September 2011

z/OS Management Facility V1.13 Resource Requirements

Judi Bank
IBM Systems and Technology Group, Systems Hardware Development
STSM, z/OS Performance Software Performance Analyst
1-919-486-0941
judib@us.ibm.com

Measurement data captured by Kumar Chintapalli, IBM Software Test Specialist

Table of Contents

Introduction

The z/OS® Management Facility (z/OSMF) is a product that provides a browser based interface for system management for z/OS customers. It was first introduced in z/OS V1.11 and runs on z/OS V1.10 and above. It is designed for systems programmers to more easily manage z/OS by simplifying administration and day to day operation. It is especially useful for increasing the productivity of new or experienced systems programmers. It features integrated wizards and help panels. z/OSMF is offered as a separately licensed product to z/OS customers at no additional charge. The program number for z/OS Management Facility is 5655-S28. This document describes z/OSMF V 1.13.

z/OSMF is implemented as a Web 2.0 application using a special version of WebSphere® Application Server called WebSphere Application Server OEM Edition for z/OS. WAS OEM is shipped with z/OSMF at no additional charge. Depending on the system management task to be performed, large portions of z/OSMF are executed in Java, and are therefore eligible for IBM System z® specialty engines. In addition, z/OSMF interfaces with other z/OS components including the Common Information Model (CIM) which is also shipped with z/OS. z/OSMF runs entirely on a z/OS host. It does not require any code on the user's client system other than a supported browser attached via a secure connection.

z/OSMF Hardware Requirements

z/OSMF requires a certain amount of CPU resource and memory for satisfactory operation during some z/OSMF scenarios that consume more CPU time than others. The target system for z/OSMF should have:

- Available CPU resource equivalent to a processor with a Processor Capacity Index (PCI) of at least 120 during periods of high CPU usages for certain scenarios. Please note that PCI is generally equivalent to millions of instructions per second (MIPS).
- Two gigabytes (2 GB) of central storage.

When calculating available CPU resource you need to consider the following. If your processor has multiple partitions, you should determine the percent of the machine available to the partition running z/OSMF. If you are using dedicated CPUs in your LPAR, compute the PCI as a percent based on the number of dedicated Central Processors (CPs) relative to total number of CPs. If you are using shared CPs, this calculation can be accomplished by adding the aggregate LPAR weights of all active partitions and computing the percent of this total represented by the LPAR weight of the target partition. Remember to subtract resources in use by other applications. If your configuration includes IBM System z Application Assist Processors (zAAPs) or IBM System z Integrated Information Processors (zIIPs), you should include the available zAAP and/or zIIP CPU time in your calculation. If you are considering purchasing a specialty processor, it is preferable to have a zIIP because new support in z/OS V1.9 allows an installation to run zAAP eligible workloads on zIIP processors.

The above z/OSMF system requirements should not be misconstrued as the steady state requirements for z/OSMF, which are very minimal (less than 0.5% of a z196 CPU). It is necessary to provide enough CPU resource to allow for periods of high z/OSMF CPU usage (spikes). Otherwise, z/OSMF may experience time outs and ABENDs which cause the WAS OEM servant address space to be aborted. Customers will find that most of the time, resources defined for the LPAR in which z/OSMF is running will be available to other programs or other LPARs.

CPU time accumulated on zAAP or zIIP processors is not included when computing the Monthly License Charge for MLC software pricing.

For help with determining the PCI rating of your processor or processor LPAR, see the *Large Systems Performance Reference* (LSPR), SC28-1187. This document is available on the IBM Resource Link® web site: https://www.ibm.com/servers/resourcelink

You can also choose to download the Processor Capacity Reference tool (zPCR) which is available to customers at no charge from the following Web site:
http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS1381

This tool allows you to import RMF™ reports and will display the PCI of your LPAR(s).

