibm.com/support/knowledgecenter/SSYGQH\\_5.5.0/admin/troubleshoot/ts\\_c\\_welcome.dita?lang=en](http://www-01.ibm.com/support/knowledgecenter/SSYGQH_5.5.0/admin/troubleshoot/ts_c_welcome.dita?lang=en)
- WebSphere Application Server Performance  
<http://www.ibm.com/software/webservers/appserv/was/performance.html>
- developerWorks IBM® Cognos 10 BI Community  
<https://www.ibm.com/developerworks/mydeveloperworks/groups/service/html/communityview?communityUuid=8d7e4a2b-2364-4719-9f4e-aa9e24db7465>



IBM® Connections Version 5.5 – Social Software for Business IBM® Connections  
Performance Tuning Guide

- developerWorks Technical Library: IBM® Cognos Proven Practices  
[http://www.ibm.com/developerworks/views/data/libraryview.jsp?site\\_id=1&contentarea\\_by=AllZones&sort\\_by=Date&sort\\_order=2&start=1&end=100&topic\\_by=-1&product\\_by=-1&type\\_by=AllTypes&show\\_abstract=true&search\\_by=IBM Cognos Proven Practices&industry\\_by=-1&series\\_title\\_by=](http://www.ibm.com/developerworks/views/data/libraryview.jsp?site_id=1&contentarea_by=AllZones&sort_by=Date&sort_order=2&start=1&end=100&topic_by=-1&product_by=-1&type_by=AllTypes&show_abstract=true&search_by=IBM+Cognos+Proven+Practices&industry_by=-1&series_title_by=)
- Performance Tuning for IBM® Cognos 10  
<http://www-01.ibm.com/support/docview.wss?uid=swg21618119>



## Appendix B - Oracle® Initialisation Parameters Summary

Note that the suggested initialisation values below may be changed directly using the EM Console or using the “alter system” command.

Table 12: Oracle® Option A: Automatic Shared Memory Management (ASMM)

Parameter	Recomm. Value	Command
memory_target	0	alter system set memory_target =0 scope = spfile;
sga_max_target	28G	alter system set sga_max_target = 12000M scope = spfile;
sga_target	26G	alter system set sga_target = 12000M scope = spfile;
pga_aggregate_target	2G	alter system set pga_aggregate_target = 1200M scope = spfile;

Table 13: Oracle® Option B: Automatic Memory Management (AMM)

Parameter	Recomm. Value	Command
memory_max_target	30G	alter system set memory_max_target =14464M scope = spfile;
memory_target	30G	alter system set memory_target = 12000M scope = spfile;
sga_target	24G	alter system set sga_target = 0 scope = spfile;
pga_aggregate_target	0	alter system set pga_aggregate_target = 0 scope = spfile;

Table 14: Oracle® Initialisation Parameters

Parameter	Recomm. Value	Command
db_securefile	0	alter system set db_securefile='ALWAYS' scope = spfile;
processes	2000	alter system set processes = 2000 scope = spfile;
sessions	3024	alter system set sessions =3024 scope = spfile;
open_cursors	1000	alter system set open_cursors = 1000 scope = spfile;
db_writer_processes	3	alter system set db_writer_processes = 3 scope = spfile;



## Appendix C – DB2 Compress Homepage

Cut and paste the lines between “**Begin and End of Compress\_homepage\_DB2.sql – Do NOT copy this line**” into a file that you can call **Compress\_homepage\_DB2.sql**

From a DB2 command window on the DB2 server, issue the following commands:

```
set DB2INSTANCE=<DB2_instance_name>
db2 -td@ -vf Compress_HP_DB2.sql
```

**Note:** <DB2\_instance\_name> is the name of the DB2 instance serving the Homepage database. The default DB2 instance is called “DB2”. If you use multiple instances e.g. one for each database, you need to identify the one serving the Homepage database.

