



Idle WebSphere Tuning Considerations

Now updated for WebSphere v8.5.5



WebSphere software

Beena Hotchandani, IBM
Performance Analyst

Stephen Kinder, IBM
Senior Technical Staff, WebSphere Virtualization Architect

Table of Contents

Overview:	3
Introduction:	3
Testing Methodology and Configuration	4
WebSphere Idle Server Tuning Considerations	4
Application Tuning Options.....	5
Application Server Tuning Options	6
Node Agent Tuning Options.....	9
High Availability Manager and Intelligent Management Tuning Options.....	10
Liberty Profile Tuning	12
Conclusion:	13
Acknowledgements:.....	13
Document History:	13

Overview:

Over the years, WebSphere Application Server has been engineered to keep itself ready and optimized for incoming work; doing housekeeping when possible in the background. When WebSphere Application Server is deployed on a single physical server, this is effective use of “whitespace” or unused capacity on the box to ensure that real requests are processed optimally. More recently, server consolidation and cloud projects attempt to leverage this same “spare” capacity to run other workloads. Server consolidation and cloud scenarios typically use virtualization to *over-commit* resources such as CPU and memory – a condition where the amount of virtual resources defined exceeds the actual physical resources available. In scenarios where WebSphere is sharing resources with other workloads, trade-offs may need to be made between optimal “WebSphere” configuration and optimal usage of the overall system resource.

This paper provides a list of WebSphere configurable options that can be used to reduce background processing within the application server and other components in the Network Deployment (ND) infrastructure (i.e. Node-agent, DMGR, etc.) that consume CPU cycles when WebSphere would typically be considered “idle”. The paper details some of the tuning options, their implications and some additional recommendations, which may help conserve resources during idle periods.

Tuning for minimizing WebSphere startup times are also presented. WebSphere server start typically is CPU intensive and can have adverse effects on other workloads running in the same shared environment. Server startup times are also relevant to idle workloads when Virtual Enterprise is used to manage dynamic clusters, and also when many environments need to be restarted at once – such as after a failure, or during periodic recycles of infrastructure. Since the initial writing of this paper, the features related to dynamic clustering, formally offered in Virtual Enterprise, are now included in WebSphere Application Server Network Deployment.

In addition, information on tuning the Liberty Profile has also been included. The Liberty Profile of WebSphere has been specifically designed with virtualization in mind, as fit-for-purpose and lazily initialized components yield significantly better memory and startup profiles which, of course benefit server consolidation and cloud multi-tenant environments.

Introduction:

WebSphere provides a highly resilient application server infrastructure, complete with fail-over, high availability, distributed cache management, and most recently dynamic clustering capabilities. This resiliency and robustness comes with a price. WebSphere schedules timed events to check for availability of other servers, manage caches and otherwise keep the entire infrastructure up-to-date and running flawlessly. There are several configuration options, which can be leveraged to minimize background processing during idle periods. In production environments like those seen in server consolidation deployments, careful examination, understanding, and configuration of these features can help the underlying infrastructure manage the shared resources more effectively.

Testing Methodology and Configuration

Performance Measurements for tuning the idle server were done in a virtual environment, using z/VM and running WebSphere Application Server v8.5.5.5 on System z Linux (SLES 11.2). Even though these tests were performed under a z/VM virtual environment, the tuning recommendations apply to WebSphere Application Server under any shared environment, native operating system or any other virtualization platform including Pure App Systems, PowerVM, VMWare, KVM, etc.

The test configuration consisted of one Deployment Manager, two Node Agents, and three application servers with the DayTrader 1.2 Benchmark installed. DayTrader is a Java Enterprise Edition (Java EE) benchmark modeling an online stock brokerage application.

WebSphere Idle Server Tuning Considerations

The following table describes the configuration options that were tested. Choosing an appropriate setting for each option may involve a functional trade-off. Consequently, each option should be carefully considered in your environment to ensure that required capabilities are not fully or partially disabled. Furthermore, each option should be thoroughly tested with your application and environment to ensure that both functional and performance expectations are met.