It may be possible to execute some z/OSMF applications and scenarios on systems that do not conform to the requirements above, depending upon the number of objects to be retrieved and/or displayed by z/OSMF. However, under these circumstances, z/OSMF may sometimes exhibit behaviors such as long response times, time outs, blank response screens, or abnormal terminations (ABENDs). Some customers may still choose to install z/OSMF because they plan to use the less compute intensive applications and scenarios. The information below should help you understand which applications and scenarios are most sensitive to the availability of CPU resource. Some applications may also run successfully in less memory, however, we recommend at least 2 GB. We do not recommend using Software Deployment or Configuration Assistant in systems or LPARs with significantly fewer resources than the recommended system described above.

You should try to avoid paging to auxiliary storage, as this will elongate response times and sometimes cause time outs.

Laboratory Performance Measurement Results

Performance tests in the IBM laboratory were executed using a variety of hardware configurations, including a system with resources equal to the minimum system requirements for z/OSMF. Data was collected for a single user by repeating a given scenario multiple times within a fixed test interval. Many scenarios can be repeated easily by clicking a refresh button, although others require multiple mouse clicks to re-drive the desired action.

Most tests were executed on a 2817-701 processor (z196) LPAR with a single dedicated CPU. Our LPAR had 8 GB of memory. However, to emulate a configuration meeting minimum system requirements, we also tested a capped LPAR having 1/10th of a z196 CPU and 2GB of memory.

The following software versions were installed:
1. z/OS V 1.13 and z/OSMF V 1.13
2. WebSphere Application Server OEM V7.0.0.17 64-bit
3. RMF Version V1R13

Browser considerations are described later in this paper. Supported browsers for z/OSMF can be found at the following Web address:
http://www-03.ibm.com/systems/z/os/zos/zosmf/browser_notes.html

The client browser can run in the following operating systems
- Microsoft® Windows® XP® 32-bit
- Windows 7® Professional 32-bit
- Windows Vista® Business Edition 32-bit

Workload Description

The workload was composed of a broad mixture of transactions using the following z/OSMF applications:
- z/OSMF Configuration Assistant (for the z/OS Communications Server)
- z/OSMF Capacity Provisioning
- z/OSMF Incident Log
- z/OSMF Resource Monitoring (a GUI for RMF)
- z/OSMF ISPF
- z/OSMF Application Linking and Launching
- z/OSMF Software Deployment (SMP/E information)
- z/OSMF WLM (Policy Editing and Display)
- z/OS jobs REST interface

This paper will focus on performance and configuration considerations for these z/OSMF tasks. For more information on what they do and their value, please see:
http://www-03.ibm.com/systems/z/os/zos/zosmf/

For example, these z/OSMF tasks can not only replace, older, repetitive, manual intensive processes, but also often accomplish these tasks faster and more effectively (with IBM best practices) than before.

The test workload did not include administration tasks, such as adding or deleting users because these scenarios are not often invoked on a regular basis. The objective was to find a variety of transactions to represent a typical steady state environment.

In all cases, a portion of the CPU time was eligible for execution on a System z Application Assist Processor (zAAP) or a System z Integrated Information Processor (zIIP). The eligible percentage varies with specific management tasks being executed and was observed to be between 4% and 95% of the total CPU time.

Applications that retrieve large amounts of data from files residing on the host, generally have less zAAP or zIIP eligibility. The potential varies not only with the individual z/OSMF application, but also with the environment in which it is executed. For example, the number of incidents in the Incident Log display may influence the percent of CPU time that can be eligible.

Note:	No other applications were executing during the performance test. RMF was running, because it is required to enable the z/OSMF Resource Monitoring application. Therefore, the CPU time for RMF is included in the totals. No other CPU monitoring was active. z/OSMF requires very minimal CPU resource when there are no active users. The z/OS operating system has timer driven processes, such as Missing Interrupt Handler, which cause a low level of activity even when no work is running in the system. z/OSMF itself uses less than 0.5% of the CPU (in the minimum configuration).

Resources for Specific z/OSMF Scenarios

The following information is provided to help you better understand the comparative CPU intensity of a representative subset of scenarios for z/OSMF applications.