### Begin of Compress\_homepage\_DB2.sql – Do NOT copy this line

```
-- *****
--
-- Licensed Materials - Property of IBM
--
-- 5724_S68
--
-- Copyright IBM Corp. 2001, 2011 All Rights Reserved.
--
-- US Government Users Restricted Rights - Use, duplication or
-- disclosure restricted by GSA ADP Schedule Contract with
-- IBM Corp.
--
-- *****

CONNECT TO HOMEPAGE@

ALTER INDEX HOMEPAGE.PROFILES_READERS_DEL_SERV_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_C_CONTENT_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.STATUS_READERS_DEL_SERV_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NOT_REC_READERS_RLL_BRD_VVIS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.MENTIONS_READERS_COMM_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.USER_PREF_NAME_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PROFILES_VIEW_ITM_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.TOPICS_READERS_COMM_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.SR_INDEX_DOCS_ACS_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_SRC_TYPE_DEF_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.SR_INDEX_DOCS_FRID_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_ENTRIES_ITEM COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BRD_COMM_ITEM_COR_PUB COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NOT_REC_READERS_RLL_IIX COMPRESS YES@
```



IBM® Connections Version 5.5 – Social Software for Business IBM® Connections  
Performance Tuning Guide

```
COMMIT@
ALTER INDEX HOMEPAGE.WIKIS_READERS_CD_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.SEEDLIST_BRD_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PROFILES_VIEW_STR_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.FILES_READERS_ITM_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_OPT_TASK_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.ITEM_ID_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.EMD_EMAIL_PREFS_PER COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PERSON_USER_MAIL_LWR COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_LABEL_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_SL_STR_UNIQUE COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.EXTERNAL_READERS_ITM_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_NEWS_COMMENT_CONTENT_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.HP_WIDGET_INST_TAB_INST_ID_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.IMG_CACHE_URL_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.FILES_READERS_STR_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_AS_STATUS_PK COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BLOGS_READERS_ROLLUP_IIIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.WIKIS_READERS_DEL_SERV_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.FORUMS_READERS_ROLLUP_IIIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.TOPICS_READERS_DEL_SERV_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_ENTRIES_CONT COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.EXTERNAL_READERS_STR_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NOTIFICA_READERS_COMM_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_INDEX_TASK_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BRD_MEN_ITEM_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.RESPONSES_READERS_DEL_SERV_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.AGGREGATED_READERS_RLL_BRD_VVI COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.SR_FEEDBACK_PERSON_ID_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.UNIQUE_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_EMD_RES_PRF_PER_FRQ_RES_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_SL_UD_STR COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.EXTERNAL_READERS_COMM_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NOT_REC_READERS_COMM_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.UNIQUE_TASK_ID_FC COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PERSON_MEM_SAND_IX COMPRESS YES@
COMMIT@
```



## IBM® Connections Version 5.5 – Social Software for Business IBM® Connections Performance Tuning Guide

```
ALTER INDEX HOMEPAGE.TOPICS_READERS_RLL_IIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_PROPERTIES_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.ACTIVITIES_READERS_COMM_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.MENTIONS_READERS_ITM_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.TAB_INST_TAB_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.WIKIS_READERS_COMM_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NEWS_BRD_SL_UPDATE COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.FILES_READERS_RLL_BRD_VVIS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.CLNT_NAME_UK COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.UNIQUE_LCT_NAME COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BOOKMARKS_READERS_COMM_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_FOLLOWS_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BRD_RECOMMENDER_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_RESOURCE_TYPE_CONT COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_REPLYTO COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_SRC_TYPE_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.MENTIONS_READERS_STR_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_MT_CFG_DEFINITION COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_NEWS_SN_READER_NETW COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.ACTIVITIES_READERS_DEL_SERV_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.AGGREGATED_READERS_CD_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_SL_CD_DELETED COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.SR_NUMBER_STAT_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.UNIQUE_TASK_ID_BKP COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_SN_ITEM_READ_UNQ COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.SR_INDEX_DOCS_LLT4_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BLOGS_READERS_STR_RDR COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_CONTENTSTATS_PK COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_STORY_CONT_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NOT_REC_READERS_CD_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BRD_ENTRY_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PERSON_LAST_UPDATE COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BRD_E_CONTAINER_ID_UIDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NOT_REC_READERS_ROLLUP_IIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.COMMUNITIES_VIEW_STR_RDR COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.COMM_READERS_STR_RDR COMPRESS YES@
```



