Tivoli® OMEGAMON XE on z/OS

Version 4.2.0

User's Guide

SC27-2353-00

Tivoli® OMEGAMON XE on z/OS

IBM

Version 4.2.0

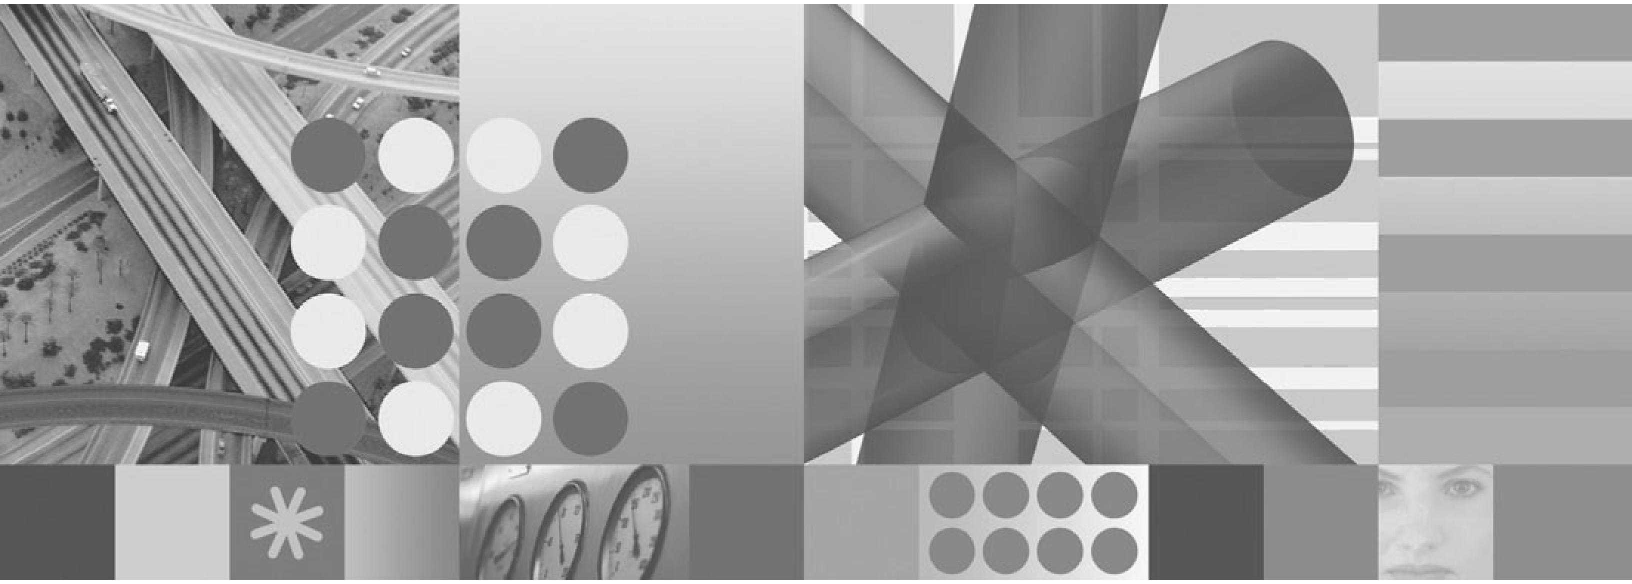

User's Guide

SC27-2353-00

Note

Before using this information and the product it supports, read the information in “Notices” on page 333.

First Edition (March 2009)

This edition applies to Version 4 Release 2, Modification 0 of IBM Tivoli OMEGAMON XE on z/OS (program number 5698-A33) and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

Figures

Tables

Part 1. Getting started

IBM® Tivoli® OMEGAMON® XE on z/OS® is a tool for monitoring, analyzing, and managing operating systems, workloads, and shared resources in a z/OS environment. The topics in this section describe the product and provide monitoring scenarios that illustrate how the product can be used to monitor and manage the performance and availability of z/OS systems.

Chapter 1, "Introducing IBM Tivoli OMEGAMON XE on z/OS," on page 3 contains an overview of the product and describes the new features in this release.

Chapter 2, "Using and customizing OMEGAMON XE on z/OS," on page 17 provides information about using the predefined workspaces and situations, configuring historical data collection for the product-provided monitored attributes, issuing UNIX® commands from the Tivoli Enterprise Portal, and configuring the launch of the Resource Measurement Facility Monitor III data portal.

Chapter 3, "Using the Inspect function," on page 35 contains instructions for using the Inspect function, which lets you determine where in an address space code is spending its time.

Chapter 4, "Using dynamic links to OMEGAMON for MVS," on page 43 discusses the predefined links and workspaces that let you access related OMEGAMON for MVS™ screens and describes how you can use them as the basis for creating your own context-sensitive links.

This guide is *not* intended to describe the features and functions of the Tivoli Enterprise Portal and the services it provides. Where appropriate, this guide refers you to the appropriate documents.

Chapter 1. Introducing IBM Tivoli OMEGAMON XE on z/OS

OMEGAMON XE on z/OS is a member of the latest generation of OMEGAMON XE on z/OS monitoring agents. Using this product, you can monitor and manage workload performance and resource utilization of entire sysplexes, as well as the individual z/OS systems that participate in them.

OMEGAMON XE on z/OS monitoring agents located on each z/OS image provide extensive system-level performance and usage information. In addition, they monitor the status and configuration of IBM cryptographic coprocessors installed in Series z servers and provide data on UNIX System Services hosted on z/OS systems.

OMEGAMON XE on z/OS also provides comprehensive usage information for sysplex-level resources such as coupling facilities, global enqueue, global resource serialization (GRS) ring systems, shared DASD groups, and cross-system coupling facilities (XCFs), as well as performance information for the service classes, report classes, and resource groups that use those resources. The product provides both system-level and sysplex-wide reporting of actual and potential usage of special processor resources.

New and migrated features allow users of 3270-based OMEGAMON products to migrate to the XE architecture without losing significant function. In addition, however, OMEGAMON XE on z/OS offers you continued access to the 3270 menu system (OMEGAMON for MVS) and CUA® (OMEGAMON II® for MVS) implementations.

Used in conjunction with other OMEGAMON XE monitoring products, the data, analyses, and alerts presented by OMEGAMON XE on z/OS help you develop a holistic view of your entire computing enterprise from a single console.

This chapter explains how OMEGAMON XE on z/OS works, describes the resources it provides, provides examples of how it can be used to monitor, analyze, and manage operating systems, workloads, and shared resources in a sysplex environment, and describes new features in this release.

How OMEGAMON XE on z/OS works

OMEGAMON XE on z/OS monitoring agents collect sysplex- and LPAR-level performance data on z/OS systems. The data can be viewed through a graphical user interface called Tivoli Enterprise Portal (TEP). A subset of data can also be viewed through a 3270 menu-driven interface.

The OMEGAMON XE on z/OS product takes advantage of the Tivoli Management Services infrastructure. Tivoli Management Services components provide security, data transfer and storage, notification mechanisms, user interface presentation, and communication services for a number of products, including IBM Tivoli Monitoring and OMEGAMON XE monitoring agents, in an agent-server-client architecture (Figure 1 on page 4). For more information on the Tivoli Management Services components, see the *OMEGAMON XE and Tivoli Management Services on z/OS: Common Planning and Configuration Guide* and *IBM Tivoli Monitoring: Installation and Setup Guide*.