Generally, scrolling actions in the browser view are executed on the client machine and use minimal CPU resources on the host.

z/OSMF uses Secure Socket Layer (SSL) communication between the browser client and the host. Many IBM System z mainframe processors can be configured with optional hardware based encryption that provides better physical security and improved performance.

The charts below show average CPU time per browser click for some representative z/OSMF scenarios based on tests made in the laboratory. Please note that your results may differ because of differences in the configuration of your system. For example, if your installation has less than 150 dumps or incidents, you may experience shorter response time and lower CPU resource consumption for Incident Log commands.

In Chart 1 below the following abbreviations are used

IL	Incident Log
Config Asst	Configuration Assistant
Cap Prov	Capacity Provisioning
SD	Software Deployment

The second chart contains some of the same data from the first chart, with Software Deployment data added. You can see that including Software Deployment causes the chart to be scaled differently than Chart 1. This illustrates how CPU requirements for Software Deployment compare with other applications. Note however, that many SD transactions use WebSphere asynchronous threads which can be classified separately in your WLM policy. This gives you the ability to adjust the priority of the asynchronous threads to avoid impacting other applications. You will find instructions for this in the z/OSMF Configuration Guide.

Our laboratory measurements were made on a system with the following characteristics:

150 total incidents
2 SMPE target zones, each with 631 data sets
Large WLM Policy is 1 MB, Small WLM Policy 35 K bytes

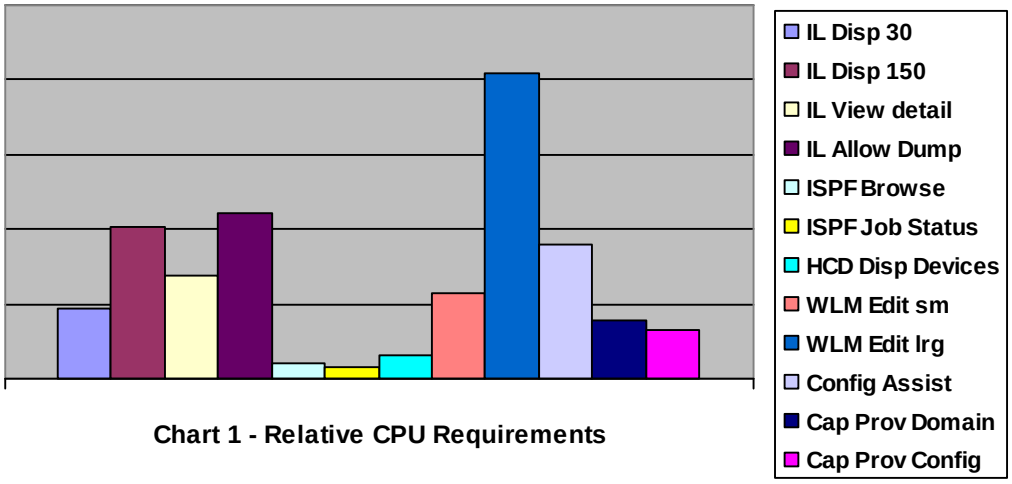

Chart 1 - Relative CPU Requirements

SDSF Status Display was 27 lines

ISPF Browsed Dataset was 86 lines

The scenarios shown for Software Deployment are priming the configuration, with one SMPE target zone and again with two target zones, as well as display catalogs for target data sets, display volumes for target data sets in a deployment and checking for missing SYSMODs