IBM® Connections Version 5.5 – Social Software for Business IBM® Connections  
Performance Tuning Guide

```
COMMIT@
ALTER INDEX HOMEPAGE.NR_SRC_TYPE_DEF_UIDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PROFILES_READERS_ROLLUP_IIIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BOOKMARKS_READERS_RLL_BRD_VVIS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_SL_UD_DELETED COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.DISCOVERY_VIEW_ITM_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.STATUS_READERS_ROLLUP_IIIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.SR_STRING_STAT_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.ITEM_CORR_CREATION_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PROFILES_READERS_RLL_BRD_VVIS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.DISCOVERY_VIEW_STR_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.ACTIONABLE_READERS_ITM_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.COMMUNITIES_VIEW_CD_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.COMM_READERS_CD_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_PROF_READERS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.CREATION_ITEM_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.AGGREGATED_READERS_COMM_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.TAB_INST_UI_ID_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.FILES_READERS_CD_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_AS_COUNT_PER_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.ACTIONABLE_READERS_STR_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_ACT_READERS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_NEWS_STATUS_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_BRD_ENTRIES COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.TOPICS_READERS_CD_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.SAVED_READERS_CD_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_CV_READERS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.RESPONSES_READERS_ITM_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_TRANCHE_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.EMD_RESOURCE_PREF_UNQ COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_COMM_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_COUNT_PER_VIEW_UNQ COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.WIKIS_READERS_ROLLUP_IIIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.TOPICS_READERS_ROLLUP_IIIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NOTIFICA_READERS_DEL_SERV_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_FOLLOWS_PERS COMPRESS YES@
COMMIT@
```



## IBM® Connections Version 5.5 – Social Software for Business IBM® Connections Performance Tuning Guide

```
ALTER INDEX HOMEPAGE.NR_SRC_DEF_UIDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.RESPONSES_READERS_STR_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.RESPONSES_READERS_ROLLUP_IIIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NOTIFICA_READERS_STR_RDR COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.REPLYTO_RECIP_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_PERSON_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.MENTIONS_READERS_RLL_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BRD_MEN_PER_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_WIDGET COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.EXTERNAL_READERS_RLL_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_BRD_COMMENT_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PERSON_ROLE_PER_R_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_SL_ITEM_ID_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BRD_AUTHOR_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.OEMBED_SCHEDULER_TASK_IDX1 COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.OEMBED_SCHEDULER_TASK_IDX2 COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.FORUMS_READERS_STR_RDR COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_ENTRIES_AR_ITEM COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_GSP_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.LOGINNAME_UNIQUE COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_NETWORK_PERS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NOTIFICA_READERS_ITM_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_SENT_NOT_READ COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.STATUS_READERS_ITM_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PROFILES_READERS_CD_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.ACTIVITIES_READERS_RLL_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.SR_FILESCONTENT_STATUS_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.UNQ_ORG_EXID
  COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.SAVED_READERS_COMM_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.SR_INDEX_DOCS_SUID_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_SRC_TYPE_ORG_UIDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NOTIFICA_READERS_STR_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.ACTIVITIES_READERS_ROLLUP_IIIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_ENTRIES_AR_CONT COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.STATUS_READERS_STR_IX COMPRESS YES@
COMMIT@
```



IBM® Connections Version 5.5 – Social Software for Business IBM® Connections  
Performance Tuning Guide

```
ALTER INDEX HOMEPAGE.NR_NEWS_STATUS_NETWORK_ITEM_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_BKUP_TASK_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PERSON_SND_SR COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_LOOKUPKEY COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PERSON_EXID_PER_ID_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_TRC_INFO_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_EMAIL_PREFS_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.CONTAINER_ID_UNQ COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.TAGS_READERS_RLL_IIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_FC_TASK_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_TMR_STAT_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_PERSON_ROLE_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.EXTERNAL_READERS_DEL_SERV_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PERSON_STATE_MEM_LU_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_ORGANIZATION_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PERSON_STATE_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PROFILES_VIEW_RLL_IIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.ACTOR_ENTRY COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_HP_WIDGET_INST COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_STATUS_UPDATE_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_ENTRY_AR_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NOTIFICA_READERS_RLL_BRD_VVIS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.STATUS_COMMENT_ITEM_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.RESPONSES_READERS_RLL_BRD_VVIS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.DISCOVERY_VIEW_RLL_IIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_COMM_READERS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.UNQ_SOURCE_TYPE_SRC_ORG COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BOOKMARKS_READERS_STR_RDR COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NEWS_STATUS_NET_R_U COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.SEQ_NUMBER_UNIQUE COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PROFILES_READERS_STR_RDR COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_TAGS_READERS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.UNQ_NAME COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PERSON_ROLE_ROLE_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.STATUS_READERS_STR_RDR COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.DISCOVERY_VIEW_ACT_CD_VIS_IDX COMPRESS YES@
```