	<i>Our Test Results</i>	<i>Pros</i>	<i>Cons</i>	<i>Detail page</i>
Application Tuning				
Class loading and update detection	Measureable reduction of CPU consumption during idle periods.	Idle CPU reduction.	Manual update required to reflect application changes	Page 5
JSP and JSF options	Although our application did not have a JSP, we have seen measureable improvements in customer environments.	Idle CPU reduction.	Manual update required to reflect application change.	Page 6
Application Server Tuning				
Start components as needed	Measureable reduction of CPU during startup.	Startup CPU Reduction.	Some component level startup processing is deferred, until first usage	Page 6
Dynamic cache service background processing	Not measurable	Reduces frequency of batched updates and reduces scan time for unused objects.	Possibly some impact (see below)	Page 7



Node Agent Tuning				
Automatic File Synchronization	Measureable Idle CPU Reduction	Eliminates node sync polling for nodes which are out of sync.	Manual synchronization required.	Page 9
High Availability Manager and Intelligent Management Tuning				
Core group service	Measureable Idle CPU Reduction	Reduces Overall CPU consumption.	Possible functional risk, if HA features are required, depending on your overall infrastructure topology.	Page 10
Intelligent Management Disable Intelligent Mgt. in WAS 8.5.5	Measurable CPU Reduction	Idle CPU equivalent to 8.0 (without Intelligent Management)	No Intelligent Management functionality	Page 10
Liberty Profile Tuning				
Turn Off File monitoring	Measurable CPU reduction	Reduces Idle CPU considerably	No functionality is lost but if changes are made, the server needs to be restarted.	Page 12

Application Tuning Options

This section discusses WebSphere Application tuning properties that can be used to reduce CPU consumption during idle periods that have an effect on application behavior. They effect how often the WebSphere Application server will check for updates to EJBs, web modules, and JSPs. In a production environment where changes to these components are made infrequently, this level of monitoring may not be necessary and can be reduced or eliminated.

Class loading and update detection

Description: WebSphere Application Server provides the ability to control how often application classes in EJB and web modules are polled for updates. For large applications, this scanning can consume considerable CPU cycles on a frequent basis if left unchanged.

Consideration(s): This feature is primarily needed in development environments (i.e. Rational Application Developer) where application updates are expected frequently. In production environments where individual class file updates are very infrequent, update checking can safely be disabled. This can be done by setting the class reloading setting for Web and EJB

modules to true and by setting the polling interval to zero as described in the WebSphere for Application Server InfoCenter at:

http://www14.software.ibm.com/webapp/wsbroker/redirect?version=phil&product=was-nd-zos&topic=urun_rapp_classload

Trade off(s): If this feature is disabled, the application must be stopped and restarted manually to reflect changes made in the updated classes.

Java Server Pages (JSP) Reloading

Description: WebSphere Application Server provides the ability to translate and compile JavaServer Pages (JSPs) at runtime when the JSP or its dependencies are modified. This is known as JSP reloading. This feature can be disabled by setting the reloadEnabled JSP engine parameter to false in the WEB-INF/ibm-web-ext.xmi or WEB-INF/ibm-web-ext.xml file, as described in the WebSphere Application Server InfoCenter :

http://www14.software.ibm.com/webapp/wsbroker/redirect?version=phil&product=was-nd-zos&topic=rweb_jspreloading

Consideration(s): This feature is primarily intended for development environments (i.e. RAD) where frequent application updates are expected. In production environments where JSP updates are much less frequent, update checking can be safely disabled.

Trade off(s): If this feature is disabled, the application will have to be stopped and restarted manually to reflect changes made to JSP(s)

Application Server Tuning Options

WebSphere Application Server provides a variety of settings that can be used to configure caching and startup configuration of an application server, cluster as well as the Java Virtual Machine which can effect how well WAS performs in a virtualized environment. This section discusses these settings and their effects on idle CPU consumption.