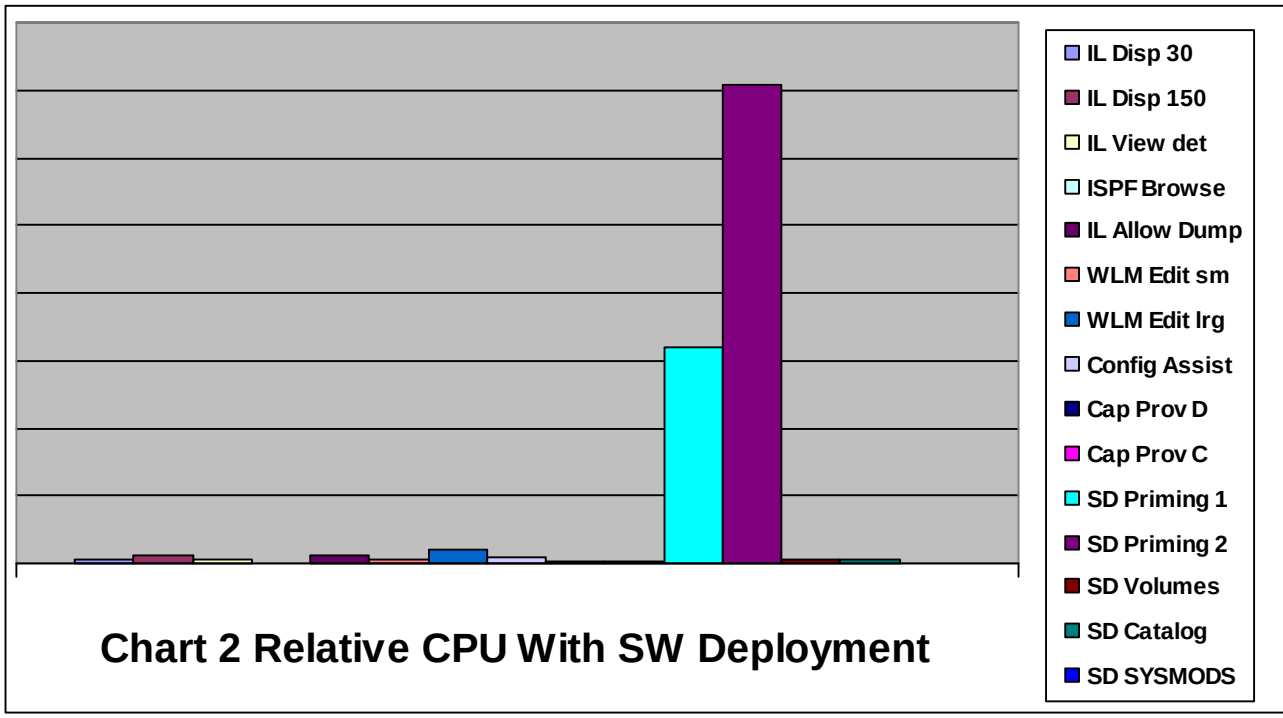

Chart 2 Relative CPU With SW Deployment

Chart 3 below illustrates the response time for selected z/OSMF scenarios when tested on a IBM zEnterprise™ 196 (z196) processor with one dedicated CPU and again with a z196 LPAR meeting recommended requirements. Internet Explorer Version 8 was used for these tests. WebSphere start up elongated from two minutes to six minutes and is not shown in the chart. See the section on Startup for more details.