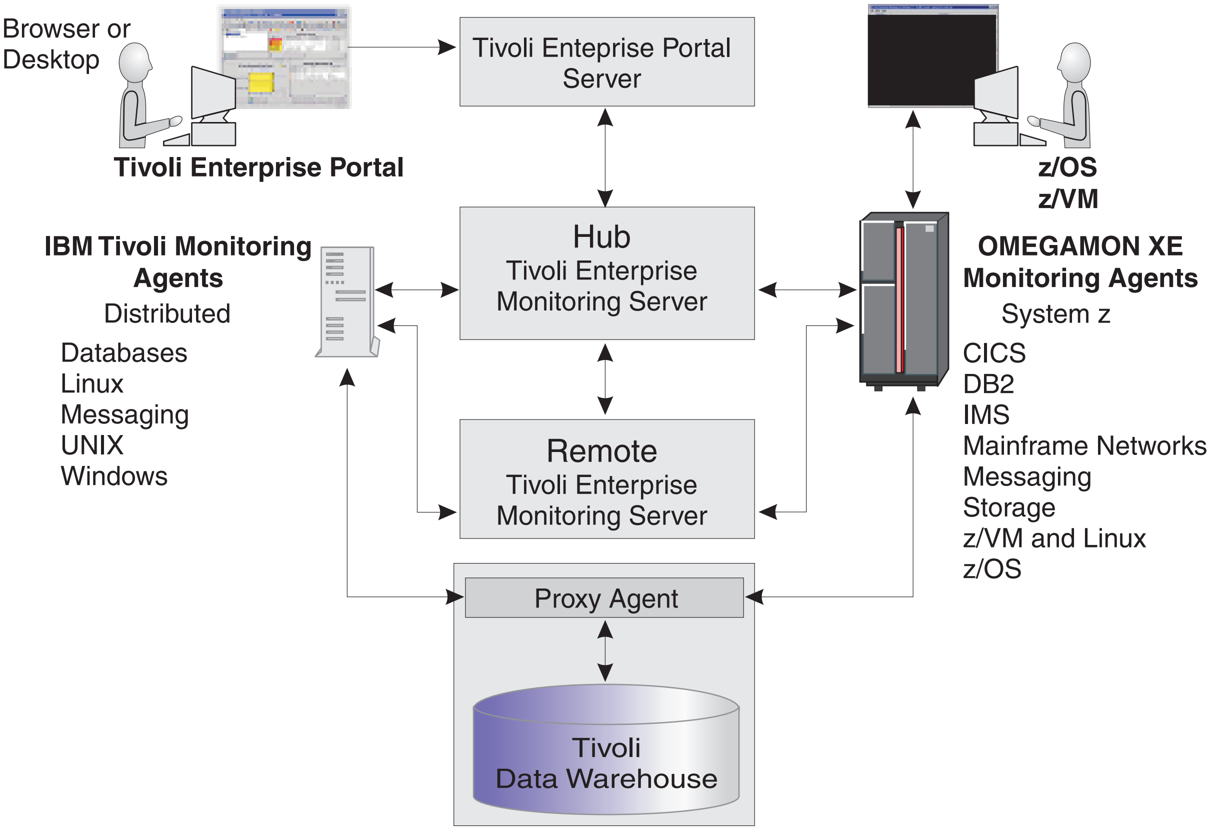

Figure 1. Agent-server-client architecture

In the Tivoli Enterprise Portal, real time and historical data collected by OMEGAMON XE on z/OS monitoring agents is displayed in tabular and graphic *views* in a set of predefined *workspaces* (see the example in Figure 2 on page 5).

When you open a workspace, Tivoli Enterprise Portal retrieves monitored data from the monitoring agents via the hub monitoring server and sends the results to the workspace. Chart and table views in the workspace use *queries* to specify what data the Tivoli Enterprise Portal requests.

The characteristics or properties of the logical and physical objects monitored by OMEGAMON XE on z/OS (for example, the amount of virtual storage allocated to a task) are known as *attributes*. These attributes are used to define the queries that specify the data to be displayed in the workspaces.

With the proper user authority, you can tailor these views to display critical or warning indicators when monitored values reach specified thresholds, and or to filter incoming data so you see only the information you are interested in at any given time. You can add additional views to existing workspaces or create your own workspaces and define your own queries using OMEGAMON XE on z/OS attributes.

Attributes are also used to describe *situations*, or conditions, that can trigger *events*. When situation events occur, event indicators are displayed in the Tivoli Enterprise Portal Navigator view. Situations can also trigger automated actions and policies.

You can link from a event indicator in the Navigator to a situation event workspace that provides information about conditions prevailing at the time the event occurred and about current conditions, as well as expert advice on how to handle the situation.

OMEGAMON XE on z/OS provides a set of predefined situations that you can run to monitor a wide range of conditions. You can also create your own situations using the OMEGAMON XE on z/OS attributes.

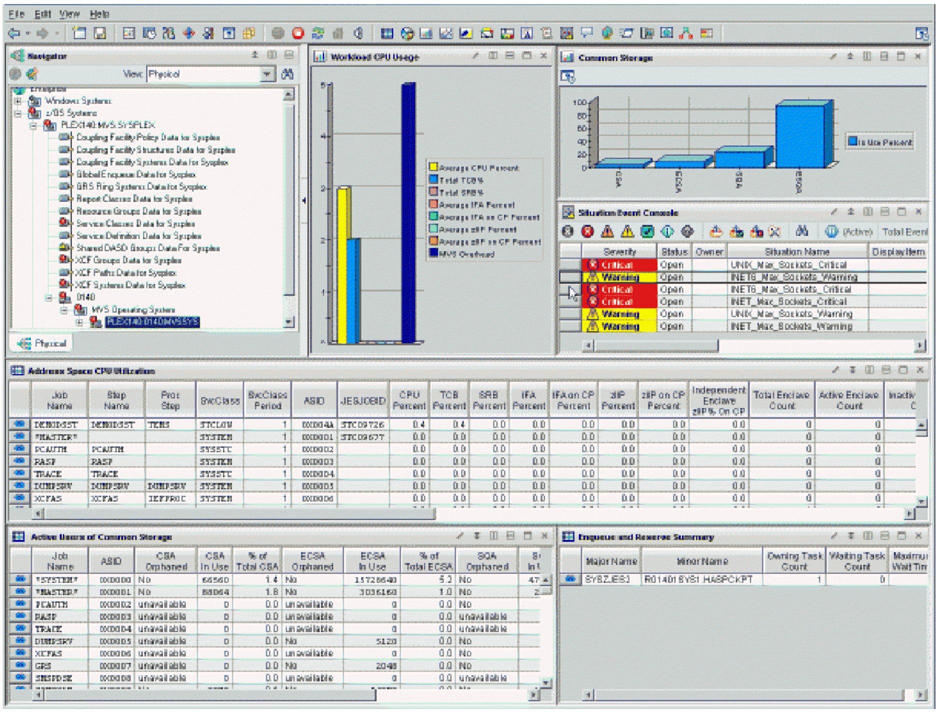
File
Edit
View
Help
Navigator
View: Physical
Enterprise
Windows Systems
z/OS Systems
PLEX140:MVS:SYSPLEX
Coupling Facility Policy Data for Sysplex
Coupling Facility Structures Data for Sysplex
Coupling Facility Systems Data for Sysplex
Global Enqueue Data for Sysplex
GRS Ring Systems Data for Sysplex
Report Classes Data for Sysplex
Resource Groups Data for Sysplex
Service Classes Data for Sysplex
Service Definition Data for Sysplex
Shared DASD Groups Data For Sysplex
XCF Groups Data for Sysplex
XCF Paths Data for Sysplex
XCF Systems Data for Sysplex
0140
MVS Operating System
PLEX140:0140:MVSSYS
Physical
Workload CPU Usage
5
4
3
2
1
0
Average CPU Percent
Total TCB%
Total SRB%
Average IFA Percent
Average IFA on CP Percent
Average zIIP Percent
Average zIIP on CP Percent
MVS Overhead
Common Storage
100
80
60
40
20
0
CSA
ECSA
SQA
ESQA
In Use Percent
Situation Event Console
(Active)
Total Event