IBM® Connections Version 5.5 – Social Software for Business IBM® Connections  
Performance Tuning Guide

```
COMMIT@
ALTER INDEX HOMEPAGE.NR_NEWS_SC_CD COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NOTIFICA_READERS_ROLLUP_IIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_FOLLOWS_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.ACTIONABLE_READERS_DEL_SERV_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.FORUMS_READERS_ITM_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_ACTION_READERS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_RESOURCE_TYPE_CONT_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_NETWORK_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.ACTIVITIES_READERS_RLL_BRD_VVI COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_BRD_MEN_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.UNIQUE_GSP_NAME COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.TOPICS_READERS_ITM_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BOARD_OWNER_ASSOC_UIDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.HP_WIDGET_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.FORUMS_READERS_STR_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NEWS_BRD_ITEM_CORR_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_COMM_FOLLOW_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BRD_REC_RECOM_ITEM_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PERSON_EXID_ORG COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.STATUS_READERS_RLL_BRD_VVIS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.TAGS_READERS_COMM_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.DISCOVERY_VIEW_CD_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_NEWS_COMMENT_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.TOPICS_READERS_STR_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.UNQ_SETTING COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_SR_STAT_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_BRD_RECOMM_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.RESPONSES_READERS_STR_RDR COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.SR_FILESCONTENT_VERSION_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.DISCOVERY_VIEW_DEL_SERV_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BOOKMARKS_READERS_CD_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_PREREQ COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_STATUS_COMMENT_DATE COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_HP_TAB_INST COMPRESS YES@
COMMIT@
```



IBM® Connections Version 5.5 – Social Software for Business IBM® Connections  
Performance Tuning Guide

```
ALTER INDEX HOMEPAGE.COLL_PERSON_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_HP_UI COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PROFILES_VIEW_RLL_BRD_VVIS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_ENTRY_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.FILES_READERS_DEL_SERV_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BRD_ENTRIES_ITEM COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.STORIES_ITEM_CORR_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BLOGS_READERS_CD_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_S_CONTENT_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.EMD_EMAIL_PREFS_TR COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_SL_UD_DELETED_VIS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NOT_REC_READERS_ITM_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_FBK_CTXT_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.UNIQUE_TASK_ID_ST COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.UNIQUE_TASK_ID_OPT COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.FILES_READERS_STR_RDR COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.SAVED_READERS_DEL_SERV_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_FBK_PARAMS_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NOT_REC_READERS_STR_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.SAVED_READERS_STR_RDR COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.DISCOVERY_VIEW_RLL_BRD_VVIS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PROFILES_READERS_COMM_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_SCHEDULER_TASK_IDX1 COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_SCHEDULER_TASK_IDX2 COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_NEWS_SN_READER_FOLL COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.EXTERNAL_READERS_ROLLUP_IIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.SAVED_READERS_ITM_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_FILESCONTENT_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BOARD_PK COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.UNIQUE_TASK_NAME COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.OEMBED_SCHEDULER_LMPR_IDX1 COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_METRIC_STAT_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BLOGS_READERS_RLL_BRD_VVIS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.TAGS_READERS_ITM_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.COMMUNITIES_VIEW_DEL_SERV_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.COMM_READERS_DEL_SERV_IX COMPRESS YES@
```