Start components as needed

Description: This application server property controls how components and services are started within the application server. If enabled, components are started as they are needed by the application. Otherwise, all components are started during the server startup phase. Details on how to select this property can be found at:

http://www14.software.ibm.com/webapp/wsbroker/redirect?version=phil&product=was-nd-zos&topic=urun_rappsvr

Consideration(s): Enabling this feature allows server components and services to start as needed. For instance if the applications installed on a server do not leverage EJBs or SIP, the EJB and SIP components will not startup. Selecting this option can improve startup time and

reduce the memory footprint of the server, because fewer components are started during the startup process. With this feature, it is important to group similar apps together on a server based on the components they will use (and therefore start).

Trade off(s): Customers may want to consider “warming up” an application to avoid any delay that might occur when the first application unit of work is processed.

Dynamic cache service background processing

Description: The dynamic cache service is used to provide generalized caching services for servlets, portlet fragments, web services, application data, etc. Several background processing tasks are needed to manage these caches effectively. Since the dynamic cache service cannot be disabled in Version 7, these background tasks are always active even if the dynamic cache service is not being used. Following is the list of custom properties that can be used to reduce the overhead of these processes.

com.ibm.ws.cache.CacheConfig.batchUpdateMilliseconds: This option controls the frequency of batched cached updates between distributed caches in a cluster via DRS (Data Replication Service). The default is 1000 ms.

com.ibm.ws.cache.CacheConfig.timeGranularityInSeconds:

com.ibm.ws.cache.CacheConfig.lruToDiskTriggerTime: These two settings control how frequently the caches are checked for expiring entries. The `lruToDiskTriggerTime` should always be a multiple of the `timeGranularityInSeconds` setting. The default setting for both is 5 seconds.

com.ibm.ws.cache.CacheConfig.timeHoldingInvalidations: This option controls how long cache invalidations are held in memory before being committed to the cache. The default is 300,000 ms.

Detailed instructions on how to set these properties can be found on the InfoCenter at:

http://www14.software.ibm.com/webapp/wsbroker/redirect?version=phil&product=was-nd-zos&topic=rdyn_tunediskcache

Consideration(s): For Deployment Manager and Application Server instances where the dynamic caching service is not specifically being used, these parameters can be configured to higher values to reduce the number of associated background operations.

Before making permanent modifications to these custom properties you should:

1. Verify the dynamic cache service is not in use by enabling the dynamic cache service trace and monitoring the trace logs for activity or by installing and monitoring the Cache Monitor application in a test environment.

2. Thoroughly test the application and environment from a functional and performance perspective to ensure no impact prior to making these configuration changes in production.

Trade off(s): If the dynamic caching service is actually being used by your application, increasing these values could have negative impacts on performance and/or cache behavior and you may experience out of memory conditions.

EJB cache and pool background processing

Description: The WebSphere EJB Container maintains a number of caches and pools for performance reasons. Much like the dynamic cache service, these caches and pools are monitored and cleaned frequently by background tasks. The following properties can be modified to reduce or eliminate the overhead of these processes.

Cleanup interval: Default 3000 (ms)

This is the interval at which EJB entity caches are scanned for unused objects that can be removed from the cache. Details on how to set this property can be found on the InfoCenter at:

<http://www14.software.ibm.com/webapp/wsbroker/redirect?version=phil&product=was-nd-zos&topic=uejbecnt>

Inactive pool cleanup interval: Default 30000 (ms)

This is the interval at which the EJB Stateless Session Bean pools are scanned for inactive instances that can be removed from the pool. Details on how to set this property can be found at:

<http://www14.software.ibm.com/webapp/wsbroker/redirect?version=phil&product=was-nd-zos&topic=uejbcash>

com.ibm.websphere.bean.delete.sleep.time: Default 4200 (sec). *z/OS-Only.*

This is the interval at which Stateful Session Bean pool is scanned for timed out beans instances. Note that this setting can only be modified on z/OS. Details on how to set this property can be found at:

http://www14.software.ibm.com/webapp/wsbroker/redirect?version=phil&product=was-nd-zos&topic=xrun_jvm

Considerations(s): For DMGR and App Server instances where the EJB container is not specifically being used, these parameters can be configured to higher values to prevent the associated background operations from occurring frequently (as dictated by the defaults). Increasing these values to extremely high levels will essentially disable this behavior.