	Severity	Status	Owner	Situation Name	Display Item
	Critical	Open		UNIX_Max_Sockets_Critical	
	Warning	Open		INET6_Max_Sockets_Warning	
	Critical	Open		INET6_Max_Sockets_Critical	
	Critical	Open		INET_Max_Sockets_Critical	
	Warning	Open		UNIX_Max_Sockets_Warning	
	Warning	Open		INET_Max_Sockets_Warning	

Address Space CPU Utilization

Job Name	Step Name	Proc Step	SrvClass	SrvClass Period	ASID	JESJOBID	CPU Percent	TCB Percent	SRB Percent	IFA Percent	IFA on CP Percent	zIIP Percent	zIIP on CP Percent	Independent Enclave zIIP % On CP	Total Enclave Count	Active Enclave Count
DEMODSST	DEMODSST	TEMS	STCLOW	1	00004A	STC09726	0.4	0.4	0.0	0.0	0.0	0.0	0.0	0.0	0	0
MASTER			SYSTEM	1	000001	STC09677	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0
PCAUTH	PCAUTH		SYSSTC	1	000002		0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0
RASP	RASP		SYSTEM	1	000003		0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0
TRACE	TRACE		SYSSTC	1	000004		0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0
DUMPSRV	DUMPSRV	DUMPSRV	SYSTEM	1	000005		0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0
XCFAS	XCFAS	IEFPROC	SYSTEM	1	000006		0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0

Active Users of Common Storage

Job Name	ASID	CSA Orphaned	CSA In Use	% of Total CSA	ECSA Orphaned	ECSA In Use	% of Total ECSA	SQA Orphaned	S In
SYSTEM	000000	No	66560	1.4	No	15728640	5.2	No	47
MASTER	000001	No	88064	1.8	No	3036160	1.0	No	2
PCAUTH	000002	unavailable	0	0.0	unavailable	0	0.0	No	
RASP	000003	unavailable	0	0.0	unavailable	0	0.0	unavailable	
TRACE	000004	unavailable	0	0.0	unavailable	0	0.0	unavailable	
DUMPSRV	000005	unavailable	0	0.0	No	5120	0.0	No	
XCFAS	000006	unavailable	0	0.0	unavailable	0	0.0	No	
GRS	000007	unavailable	0	0.0	No	2048	0.0	No	
SMSPDSE	000008	unavailable	0	0.0	unavailable	0	0.0	unavailable	

Enqueue and Reserve Summary

Major Name	Minor Name	Owning Task Count	Waiting Task Count	Maximu Wait Tim
SYSZJES2	R0140I SYS1.HASPCKPT	1	0	

Figure 2. System Overview workspace

Resources provided by OMEGAMON XE on z/OS

The following sections provide an overview of the resources provided by OMEGAMON XE on z/OS.

OMEGAMON XE on z/OS attributes

OMEGAMON XE on z/OS monitors over 50 groups of attributes, providing a wealth of sysplex- and system-level data. You can use these attributes to tailor the information presented in workspaces, or to define situations that target specific threshholds, events, or performance problems you want to monitor.

For descriptions of the attribute groups and the individual attributes monitored by OMEGAMON XE on z/OS, see Chapter 10, “Attributes,” on page 95.

Predefined workspaces

OMEGAMON XE on z/OS provides two sets of predefined workspaces: sysplex-level workspaces and system-level workspaces. Each workspace or group of workspaces displays a specific set of data.

You can use the information provided by these workspaces to manage the performance and availability of systems and their resources, to identify potential problems, to trace the causes of alerts or exceptions, to make tuning and resource distribution decisions, and to identify particular conditions you want to monitor.

For information on locating and navigating OMEGAMON XE on z/OS workspaces, see “Using the predefined workspaces” on page 17. For a complete list and descriptions of the predefined workspaces, see Chapter 11, “Workspaces,” on page 219. For information about creating or customizing workspaces, see *IBM Tivoli Monitoring: Administrator’s Guide* and *IBM Tivoli Monitoring: User’s Guide* or the Tivoli Enterprise Portal online help.

Predefined monitoring situations

Situations are descriptions of conditions you want to monitor for, such as rapid growth in usage of the common service area (CSA). Situations periodically verify the values of attributes used in the situation description. When they are distributed to systems monitored by OMEGAMON XE on z/OS agents, situations can, for example, alert you to a coupling facility structure that has failed, or to a service class that is failing to meet its goal. Situations can also trigger simple (reflex) actions, or complex automation policies.

If situations are associated with Navigator items, they can generate auditory or visual event indicators, which provide access to special event workspaces containing more information about the event and guidance for how it should be handled.

OMEGAMON XE on z/OS provides an extensive set of predefined situations. These situations check for conditions that are typically considered to be problematic or noteworthy. They can also serve as templates for creating customized situations of your own. All these situations include expert advice for handling these conditions should they arise.

For information on activating and customizing the predefined situations, see "Using the predefined workspaces" on page 17. For descriptions of the predefined situations, see Chapter 12, "Situations," on page 279. For more information on creating, editing, and distributing situations, see the Tivoli Enterprise Portal online help and documentation. For information on migrating an OMEGAMON II profile, see *Tivoli OMEGAMON XE on z/OS Planning and Configuration Guide*.

Historical data collection and reporting

In addition to providing real time data, OMEGAMON XE on z/OS also lets you collect data over extended periods of time. By studying the information gathered from a historical perspective, you can, for example, identify trends and determine whether current performance is typical or exceptional, or evaluate the effect of tuning decisions.

You can view the historical data collected by OMEGAMON XE on z/OS in Tivoli Enterprise Portal workspaces or in reports generated by third-party reporting tools.

For more information about historical data collection for OMEGAMON XE on z/OS, see "Using historical data collection and reporting" on page 26, *IBM Tivoli OMEGAMON XE on z/OS: Planning and Configuration Guide*, *OMEGAMON XE and Tivoli Management Services on z/OS: Common Planning and Configuration Guide*, and *IBM Tivoli Monitoring: Administrator's Guide*.

Note: The OMEGAMON II for MVS component provides separate historical data collection and reporting through the EPILOG® collector. (The OMEGAMON II for MVS component of OMEGAMON XE on z/OS version 550 and later supports only system-level EPILOG data.) For information about historical data collection and reporting for OMEGAMON II for MVS, see *IBM Tivoli OMEGAMON XE on z/OS: EPILOG User's Guide*.

What you can do with OMEGAMON XE on z/OS

The data and resources provided by OMEGAMON XE on z/OS enable you to monitor and manage workload performance and resource usage in a variety of ways. The following sections describe some of them.