IBM® Connections Version 5.5 – Social Software for Business IBM® Connections  
Performance Tuning Guide

```
COMMIT@
ALTER INDEX HOMEPAGE.PK_MEN_READERS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.REPLYTO_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.SAVED_READERS_STR_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.LOTUSCONNECTIONSTASK_IDX1 COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.ACTIONABLE_READERS_COMM_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.LOTUSCONNECTIONSTASK_IDX2 COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.AGGREGATED_READERS_STR_RDR COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PERSON_DISPLAYNAME_LWR COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BLOGS_READERS_RLL_IIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PROFILES_READERS_ITM_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.SR_INDEX_CRAWL_VERSION_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.TAGS_READERS_STR_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.STORIES_ITEM_ID_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.USER_PREF_PER_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.ACTIONABLE_READERS_STR_RDR COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.AGGREGATED_READERS_ITM_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.DISCOVERY_VIEW_ORG_IIDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PERSON_ROLE_PER_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.HP_UI_R COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.FORUMS_READERS_RLL_BRD_VVIS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_PFS_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PROFILES_READERS_STR_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.UNIQUE_LABEL_NAME COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_RESP_READERS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.AGGREGATED_READERS_STR_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.MENTIONS_READERS_RLL_BRD_VVIS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.WIKIS_READERS_STR_RDR COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.COMMUNITIES_VIEW_RLL_IIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.COMM_READERS_RLL_IIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BLOGS_READERS_ITM_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.DELETED_STORY_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_AGG_READERS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_NETWORK_PER_COL_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_REC_NOT_READ COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.EMD_RES_PREF_PER_ID COMPRESS YES@
COMMIT@
```



IBM® Connections Version 5.5 – Social Software for Business IBM® Connections  
Performance Tuning Guide

```
ALTER INDEX HOMEPAGE.WIKIS_READERS_ITM_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.ACTIONABLE_READERS_ROLLUP_IIIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BLOGS_READERS_STR_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.MENTIONS_READERS_DEL_SERV_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.STORIES_CONTENT_STORY COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.WIKIS_READERS_STR_IX COMPRESS
YES@
COMMIT@
ALTER INDEX HOMEPAGE.SAVED_READERS_RLL_BRD_VVIS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BOOKMARKS_READERS_DEL_SERV_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_STR_STAT_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_COMM_PERSON_F COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.UNIQUE_STAT_KEY COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_HP_TAB COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NOT_REC_READERS_STR_RDR COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_STORY_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.DISCOVERY_VIEW_ROLLUP_IIIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_COMM_PER_FOLLOW_PER_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NOTIFICA_READERS_RLL_IIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.DISCOVERY_VIEW_COM COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_PROFILES_VIEW COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_TOPICS_READERS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.FILES_READERS_ROLLUP_IIIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_MIGTASKDEF_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.OH2P_CACHE_CNT COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_RESOURCE_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.UNQ_ORG_NAME COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.FORUMS_READERS_RLL_IIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.TAGS_READERS_DEL_SERV_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.SAVED_READERS_ROLLUP_IIIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.SR_INDEX_DOCS_TAW_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.FORUMS_READERS_CD_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PERSON_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PERSON_LAST_CONN_VISIT COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_DISCOVERY_VIEW COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_STORIES_REL_COMM COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_TEMPLATE_ID COMPRESS YES@
COMMIT@
```



IBM® Connections Version 5.5 – Social Software for Business IBM® Connections  
Performance Tuning Guide

```
ALTER INDEX HOMEPAGE.PROFILES_VIEW_DEL_SERV_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.HP_TAB_INST_UNQ COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_FILES_READERS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.RESPONSES_READERS_COMM_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.MENTIONS_READERS_CD_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.TOPICS_READERS_STR_RDR COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_NUM_STAT_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.COMMUNITIES_VIEW_ITM_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.COMM_READERS_ITM_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.ACTIONABLE_READERS_RLL_BRD_VVI COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.COMMUNITIES_VIEW_ROLLUP_IIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.COMM_READERS_ROLLUP_IIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_SAVED_READERS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_STORIES_DATE COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.DISCOVERY_VIEW_COMM_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.OA2T_GBINDING_WID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.COMMUNITIES_VIEW_STR_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.COMM_READERS_STR_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PERSON_SAND_OPT_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.COMMUNITIES_VIEW_RLL_BRD_VVIS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.COMM_READERS_RLL_BRD_VVIS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_AS_CONFIG_PK COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.SETTINGS_BY_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_FORUMS_READERS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_SC_UPDATE COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_BRD_OBJ_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_QUARANTINE COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_ST_TASK_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.STORIES_ITEM_ENTRY_CORR_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PERSON_SAND_LAST_UPDATE COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_STORY_CD_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_TOKEN_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.STATUS_READERS_COMM_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.ITEM_ITEM_CORR_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BRD_SL_UPDATED_DEL COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_WIKIS_READERS COMPRESS YES@
```