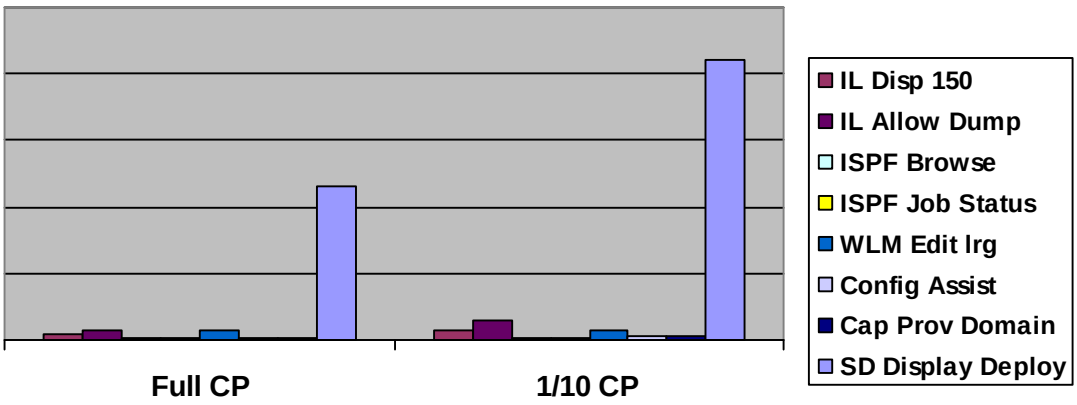

Chart 3 Response Time

Chart 4 below illustrates some of the same performance results as in Chart 1 and includes a visual indication of the percent of CPU time that is eligible for execution on a Specialty CP. The ZAAP/ZIIP system parameter enables the zAAP eligible work to execute in a zIIP processor. CPU time accumulated in the CIM server is directly eligible for a zIIP. However, since the percent of CIM CPU time is usually negligible, it has not been shown in the chart.

The percent of eligible Specialty CPU time may differ in your installation from those shown below due to variations in z/OSMF workloads under different circumstances.

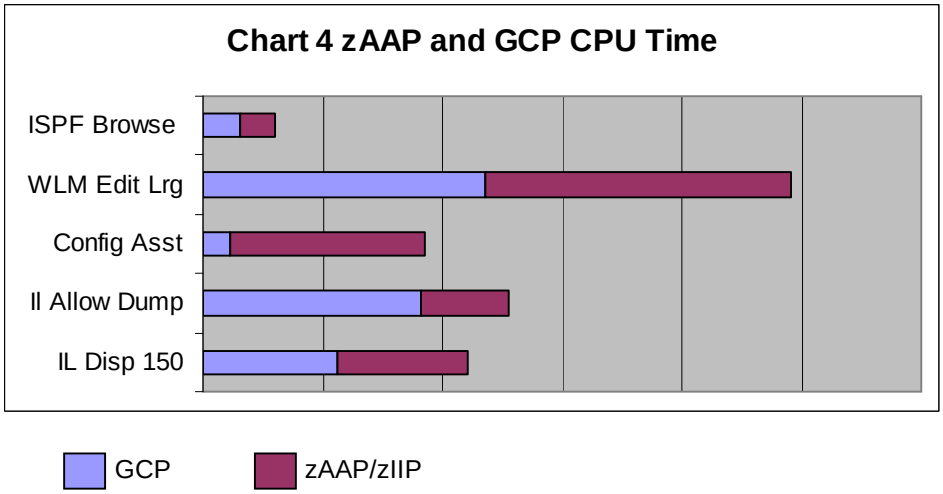

Browser Considerations

You may notice different performance characteristics depending upon the browser you select and the characteristics of the network and the client machine. Response time with Firefox is consistently better than with Internet Explorer. z/OSMF uses Javascript which takes longer to execute in Internet Explorer. The response time for z/OSMF also depends upon the speed of the user client computer and the efficiency of the network to which it is connected. Generally we observe response times of less than two seconds, except for scenarios that manipulate large amounts of data from the host.

Another characteristic of browsers is the manner in which they implement multiple browser instances and multiple browser tabs. Multiple browser instances appear to be independent, but they actually execute within the same process.

A process is also called an address space. It is an entity that runs on a computer and has its own code, memory and addressing range. Therefore, some z/OSMF applications can cause browser out of memory conditions when they are executed in more than one browser instance running simultaneously in the same Windows operating system image.

It is good practice to close browser tabs when you have finished working with them. This helps reduce the retained memory in the browser processes.

We recommend at least 2 GB of memory configured on your Windows Client systems. Our client systems in the laboratory have 3 GB to 4 GB of memory.

z/OSMF Startup Time

z/OSMF startup time varies considerably depending upon many factors. These include CPU resource, WebSphere configuration, number of servants, number of starting applications, type of applications, etc. At the present time, z/OSMF is configured with only one WebSphere servant address space, which results in faster startup time.

In the laboratory, we have observed z/OSMF startup time to be approximately one minute in a z196 LPAR consisting of one CPU however, the time elongated to 12 minutes with a minimally configured system. z/OSMF startup time refers to the time interval shown in SYSLOG between the timestamp on the START command and the BBOO0247I message, control region initialization complete. The control region does not fully initialize until all servants have completed initialization.