Monitor processor usage at the system and address space level

System CPU utilization is the percentage of time that all processors available to a system were busy dispatching work. The System CPU Utilization workspace provides information about processor usage for each monitored z/OS system. In an LPAR environment, it also provides partition management statistics.

The System CPU Utilization workspace shows the number of physical processors reported on, number of processors online, the average percentage of time that all processors collectively available in this z/OS system were busy dispatching work, and other information specific to CPU workload and partition workload.

If you use defined capacity as a basis for pricing, this workspace also shows the long-term average processor service used by this system or LPAR in millions of service units (MSUs) per hour.

For each address space in a z/OS image or LPAR, the Address Space CPU Utilization workspace provides basic identifying information such as job name and ASID, basic Workload Manager information such as service class and service class period, as well as various CPU statistics. It also provides enclave data, such as the total number of dependent and independent enclaves owned by the address space that are currently active or inactive.

Manage resources and workloads across LPAR clusters

OMEGAMON XE on z/OS provides an overview of the LPAR clusters in your central processor complex/cluster processor complexes (CPCs) and allows you to look at the system level details for the members of individual clusters.

Identify service classes that are missing their goals

OMEGAMON XE on z/OS allows you to determine whether a workload is meeting its goal, determine which resources a service class is using, and which resources are affecting service class performance.

Monitor enqueues across systems and sysplexes

OMEGAMON XE on z/OS provides information on enqueues shared across systems in a sysplex (global enqueues) in several related workspace. In environments in which enqueue management spans two or more sysplexes using Unicenter® CA-MIM™ Resource Sharing, it can also provide information on resources in conflict between users in multiple sysplexes, using the concept of an enqplex. (For information on assigning sysplexes to an enqplex, see *Tivoli OMEGAMON XE on z/OS Planning and Configuration Guide*.)

Identify and ease bottlenecks

Bottleneck analysis is a performance monitoring technique that identifies execution states of a workload and the frequency of each state. When the results of this analysis are averaged over time, it is possible to find what states (such as waiting for CPU) prevent a workload from achieving its service goal. Identifying and easing bottlenecks are a key part of performance management.

OMEGAMON XE on z/OS provides summary bottleneck data for all monitored address spaces or for a selected service class period, displaying over 20 execution states. It also provides information for over 50 execution states for selected address spaces.

Impact analysis helps you determine how various workloads are interfering with each other, by showing which workloads are using the resources that an impacted workload needs. This helps you to reduce degradation of the monitored workloads.

The Address Space Bottlenecks and Impact Analysis workspace displays resource contention statistics for a selected address space. Bar charts let you quickly determine which address spaces are impacting the favored address space, how they are impacting it, and to what extent they are impacting it.

Monitor common storage areas usage by address space

Common storage comprises:

- Common Service Area (CSA)
- Extended Common Service Area (ECSA)

- System Queue Area (SQA)
- Extended System Queue Area (ESQA)

Common storage area reporting at the address space level includes common storage area elements in use by active address spaces and orphaned elements, and usage history (trend details) by address space.

Determine processor capacity requirements

OMEGAMON XE on z/OS provides cumulative central processor usage in seconds and percentages for address spaces (job-level CPU). The Address Space CPU Utilization attribute group provides CPU, task control block (TCB), service request block (SRB), and preemptive SRB CPU times and percentages for the job as a whole, as well as job start date and time and elapsed time. These values are displayed by default in the Address Space CPU Utilization workspace.

The On/Off Capacity on Demand feature allows z/OS users to temporarily add processors to their configuration. The job-level CPU times and percentages provided by OMEGAMON XE on z/OS help you understand what workloads are using the processors and ensure the resource is getting to the correct workloads in a timely manner.

Plan and monitor special processor resources

IBM z/Series Application Assist Processors (zAAPs) are a special class of assist processors designed to run Java™ workloads. IBM System z9® Integrated Information Processors (zIIPs) handle several types of workload for DB2® version 8 and will eventually handle other types. OMEGAMON XE on z/OS provides zAAP and zIIP data in various workspaces that can help you determine how well the assist processors configured to your environment, and the workloads running on them, are performing.

OMEGAMON XE on z/OS can also help you decide how much zAAP or zIIP resource you need, based on the performance of workloads on your regular processors.

Identify inefficient or looping code

The Inspect function allows you to observe where in the executable code a z/OS address space is spending its time. Inspect provides processor usage data for a selected address space drilled down to the agent-selected level of granularity within each control sector (CSECT), for the most active task control blocks (TCBs). This information helps identify inefficient code, or where in an address space code may be looping. The Inspect CPU Usage workspace contains sampling statistics and messages sent by the Inspect agent, which enables you to evaluate the statistical accuracy of the Inspect data.

Identify looping tasks

Address spaces running within the z/OS operating system can occasionally fall into a loop executing the same set of instructions repeatedly. This looping can absorb large amounts of precious CPU resources and prevent other work from processing efficiently. The z/OS operating system, and its Workload Manager (WLM) component in particular, tries to distribute CPU resources to address spaces in an intelligent way by honoring goals established by the system administrator. The algorithms used to distribute resources will interrupt a looping task and give the CPU to other work.

While this does reduce the impact of a looping job, it also masks the job, making it hard to identify the looping task. Since the looping task continues to be given the CPU on a periodic basis, the task will continue to squander resources that could have been given to “well behaved” workloads. It is important that system administrators discover looping jobs as soon as possible. Current monitoring technology can be set to raise alerts when a task uses an excessive amount of CPU, but because of the Workload Manager's actions, looping jobs may not clearly stand out. Setting an arbitrary threshold such as 30%, 40%, or 50% CPU utilization in an attempt to detect looping tasks will often miss looping tasks or worse produce high numbers of false positives.

The Looping CPU Index, which is one of the Address Space Bottlenecks attributes, helps you identify looping tasks.

Interoperability and integration with other products

OMEGAMON XE zSeries® products are designed to integrate with each other and with other products that use Tivoli Management Services. These products exploit the power of the Tivoli Enterprise Portal to integrate and correlate performance and availability information from a variety of sources. For example, the Tivoli Enterprise Portal allows you to create custom workspaces composed of data from a range of Tivoli service availability and performance management solutions (IBM Tivoli Monitoring, IBM Tivoli Composite Application Management, and IBM Tivoli NetView® for z/OS, as well as OMEGAMON XE). You can create context-sensitive links between product workspaces to obtain additional information about systems, subsystems, resources or network components that are being monitored by other monitoring agents, or to TN3270-based applications.

In addition, OMEGAMON XE products are being integrated with an increasing number of other Tivoli and IBM products. Situation events reported by OMEGAMON XE monitoring agents can be forwarded to Tivoli Event Console or Tivoli Netcool®/OMNIbus™ for event correlation and management. From the Tivoli Enterprise Portal you can launch in context into other Web-based or Web-enabled Tivoli applications like Tivoli Business Services Management without having to re-enter user credentials, and you can launch in context into the Tivoli Enterprise Portal from other Tivoli applications.

For more information on cross-product linking, see “Launching the RMF Monitor III data portal” on page 32, “Using cross-product workspace links” on page 21. For information on linking in context to OMEGAMON for MVS screens, see Chapter 4, “Using dynamic links to OMEGAMON for MVS,” on page 43.