Before making permanent modifications to these properties you should:

1. Verify that the EJB container is not in use by enabling the EJB container trace and monitoring the trace logs for activity in a test environment.
2. Thoroughly test the application and environment from a functional and performance perspective to ensure no impact.

Trade off(s): If the EJB container service is being used by your application, increasing these values could have negative impacts on performance and/or container behavior which may cause out of memory conditions.

Node Agent Tuning Options

This section briefly describes the Automatic File Synchronization function and how to modify property settings that can reduce the overhead of this function.

Automatic File Synchronization

Description: This function controls whether or not synchronization is performed automatically between the Node Agents and DMGR and how often that synchronization is performed.

Considerations(s):

1. *Security*

The security run time depends on node synchronization to propagate updated certificates during automated replacement processes. The security runtime also depends on node synchronization for Lightweight Third Party Authentication (LTPA) key changes.

If you are *not* using LTPA key regeneration, then you may consider turning off or increasing the file synchronization interval. Automatic LTPA key regeneration requires a very short node sync interval at the time the keys are regenerated, or you could cause a production outage. LTPA keys are not like certificates, where there can be multiple concurrent certs. It's possible for one process to present a new LTPA key that another process knows nothing about. So it's crucial that all the keys are in sync. Since any number of things can go wrong when keys are regenerated, we recommend that you disable the auto-regeneration feature if you increased the node synch interval.

2. *File synchronization for application deployment*

If node synchronization is disabled, you need to sync to all nodes after making a configuration change, either manually through the admin console, or programmatically through scripting. The main reason is that Serverindex.xml is changed whenever you create new servers, cluster members, or deploy applications. If it's not propagated everywhere, it can prevent HAManager from stabilizing, especially when you add new cluster members. So you should have a best practice to sync to all nodes after a configuration change, or rely on short node sync interval.

It may be beneficial to *temporarily* disable node sync during application update in order to control when the nodes receive

Details on how to modify this setting can be found at:

http://www14.software.ibm.com/webapp/wsbroker/redirect?version=phil&product=was-nd-zos&topic=urun_rsynchservice

Recommendation: If you change the node synchronization interval to something other than the default, tune to an interval of 10 minutes maximum.

Trade off(s):

- You *cannot* use automatic LTPA key exchange if you permanently disable automatic file synchronization.
- If automatic synchronization is disabled (*not recommended*), forced synchronizations must be performed via the admin console or wsadmin scripting once config changes have been made. Otherwise changes will not be reflected on the relevant nodes.
- As the WebSphere Application Server runtime evolves other components may rely on node synchronization, and disabling node synchronization could result in a breakage and a loss of service.

The WebSphere Application Server Information Center article cited above has this important notice about node synchronization intervals:

Important: Do **not** routinely disable the synchronization process as various portions of the application server run time depend on synchronization. For example, the security run time depends on node synchronization to propagate updated certificates during automated replacement processes. Also, the security runtime also depends on it for Lightweight Third Party Authentication (LTPA) key changes. Other portions of the run time are dependent on synchronization. However, a complete list is not available. If you permanently disable synchronization, nodes might not be synchronized and an outage will result. The only exception is the configuration save operation. To avoid synchronizing the configuration when the configuration repository is being updated by a save operation, it might be beneficial to temporarily disable the synchronization process, save the configuration, and then re-enable the synchronization process. This process ensures that changes are fully committed to the configuration repository before the node synchronization.