Here are some general recommendations for improving z/OSMF startup time. One of the most important things you can do to optimize z/OSMF initialization time and reduce the memory footprint is to configure z/OSMF with only those applications you plan to use. CPU constraint and paging activity during z/OSMF initialization are usually the most important factors contributing to long start time.

z/OSMF initialization requires the execution of various shell scripts, which run in BPXAS address spaces. Make sure the WLM policy used during z/OSMF initialization classifies BPXAS address spaces with appropriate priority. If they are allowed to default to a discretionary service class, they may run intermittently at low priority and may delay z/OSMF startup. Also make sure WAS OEM address spaces are properly classified in the WLM policy.

If your configuration includes System z Application Assist Processors (zAAPs) or System z Integrated Information Processors (zIIPs), you should be aware of the effects of specifying xxxHONORPRIORITY=NO in the IEAOPTxx member of PARMLIB. If you configure xxxHONORPRIORITY=NO, programs such as z/OSMF which are eligible to run on a Specialty CP are confined entirely to the Specialty CP. Thus if another program is utilizing all the available zAAP resources, z/OSMF will not be allowed to run on a general purpose processor. This can affect WebSphere startup time. In the above, xxx can be either IFA or IIP.

Some problems with the Domain Name Server (DNS) can also affect z/OSMF startup time. WAS OEM invokes the gethostbyname function, which returns an IP address when a hostname is specified. If your DNS is not working, or you are using a DNS on a network that requires multiple hops, z/OSMF initialization can be delayed. Also if there are entries for non-existent DNSs in the list of configured DNSs, TCP/IP will be delayed waiting for reply messages from these name servers.

Using the New z/OS Jobs REST Interface

z/OSMF now supports a z/OS jobs interface which uses Representative State Transfer (REST) and makes batch processing more accessible. REST services run as unauthorized programs and compete for z/OSMF resources and system resources with other users of z/OSMF Web browser services. Thus, concurrent high usage of the z/OS jobs REST interface can affect response time for users of the z/OSMF browser interface.

Applications receiving CIM connection errors while making heavy concurrent use of the cancel job, purge job or change job class requests types, may see benefit when using an increased value for the z/OSMF property IZU_CCP_USER_MAX. Increasing this option may cause unexpected consequences for other users and it should not be adjusted without first contacting IBM support.

Summary

z/OSMF can facilitate system operation, in some cases, by combining multiple manual tasks into a few clicks on a browser. In many cases, the manual steps consume more system resources than equivalent z/OSMF scenarios; but, because the manual steps are divided over a longer elapsed interval, their aggregate resource requirements are not obvious. For example, to complete some of the incident log scenarios might require multiple IPCS panels etc. In addition, there are some tasks that are very useful and either error prone or impossible to accomplish without z/OSMF. For example, using Software Deployment, you can identify missing requisite SYSMODs in your installed software by analyzing and comparing multiple SMPCSI data sets across multiple systems. This capability to analyze and compare SMPCSI data sets across systems is only available in z/OSMF and not in base SMP/E.

You should expect z/OSMF to consume modest amounts of CPU time as described above. The following are recommended best practices:

Windows clients should have at least 2 GB of memory

1. z/OSMF needs to run in a partition or system with a PCI of at least 120 and at least 2GB of memory which is required during periods of high CPU usage for some scenarios
2. Firefox browser performance is generally better then Internet Explorer
3. If using Internet Explorer, be sure to close unused browser tabs and browser windows
4. WebSphere OEM startup time can be tuned, if necessary – for more details see SG24-7816-00, System z Mean Time to Recovery Best Practices Redbook

z/OS Management Facility V1.13 Resource Requirements

IBM Systems and Technology Group
Route 100
Somers, New York 10589
U.S.A.
Produced in the United States of America,
09/2011

ZSW03197-USEN-00