New in this release

OMEGAMON XE on z/OS provides full support for IBM Tivoli Monitoring V6.2.1 or later and compatibility and exploitation of z/OS V1.9 and 1.10. In addition, V4.2.0 incorporates new features, enhances existing functionality, and continues migration of functionality from OMEGAMON classic and OMEGAMON II for MVS.

- **IBM Tivoli Monitoring V6.2.1 support**

  OMEGAMON XE on z/OS V4.2.0 is designed to run on Tivoli Monitoring V6.2.1 or later and provides support for all Tivoli Monitoring V6.2.1 features. Read about these features in *IBM Tivoli Monitoring: User’s Guide*.

  OMEGAMON XE on z/OS supports a mixture of V3.1.0, V4.1.0 and V4.2.0 monitoring agents in your environment during the migration period so that you can deploy new V4.2.0 monitoring agents to z/OS systems and subsystems along with older monitoring agents of the same product. For more information about running in this mixed environment, refer to the *IBM Tivoli OMEGAMON XE and Tivoli Management Services on z/OS: Upgrade Guide*.

- **Merger of OMEGAMON XE and OMEGAMON II for MVS FMIDs**

  In V4.2.0, the function modification identifier (FMID) for OMEGAMON XE on z/OS and OMEGAMON II for MVS have been merged to reflect the fact that OMEGAMON II for MVS is a component and not an independent product. The FMIDs for OMEGAMON II (CUA and "Classic") 3270 components have been merged into the same FMID as the OMEGAMON XE base product to simplify product installation and maintenance. All XE agents and 3270 interface modules are now consistently delivered in a single FMID per product, across the entire OMEGAMON 420 family of products. OMEGAMON II components have been reversioned to V4.2.0 for consistency with the base product.

- **Documentation changes**

  Documentation for the 3270-based menu systems interface (sometimes called Classic) for the z/OS has been rewritten, brought up to date, and integrated with the OMEGAMON XE documentation. The OMEGAMON II books previously published for OMEGAMON II for MVS product have been sunset and should be considered out of date. Note that no enhancements have been made to these books since version 5.2 of their OMEGAMON II interfaces (released with version 3.1.0), but they are still available in the previous releases section of the OMEGAMON XE and IBM Tivoli Monitoring Information Center. The documentation for the End-to-End feature is still fully supported. The End-to-End reference manual is now considered a part of the shared documentation set for several OMEGAMON XE products.

- **Reversioning of common components**

  The following common components and SMP/E FMIDs have been reversioned to reflect the current version of Tivoli Management Services:

  - OMNIMON Base (HKOB620)
  - End to End Response Time (HKET620)
  - Shared probes (HKSB620)

- **Version numbering of components and queries**

  In the Tivoli Enterprise Portal, all components are identified by the version number in the Managed System Status table view. For example, the Version column for rows in this table view might display the following entries:

  - 04.20.00 with an entry for a monitoring agent indicates that the version of this monitoring agent is 4.2.0.
  - 06.21.00 with an entry for a monitoring server indicates that the version of this monitoring server is IBM Tivoli Monitoring 6.2.1.

  The new version numbering feature helps you to quickly identify the version of each monitoring agent in an environment that comprises multiple versions.

  In addition, the version number appears in the description of queries in the Situation editor. Multiple queries with the same name and description are required to support upgrade scenarios. The addition of a version number enables you to identify which query is appropriate for a particular agent at a particular version.

- **zIIP offload enablement**

  A portion of the OMEGAMON XE on z/OS DASD data collection processing is redirected to IBM System z® Integrated Information Processors (zIIPs) where available. This frees up the standard processors for other work and can reduce software licensing costs.

  Redirection of processing occurs by default, but you can disable the offloading by adding a KM5ZIIPOFFLOAD=NO to the &*rhilev*.&*rte*.RKANPARU(KDSENV) file using the new nonstandard parameters editing facility in the Configuration Tool.

- **Monitoring of system lock information**

  Suspend and spin locks are part of the serialization functionality in z/OS. Occasionally, tasks can be significantly delayed if a lock holder fails to release the resource in a timely manner. For z/OS images running at 1.10 and higher, OMEGAMON XE on z/OS now provides monitoring services for system locks. Resource Management Facility (RMF™) data collection must be enabled for lock data to be available.

  The suspend and spin lock data has been added to the Enqueue and Reserve Summary workspace. The name of this workspace has been changed to Enqueue, Reserve, and Lock Summary to reflect this addition.

- **Monitoring of blocked workloads**

  This enhancement provides the capability to monitor Workload Manager (WLM) service and report classes that ran work units at a promoted dispatching priority.

  It is no longer unusual for users to run zSeries systems at full capacity for extended periods. As a result, workloads may wait for extended time periods. In many cases this is not a problem: work of lower importance work must wait until resources become available. However, in some cases this may cause a priority inversion, as when, for example, low priority work obtains a resource that higher importance work is waiting on, but is blocked by a large CPU consumer of medium importance. The high priority work is now in effect blocked by the medium priority work.

Starting with z/OS V1.9, WLM addresses this issue by granting limited, or *trickle*, CPU access to work units that cannot get hold of a CPU for an extended period of time. WLM periodically examines the IN queue and identifies work units that have been CPU starved. Then, the dispatching priority of these work units is temporarily raised (*promoted*) to allow execution of a small number of instructions (so-called *trickle support*). The assumption is that such short periods of CPU access do not harm high importance work and could help low importance work to release locks and other critical resources.

OMEGAMON XE on z/OS V4.2.0 allows users to identify and monitor workloads that may be exploiting the blocked workload enhancement. The WLM Service Class Resources workspace now identifies workloads that ran work units at a promoted dispatching priority and indicates the percentage of time that the workload ran at a promoted priority.

- **New options for RMF data collection**

  Resource Management Facility (RMF) data is now available for system lock and cross-coupling facility (XCF) data, as well as for coupling facility data. A configuration variable allows you to configure collection for all RMF-supplied data, for lock data only, or for coupling facility and cross-coupling facility data only. The default is no RMF-supplied data.

- **Monitoring of UNIX System Services zFS**

  The z/OS Distributed File Service (DFS™™) zSeries File System (zFS) is a z/OS UNIX file system that can be used in addition to the Hierarchical File System (HFS). zFS provides significant performance gains in accessing files approaching 8K in size that are frequently accessed and updated. The access performance of smaller files is equivalent to that of HFS.

  OMEGAMON XE on z/OS V4.2.0 introduces monitoring of statistics for zFS kernel data, storage, metadata cache, directory cache, and user cache.

- **Monitoring of UNIX System Services socket usage**

  OMEGAMON XE on z/OS V4.2.0 also introduces monitoring of UNIX System Services socket usage:

  - The following metrics added to USS Kernel table in the UNIX Kernel workspace for each type of socket (Internet, Internet V6, and UNIX):
    - Maximum number of sockets configured
    - Number of sockets currently in use
    - High water mark for socket usage
    - Percentage of sockets in use
    - High water mark as a percentage of maximum
  - A new attribute added to the OMEGAMON USS Kernel table that identifies the zFS address space name and the OMVS address space name.
  - Two new links from the OMEGAMON USS Kernel workspace views provide filtered navigation to existing OMEGAMON XE on z/OS Address Space CPU Usage Details workspace. One link employs the OMVS address space name as a filter, the second employs the zFS address space name as a filter.
  - A new link from the USS Overview workspace USS Kernel view to a new zFS system-level activity summary workspace.