IBM® Connections Version 5.5 – Social Software for Business IBM® Connections  
Performance Tuning Guide

```
COMMIT@
ALTER INDEX HOMEPAGE.NEWS_BRD_ITEM_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.MENTIONS_READERS_ROLLUP_IIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_SCHEDULER_LMPR_IDX1 COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.UNIQUE_RES COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_MT_CFG_SETTINGS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PROFILES_VIEW_ORG_IIDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.TAGS_READERS_CD_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BOOKMARKS_READERS_RLL_IIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.AGGREGATED_READERS_DEL_SERV_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PERSON_EMD_CHECK COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_AS_COUNT_ORG_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.SRC_TYPE_ORG_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BOOKMARKS_READERS_ROLLUP_IIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.UNIQUE_TASK_ID_IND COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PROFILES_READERS_RLL_IIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PROFILES_VIEW_CD_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.LOTUSCONNECTIONSLMPR_IDX1 COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.COMMUNITIES_VIEW_COMM_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.COMM_READERS_COMM_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BRD_REC_ITEM_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.WIDGET_GADGET_UK COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.STATUS_READERS_RLL_IIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.STATUS_READERS_CD_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.FILES_READERS_COMM_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.SR_FEEDBACK_CLIENT_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BRD_CURRENT_STATUS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.SR_INDEX_DOCS_SUDCRA_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_NEWS_SC_ITEM_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.EXTERNAL_READERS_CD_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.CURRENT_STATUS_INDEX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.USER_PREF_PER_NAME_UNQ COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.ACTIONABLE_READERS_CD_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_COMPIDCLIENTID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_BLOGS_READERS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_STORIES_EIDX COMPRESS YES@
COMMIT@
```



IBM® Connections Version 5.5 – Social Software for Business IBM® Connections  
Performance Tuning Guide

```
ALTER INDEX HOMEPAGE.WIDGET_INST COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.TAGS_READERS_ROLLUP_IIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.CL_STATUS_COMMENT_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.MENTIONS_READERS_STR_RDR COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_EXT_READERS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.RESPONSES_READERS_RLL_IIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_COMM_FOLLOW_COM_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.EXTERNAL_READERS_STR_RDR COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.RESPONSES_READERS_CD_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_STATUS_NETWORK_DATE COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PROFILES_VIEW_ROLLUP_IIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_NEWS_SN_READER_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.SR_INDEX_DOCS_RPS_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.CREATION_DATE_ITEM_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.EXTERNAL_READERS_RLL_BRD_VVIS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_STORIES_ER_UUID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.FORUMS_READERS_COMM_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BLOGS_READERS_DEL_SERV_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_INDEX_MGMT_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_FOLLOWS_RES COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.FORUMS_READERS_DEL_SERV_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.METRIC_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.ACTIVITIES_READERS_STR_RDR COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_AS_COUNT_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.FILES_READERS_RLL_IIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.UNIQUE_PERS_P_COMM COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.OA2T_PERSON_SERVICE COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_CUR_ST_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.SR_INDEX_DOCS_ACT_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.UNIQUE_SERVICE_NAME COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.SAVED_READERS_RLL_IIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_LOGINNAME_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.WIKIS_READERS_RLL_BRD_VVIS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PROFILES_VIEW_COMM_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_SEEDLIST_PK COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.STORY_CONTAINED_ID COMPRESS YES@
```