High Availability Manager and Intelligent Management Tuning Options

High Availability Manager Description: A High Availability Manager instance runs on every application server, proxy server, node agent and deployment manager in a cell. A cell can be divided into multiple high availability domains known as core groups. Each high availability manager instance establishes network connectivity with all other high availability manager instances in the same core group, using a specialized, dedicated, and configurable transport channel. The transport channel provides mechanisms that allow the high availability manager instance to detect when other members of the core group start, stop, or fail. The

amount of CPU and memory resources consumed by the HA manager and associated components increases as the size of a core group increases. The procedure used to disable the HA Manager can be found at:

http://www14.software.ibm.com/webapp/wsbroker/redirect?version=phil&product=was-nd-zos&topic=trun_ha_ham_enable

Also refer to Configuring the Core Group Bridge Service at:

http://www14.software.ibm.com/webapp/wsbroker/redirect?version=phil&product=was-nd-zos&topic=trun_ha_coregroupbridge

Consideration(s): If none of the HA manager related services are being leveraged by your application or environment, the service can be disabled to eliminate the associated overhead. Depending on the topologies and services being used, the service can be disabled selectively among the DMGR, node agents and app server processes.

If you are considering disabling the HA Manager, it is strongly suggested that you review the WebSphere documentation at the link below carefully and perform adequate functional, behavioral, and performance testing to ensure that the disabled services do not negatively impact your application.

http://www14.software.ibm.com/webapp/wsbroker/redirect?version=phil&product=was-nd-zos&topic=crun_ha_ham_required

Trade off: Because of the number of WebSphere services and features linked to the HA manager, disabling this component may involve a certain degree of risk. Furthermore, some extended WebSphere technologies or stack products may also have dependencies on the HAManager. For instance, WVE v6.1.1.0 and releases prior require the HAManager features to be enabled.

Intelligent Management Description: The Virtual Enterprise product was integrated in WAS 8.5 and is now referred to as Intelligent Management. It is enabled by default. This might increase idle server CPU time considerably. Unfortunately there is no option in WAS 8.5 to disable it.

In WAS 8.5.5, a new custom property (`LargeTopologyOptimization`) was added to disable Intelligent Management for those who do not to use the functionality. Details on how to disable Intelligent Management is found in the following link:

http://www14.software.ibm.com/webapp/wsbroker/redirect?version=phil&product=was-nd-zos&topic=ragt_cell_customprops

Considerations: Disabling Intelligent Management will cause no loss in any other functionality except that Intelligent Management functions are disabled.

Trade off: No other tradeoff except we lose Intelligent Management functions

Liberty Profile Tuning

Description: Turn off File Monitoring – the Liberty profile is fully capable of responding to any changes in the configuration, application and other resources dynamically. There is basically polling going on to check if any changes being made. The updates are done automatically and server restart is not needed.

Considerations: Turning off file monitoring does not have any impact on the functionality and since the server will no longer poll for changes, there is a reduction in CPU cycles. Details on how to disable file monitoring can be found in:

http://pic.dhe.ibm.com/infocenter/wasinfo/v8r5/topic/com.ibm.websphere.wlp.zseries.doc/ae/twlp_setup_dyn_upd.html

Trade off: If any updates are made in the configuration, the server should be restarted

Conclusion:

As larger shared environments become commonplace, it is important to review your WebSphere configuration to ensure that you've adequately evaluated its effect primarily when WebSphere is idle. It was our experience that the biggest hitters were Class Loading and update detection at the application level, automatic file synchronization, and disabling HA Manager, while others used in combination provided additional savings. Turning off unneeded features and dialing back background process intervals can be beneficial to the overall resource consumption of your entire configuration and help WebSphere participate better as part of your infrastructure.

Acknowledgements:

The following people are gratefully acknowledged for their contributions to this paper.

Thomas Alcott, IBM Worldwide WebSphere Application Infrastructure
 Chris Blythe
 Jim Cunningham, IBM WebSphere Performance on z platform
 Tom Seelbach, IBM WebSphere SCA Development
 Steve Wehr, IBM systems & Technology Group, System z Linux Infrastructure
 Bruce Hayden, IBM Advance Technical Skills – z/VM System z Linux
 Romney White, IBM System z Virtualization Technology
 Andrew Low, IBM JVM development
 Juergen Doelle IBM, System z Linux performance

Document History:

April 6, 2011	Original document.
May 6, 2011	Updated node synchronization section to further clarify the importance of not disabling the node synchronization function. A tuning option is to extend the synchronization interval perhaps. But never disable entirely.
August 1, 2013	Updated for WAS v8 and v8.5.5 including the Liberty Profile

End of Document