- **Reporting of new model capacities and associated capacity ratings**

  Capacity provisioning management (CPM), introduced on the System z10™ platform, provides the ability to automatically "provision" physical processors to and from a Central Processor Complex (CPC). As processors are provisioned, the capacity of the CPC changes and two new capacity indicators, Model Temporary Capacity and Model Permanent Capacity, come in to play.

  Currently, a single Model Capacity is represented by the model number associated with the full capacity rating of an individual machine (for example, 2097-718), as well as the maximum millions of service units (MSU) capacity that can be delivered. With the introduction of On/Off Capacity On Demand (OOCoD) and Capacity Backup (CBU), the effective Model Number associated with different numbers of provisioned CPUs changes, based on the new aggregate MSU capacity that can be delivered. These changes affect software licensing charges based on the capacity increases and decreases.

OMEGAMON XE on z/OS V4.2.0 reports on all three model capacity indicators and their associated MSU capacity ratings instead of Model Capacity, reported in previous versions. Reporting of the new model capacities and associated capacity ratings provides information that can be used to determine what CPU resource is available within a CPC for delivery to workloads at any point in time, or within an historical timeframe. The identifiers reflect manual actions performed to increase/decrease CPC capacity or the effects of automatic provisioning as determined by a CPM Policy.

- **Monitoring of HiperDispatch**

  To address increasing workload demand for processor cycles and high speed memory access, z/OS on the IBM System z10 implements a new approach to dispatching work. HiperDispatch throughput improvements have been achieved by making z/OS aware of the underlying physical topology of configured processors. z/OS can use this awareness to attempt to re-dispatch a unit of work repeatedly on the same physical CPU, or collection of physically adjacent CPUs (an affinity node), to increase the chances of obtaining data from processor L1, L1.5 or L2 cache, instead of incurring a time delay by going to main memory.

  Support for HiperDispatch was introduced as an interim feature in OMEGAMON XE on z/OS V4.1.0, to coincide with the launch of System z10. A new HiperDispatch Management attribute has added to the System CPU Utilization attribute group and is displayed in the System CPU Utilization workspace. A link from the workspace connects to a new HiperDispatch Details workspace, which displays data from two new attribute groups: HiperDispatch Logical Processors attributes, which show processor type and HiperDispatch statistics like priority, share percentage, and status for each processor configured in an LPAR, and HiperDispatch Management attributes, which provide information about the weight and status of a logical processor when the LPAR is in HiperDispatch mode.

  In addition, the OMEGAMON for MVS **HDSP** command now displays HiperDispatch management metrics.

- **Monitoring of Extended Address Volumes**

  Extended Address Volumes (EAV) is a new concept, introduced in z/OS V1.10, designed to deal with the continued growth in z/OS disk storage requirements and the constraints that users currently face in that area. EAV increases the amount of addressable DASD storage per volume beyond the 65,520 cylinder limit to an architectural maximum of 268,434,453 cylinders per volume by changing how tracks on Extended Count Key Data (ECKD™) volumes are addressed. This allows the z/OS operating system, which is limited to 65,535 devices, to satisfy future disk storage requirements by increasing the individual capacity, rather than the total number, of DASD volumes.

  Support for EAV has been also been added to OMEGAMON for MVS. The SEEK, SVOL, DSN, and DSNV OMEGAMON commands have been updated.

- **Support for reusable address spaces**

  z/OS V1.9 introduced the ability to reuse address space identifiers (ASIDs). ASID reuse provides relief for z/OS users who currently have to schedule IPLs to reclaim ASIDs that have become non-reusable (that is, address spaces that produce message “IEF352I ADDRESS SPACE UNAVAILABLE” on termination), by allowing them to create reusable ASIDs for product started tasks.

  OMEGAMON XE on z/OS V4.2.0 introduces support for this function by allowing OMEGAMON XE on z/OS started tasks to be started with the new z/OS REUSASID=YES start command parameter.

- **Classic commands introduced in V4.1.0 Interim Feature**

  PTF UA37434 added support to the OMEGAMON for MVS interface for the following data:

  - Workload Manager (WLM) configuration data and performance metrics
  - Detailed enclave CPU consumption by owning address spaces
  - License Manager 4-Hour Rolling Average MSU details
  - Total indirect enclave CPU consumption for an address space

  For detailed information on new commands, see *IBM Tivoli OMEGAMON XE on z/OS: OMEGAMON for MVS User's Guide* and the online help.

- **New z/OS System Overview workspace**

A new z/OS System Overview workspace summarizes key performance aspects of the LPAR. This workspace is the default workspace for each managed system item in the Navigator.

- **New CPU Loop Index attribute**

  This attribute, added to the Address Space Bottlenecks attribute group, helps you identify looping jobs.

  The CPU Loop Index is a percentage value representing the sum of all CPU, zIIP, zIIP on CP, zAAP, and zAAP on CP using and waiting counts, divided by total sample count. For CPU looping jobs, this value is usually above 98%.

  The CPU Loop Index appears in the Address Space Bottlenecks Summary workspace, the Address Space Bottlenecks Detail workspace, and the Address Space Bottlenecks and Impact Analysis workspace. A new situation, KM5_CPU_Loop_Warn alerts you to potentially looping address spaces. (Note that very CPU intensive jobs may read high without being in a loop, so this value is a guide, not a guarantee.)

- **Support for situation event forwarding**

  If you are using IBM Tivoli Enterprise Console® (TEC) or IBM Tivoli Netcool/OMNIbus, in addition to IBM Tivoli Monitoring, to manage events in your enterprise, you can now forward events reported by OMEGAMON XE on z/OS monitoring agents to these event management products. The benefits of these products and the details of how they can be integrated with IBM Tivoli Monitoring are described in the ″Event integration scenarios″ section of the *IBM Tivoli Monitoring: Installation and Setup Guide*.

  Before situation events can be forwarded, event forwarding must be enabled on the hub monitoring server, and a default destination server must be defined. In addition, the TEC or Netcool/OMNIbus server (the event server) must be configured to receive the events, a situation update forwarding process must be installed on the event server, and, for events forwarded to TEC a baroc file for the agent must be installed and imported on the event server. The *IBM Tivoli Monitoring: Installation and Setup Guide* provides detailed instructions for enabling event forwarding from a distributed Tivoli Enterprise Monitoring Server and for configuring TEC and OMNIbus to receive the events, including the installation of the event synchronization component and installing the .baroc files. *IBM Tivoli Management Services on z/OS: Configuring the Tivoli Enterprise Monitoring Server on z/OS* provides instructions for configuring a hub monitoring server on z/OS and locating the agent baroc files.

  After event forwarding is enabled, by default all situation events are forwarded to the specified event server. However, you can customize which situation events are forwarded and to which event server, using the Situation editor in the Tivoli Enterprise Portal. You may also need to assign an event status compatible with the target event server. For information on specifying which situation events to forward, see the Tivoli Enterprise Portal online help and the *IBM Tivoli Monitoring: User's Guide*.

- **New address space performance data**

  OMEGAMON XE on z/OS now provides the following address space resource and storage information, previously available only in OMEGAMON for MVS:

  - Address Space status (wait or swap reason and wait or swap time)
  - Working set size
  - Page-ins per second
  - Page-outs per second
  - Shared page views
  - Shared pages in central storage
  - Other auxiliary storage slots
  - Dispatch priority
  - UIC statistics

  The corresponding attributes have been added to the Address Space Real Storage attribute group.