IBM® Connections Version 5.5 – Social Software for Business IBM® Connections  
Performance Tuning Guide

```
COMMIT@
ALTER INDEX HOMEPAGE.PERSON_MEM_STATE_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_RES_PREF_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PRVD_NAME_UK COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.TAGS_READERS_STR_RDR COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_WIDGET_TAB COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.AGGREGATED_READERS_RLL_IIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.HP_UI_PERSONID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.DISPLAYNAME_LOWER_MEM COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.UNIQUE_PERS_COMM COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.TOPICS_READERS_RLL_BRD_VVIS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.ACTIVITIES_READERS_ITM_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.AGGREGATED_READERS_ROLLUP_IIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.ACTIONABLE_READERS_RLL_IIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PROFILES_VIEW_STR_RDR COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_INDEX_DOCS_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_COMM_FOLLOW_PERSON_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_FEEDBACK_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_TASK_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NOTIFICA_READERS_CD_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_SU_READERS
  COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_USER_PREFS_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BOOKMARKS_READERS_ITM_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.ACTIVITIES_READERS_STR_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.WIKIS_READERS_RLL_IIX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.OA2_TOKEN_SHA1_UNQ COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BLOGS_READERS_COMM_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.IMG_CACHE_DEL_BATCH_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_MT_CFG_FILE COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.SR_TIMER_STAT_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.ACTIVITIES_READERS_CD_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_COMM_PER_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.CREATION_DATE_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.IMG_CACHE_FILE_UNQ COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.BOOKMARKS_READERS_STR_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_COMM_FOLLOW_COM_PER_ID COMPRESS YES@
```



IBM® Connections Version 5.5 – Social Software for Business IBM® Connections  
Performance Tuning Guide

```
COMMIT@
ALTER INDEX HOMEPAGE.BRD_IS_DEL_CDATE_ACT_ENT COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_BOOKS_READERS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.TAGS_READERS_RLL_BRD_VVIS COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.PK_IMG_CACHE_ID COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.STORIES_EVENT_ITEM_ACTOR_IDX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NOT_REC_READERS_DEL_SERV_IX COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.NR_STORIES_CONTENT_DATE COMPRESS YES@
COMMIT@
ALTER INDEX HOMEPAGE.UNIQUE_COMP_UUID COMPRESS YES@
COMMIT@
```

-- -----

```
ALTER TABLE HOMEPAGE.HOMEPAGE_SCHEMA COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.MT_ORGANIZATION COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.PERSON COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.PERSON_ROLE COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.HP_USER_PREFS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.LOGINNAME COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.HP_UI COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.HP_TAB COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.HP_TAB_INST COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.WIDGET COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.HP_WIDGET_TAB COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.HP_WIDGET_INST COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.PREREQ COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NT_REPLYTO COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NT_REPLYTO_RECIPIENT COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.METRIC_STAT COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.OAUTH1_TOKEN COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.OAUTH1_PROVIDER COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.OAUTH1_CLIENT COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.OAUTH1_CONTEXT COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.OAUTH2_PROVIDER COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.OAUTH2_CLIENT COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.OAUTH2_GADGET_BINDING COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.OAUTH2_TOKEN COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.OH2P_CACHE COMPRESS YES@
```



IBM® Connections Version 5.5 – Social Software for Business IBM® Connections  
Performance Tuning Guide

```
COMMIT@
ALTER TABLE HOMEPAGE.OH2P_CLIENTCFG COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.MTCONFIG COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.MT_CFG_DEFINITIONS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.MT_CFG_SETTINGS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.MT_CFG_FILES COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_SOURCE_TYPE_DEFAULT COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_SOURCE_TYPE COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_TEMPLATE COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_SCHEDULER_TASK COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_SCHEDULER_TREG COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_SCHEDULER_LMGR COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_SCHEDULER_LMPR COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_NEWS_STATUS_NETWORK COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_NEWS_STATUS_COMMENT COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_NEWS_STATUS_CONTENT COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_NEWS_COMMENT_CONTENT COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_RESOURCE COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_FOLLOWS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_COMM_FOLLOW COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_STORIES_CONTENT COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_NETWORK COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.EMD_RESOURCE_PREF COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.EMD_TRANCHE COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.EMD_TRANCHE_INFO COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.EMD_EMAIL_PREFS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_COMM_PERSON_FOLLOW COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_ENTRIES COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_STORIES COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_ENTRIES_ARCHIVE COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_AGGREGATED_READERS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_RESPONSES_READERS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_PROFILES_READERS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_COMMUNITIES_READERS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_ACTIVITIES_READERS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_BLOGS_READERS COMPRESS YES@
COMMIT@
```



## IBM® Connections Version 5.5 – Social Software for Business IBM® Connections Performance Tuning Guide

```
ALTER TABLE HOMEPAGE.NR_BOOKMARKS_READERS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_FILES_READERS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_FORUMS_READERS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_WIKIS_READERS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_TAGS_READERS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_STATUS_UPDATE_READERS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_EXTERNAL_READERS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_ACTIONABLE_READERS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.BOARD COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.BOARD_ENTRIES COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.BOARD_COMMENTS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.BOARD_OBJECT_REFERENCE COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.BOARD_RECOMMENDATIONS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.BOARD_CURRENT_STATUS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.BOARD_MENTIONS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.DELETED_STORIES_QUEUE COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_COMM_SETTINGS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_DISCOVERY_VIEW COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_PROFILES_VIEW COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_NOTIFICATION_SENT_READERS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_NOTIFICATION_RECEIV_READERS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_SAVED_READERS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_TOPICS_READERS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_MENTIONS_READERS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_COMMUNITIES_VIEW COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_AS_SEEDLIST COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_AS_COLLECTION_CONFIG COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_AS_CRAWLER_STATUS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_AS_CONTENT_INDEX_STATS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.IMG_CACHE COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.NR_AS_COUNTS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.OEMBED_SCHEDULER_TASK COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.OEMBED_SCHEDULER_TREG COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.OEMBED_SCHEDULER_LMGR COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.OEMBED_SCHEDULER_LMPR COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.SR_TASKDEF COMPRESS YES@
```



IBM® Connections Version 5.5 – Social Software for Business IBM® Connections  
Performance Tuning Guide

```
COMMIT@
ALTER TABLE HOMEPAGE.SR_INDEXINGTASKDEF COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.SR_OPTIMIZETASKDEF COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.SR_FILECONTENTTASKDEF COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.SR_BACKUPTASKDEF COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.SR_SANDTASKDEF COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.SR_FILESCONTENT COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.SR_MIGTASKDEFINFO COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.SR_INDEX_MANAGEMENT COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.SR_RESUME_TOKENS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.SR_INDEX_DOCS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.SR_FEEDBACK COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.SR_FEEDBACK_CONTEXT COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.SR_FEEDBACK_PARAMETERS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.SR_STATS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.SR_STRING_STATS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.SR_NUMBER_STATS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.SR_TIMER_STATS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.SR_GLOBAL_SAND_PROPS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.SR_ECM_DOCUMENT_TYPE_LABELS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.SR_POST_FILTERING_SERVICE COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.SR_ECM_DOCUMENT_TYPE_PROPS COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.LOTUSCONNECTIONSTASK COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.LOTUSCONNECTIONSTREG COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.LOTUSCONNECTIONSLMGR COMPRESS YES@
COMMIT@
ALTER TABLE HOMEPAGE.LOTUSCONNECTIONSLMPR COMPRESS YES@
COMMIT@
```

```
COMMIT@
FLUSH PACKAGE CACHE DYNAMIC@
connect reset@
terminate@
```

**End of Compress\_homepage\_DB2.sql .sql – Do NOT copy this line**



## Appendix D - Credits and Feedback

Thanks to the following people for their work in creating this document:

- John Mulroy
- Martin Presler-Marshall
- Roger Kutshina
- Ryan Gallagher
- Shane McCarthy
- Simon Pizzoli
- Terence A. Walker
- Thomas O'Dwyer
- Tom Healy
- Trinesh Andhurthi

Please direct any feedback on this document to the editor, Simon Pizzoli. Comments can be sent by e-mail to [simon.pizzoli@ie.ibm.com](mailto:simon.pizzoli@ie.ibm.com)



## 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](http://www.ibm.com/legal/copytrade.shtml) (www.ibm.com/legal/copytrade.shtml).

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

Linux is a registered 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, or service names may be trademarks or service marks of others.

## Terms and Conditions

Permissions for the use publications is granted subject to the following terms and conditions.

**Personal Use:** You may reproduce these publications for your personal, non commercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

**Commercial Use:** You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and



regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