  The OMEGAMON XE on z/OS Address Space Overview workspace has been restructured to improve consolidate and update workspace content, and a new Address Space Details for Job workspace has been introduced.

- **Redesigned OMEGAMON XE on z/OS Address Space Overview workspace**

The Selected Execution States view that shows bottleneck data has been eliminated to improve the overall workspace performance. Bottleneck data can still be accessed through a link from the Address Space Counts view.

The CPU Usage view has been expanded to include enclave CPU usage.

The data in the Address Space CPU Utilization Summary view is now returned by the agent in descending CPU percentage order by the agent. You can now see quickly which address spaces are using the most CPU, even when the result set is split across multiple pages within the view.

Several of the views in this workspace have been re-arranged.

- **Dynamic linking from OMEGAMON XE to OMEGAMON for MVS (classic)**

  Dynamic linking from the Tivoli Enterprise Portal to the OMEGAMON for MVS menu system interface is made possible by the support for dynamic terminal integration available with IBM Tivoli Monitoring V6.2.1. Dynamic terminal integration is an extension to the Tivoli Enterprise Portal that provides seamless access to TN3270-based applications through context-sensitive links.

  A Tivoli Enterprise Portal terminal view enables you to connect to any TN3270, TN5250, or VT100 host system with TCP/IP from inside a Tivoli Enterprise Portal workspace. For 3270 or 5250 terminal views, you also have scripting capability with record, playback, and authoring of entire scripts. By associating a terminal view with a connection script and a query that returns appropriate values, you can configure a view that opens to a specific panel of a 3270 application. This feature is useful for creating contextual workspace links for investigating issues.

  OMEGAMON XE on z/OS has taken advantage of this capability to create predefined links from several workspaces to target workspaces that contain a related OMEGAMON for MVS screen in a Terminal Emulator view. The data used to connect to the target screens is retrieved from environmental variables specified during configuration of OMEGAMON XE on z/OS monitoring agents using the Configuration Tool.

  Like the predefined situations provided with the product, these predefined links are intended as examples that you can build on to create your own links, using instructions found in the Tivoli Enterprise Portal help.

  For more information, see Chapter 4, “Using dynamic links to OMEGAMON for MVS,” on page 43.

- **Collection of sysplex-level shared DASD information turned off by default**

  Because of the large DASD volume counts that have become common in recent years, monitoring DASD devices without a filter that eliminates some of the devices can lead to high CPU or storage problems and even cause the monitoring server to fail. Consequently, the behavior of OMEGAMON XE on z/OS has been modified so that it does not collect DASD device data unless a DASD Filter is active. An auto-started warning situation (KM5_No_Sysplex_DASD_Filter_Warn) notifies you if no filtering situation is in place and no devices are being monitored.

  You can turn DASD data collection on by running a DASD filter situation. OMEGAMON XE on z/OS includes a model filter situation (KM5_Model_Sysplex_DASD_Filter), which uses the DASD Device Collection Filtering attributes Average Response Time and I/O Rate. By customizing this filter to exclude well behaved devices, you can enable monitoring of devices of particular interest and avoid being overwhelmed with unwanted data. A third product-provided situation, KM5_Weak_Plex_DASD_Filter_Warn, alerts you when too many devices are being monitored and filter criteria should be strengthened.

  See *IBM Tivoli OMEGAMON XE on z/OS: User's Guide* for instructions on creating a DASD filter situation.

  See “Filtering collection of DASD device data” on page 55 for instructions on creating a DASD filter situation.

- **New situation behavior**

  Beginning with V4.2.0, OMEGAMON XE on z/OS distributes and auto-starts only situations that are required to ensure the monitoring agents are running correctly. Situations that were distributed and auto-started in previous releases are no longer set to run automatically. The only situations currently set to run automatically at start up are KM5_No_Sysplex_DASD_Filter_Warn and KM5_Weak_Plex_DASD_Filter_Warn,

- **Reporting through Tivoli Common Reporting**

  OMEGAMON XE on z/OS includes reports that run under Tivoli Common Reporting, a reporting tool and strategy common across Tivoli products. Tivoli Common Reporting provides a consistent approach to viewing and administering reports. This reporting environment runs on Windows®, Linux®, and UNIX. For more information about Tivoli Common Reporting platforms, refer to the *Tivoli Common Reporting: User's Guide*.

  For more information about the OMEGAMON XE on z/OS reports, see *IBM Tivoli OMEGAMON XE on z/OS: User's Guide*.

  For more information, see Tivoli Common Reporting.

- **Integration with Log Analyzer**

  The product provides integration with Log Analyzer. Log Analyzer is a diagnostic tool included in the IBM Support Assistant (ISA) that links messages that are displayed in logs to information regarding workarounds and solutions that are available as APARs (Authorized Problem Analysis Reports). See the *IBM Tivoli OMEGAMON XE on z/OS: Troubleshooting Guide* for more information regarding ISA and the Log Analyzer.

Chapter 2. Using and customizing OMEGAMON XE on z/OS

This chapter is intended to familiarize you with the monitoring resources provided by OMEGAMON XE on z/OS and to help you use them to meet your specific requirements.

This chapter describes how to

- Use and customize predefined workspaces that report sysplex- and system-level data.
- Activate and customize predefined situations to enable alerts and reflex actions.
- Configure historical data collection and reporting.
- Issue UNIX commands using the Take Action feature.
- Launch Resource Management Facility Monitor III.
- Find more detailed information on workspaces, attributes, and situations.

The information on the organization and use of workspaces is useful for all OMEGAMON XE on z/OS users. The information on activating and modifying situations and configuring historical data collection is more useful for users with administrative authorities who are responsible for setting up and customizing monitoring and alerts.

Using the predefined workspaces

OMEGAMON XE on z/OS includes two sets of predefined workspaces: sysplex level workspaces and system (or LPAR) level workspaces. In the Tivoli Enterprise Portal, these workspaces are accessed either directly from the physical Navigator view or through links from other workspaces.

In the Navigator, the sysplex workspaces are listed under each managed sysplex name, and the system workspaces are listed under each z/OS managed system name. In a multiplex enterprise, the Navigator looks something like the following:

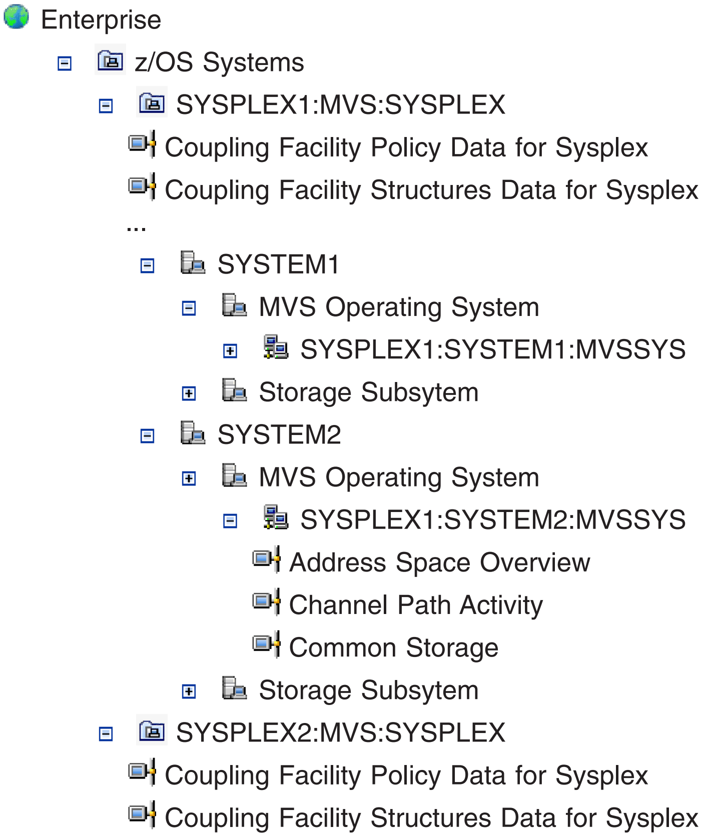

- Enterprise
  - ⊟ z/OS Systems
    - ⊟ SYSPLEX1:MVS:SYSPLEX
      - Coupling Facility Policy Data for Sysplex
      - Coupling Facility Structures Data for Sysplex
      - ...
      - ⊟ SYSTEM1
        - ⊟ MVS Operating System
          - ⊞ SYSPLEX1:SYSTEM1:MVSSYS
        - ⊞ Storage Subsytem
      - ⊟ SYSTEM2
        - ⊞ MVS Operating System
          - ⊟ SYSPLEX1:SYSTEM2:MVSSYS
            - Address Space Overview
            - Channel Path Activity
            - Common Storage
        - ⊞ Storage Subsytem
    - ⊟ SYSPLEX2:MVS:SYSPLEX
      - Coupling Facility Policy Data for Sysplex
      - Coupling Facility Structures Data for Sysplex

Most of the predefined workspaces are capable of reporting historical data. However, you must configure and start historical data collection in order for historical data to be available for reporting (see “Using historical data collection and reporting” on page 26).

Sysplex and system managed system names

From the standpoint of OMEGAMON XE on z/OS, sysplexes and systems are *managed systems*. In the Tivoli Enterprise Portal Navigator, managed systems are identified by *managed system names*.

Sysplex managed system names take the form:

```
plexname:MVS:SYSPLEX
```

where *plexname* is typically the true name of the sysplex, but might be configured to be an alias for the sysplex.

System managed system names take the form:

```
plexname:smfid:MVSSYS
```

where *plexname* is typically the true name of the sysplex, but could be configured to be an alias for the sysplex. (This part of the system managed system name typically matches the *plexname* component of its parent sysplex in the navigation tree.) The *smfid* component is the true System Management Facility (SMF) ID for the system or LPAR being monitored.

Organization of the workspaces

The physical Navigator view of the Tivoli Enterprise Portal shows an enterprise as a mapping of platforms, computers, agents, and monitored resources. Each Navigator item can be associated with one or more workspaces that provide information relevant to that level of the Navigator.

In a sysplex environment, monitored sysplexes appear between the platform and system levels of the Navigator tree, listed by their managed system names:

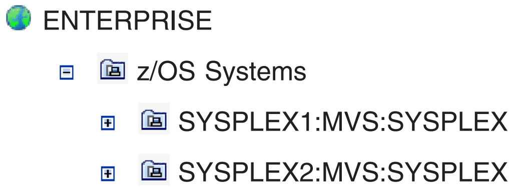

```
ENTERPRISE
  z/OS Systems
    SYSPLEX1:MVS:SYSPLEX
    SYSPLEX2:MVS:SYSPLEX
```

If OMEGAMON XE on z/OS is installed, the Sysplex Enterprise Overview is the default workspace for the **z/OS Systems** item in the Tivoli Enterprise Portal Navigator. As its name indicates, this workspace provides an overview of all the sysplexes in your enterprise. From the table views in this workspace you can link to sysplex-level workspaces for a selected sysplex. From the **z/OS Systems** Navigator item, you can also access the Cross-System Cryptographic Coprocessor Overview workspace.

Accessing the Cross-System Cryptographic Coprocessor Overview workspace

You can access the Cross-System Cryptographic Coprocessor Overview workspace in two ways, by linking from the Crypto Coprocessor Overview view of the Sysplex Enterprise Overview workspace, or from the **z/OS Systems**.

To access the workspace from the Navigator:

1. Select **z/OS Systems**.
2. After the Sysplex Enterprise Overview workspace is displayed, right-click **z/OS Systems** and select **Workspace** from the pop-up menu.

Each sysplex item in the Navigator is associated with a Sysplex Level Overview workspace, which provides summary data for the selected sysplex and links to more detailed workspaces. Below each sysplex item are items for the sysplex-level resources and components:

```
⊟ SYSPLEX1:MVS:SYSPLEX
   Coupling Facility Policy Data for Sysplex
   Coupling Facility Structures Data for Sysplex
   ...
   ⊞ SYSTEM1
   ⊞ SYSTEM2
```

Each item is associated with one or more workspaces that report information on resources shared by the sysplex or sysplex workloads. Each of these items has a default workspace, which opens when you select the item. This workspace may have may links to other related workspaces. For example, the default workspace for the **Shared DASD Groups Data for Sysplex** entry is the Shared DASD for Groups workspace, which displays information for all the groups in the sysplex. From this workspace you can link to a workspace that displays details for a selected group of DASD devices.

Beneath the sysplex items in the Navigator are items for every system (or LPAR) in the sysplex that is being monitored by an OMEGAMON XE monitoring agent. Each system has a default System Level Overview workspace. Beneath each system item in the Navigator tree is an item for each type of resource being monitored by a Tivoli OMEGAMON XE agent. For example, if you have installed both OMEGAMON XE on z/OS and OMEGAMON XE for Storage on z/OS, you will see entries for MVS Operating System and Storage Subsystem.

```
⊟ SYSTEM1
   ⊞ MVS Operating System
   ⊞ Storage Subsytem
```

If you expand **MVS Operating System** item, you see the managed system name of the system or LPAR monitored by the OMEGAMON XE on z/OS monitoring agent:

```
⊟ SYSTEM1
   ⊟ MVS Operating System
      ⊞ SYSPLEX1:SYSTEM1:MVSSYS
```

If you expand managed system entry, the workspaces that provide information about that system are listed:

```
⊟ SYSPLEX1:SYSTEM1:MVSSYS
   Address Space Summary
   Channel Path Activity
   Common Storage
```

As with the sysplex-level entries, every entry under the system name is associated with one or more workspaces. Each entry has a default workspace, which opens when you select the entry, and which may have other related workspaces you can access through links in table views in the workspace.

For more information about the organization of the product-provided workspace and the attribute groups associated with them, see “Organization of OMEGAMON XE on z/OS workspaces” on page 219.

Prerequisites for data reporting

Some workspaces or attributes display data only if specific conditions are met. See Table 1 for a list of these workspaces and conditions.

Table 1. Prerequisites for data display

Data is available in	Only if
Common storage workspaces	The Common Storage Area Analyzer (CSA Analyzer) is started. **Note:** The CSA Analyzer is shipped and installed with OMEGAMON XE on z/OS. It is configured as part of the configuration of the OMEGAMON II for MVS component and is started as a separate started task.
Channel Path Activity workspace	The Resource Measurement Facility (RMF) has been started.
GRS Ring Systems Data for Sysplex workspace	The global resource serialization (GRS) complex is in ring mode. (If the complex is in star mode, the workspace shows only the name, status, and ring acceleration of each system.)
Cryptographic workspaces	At least one IBM cryptographic coprocessor is installed and configured.
DASD MVS Workspace and DASD MVS Devices Workspace	The Resource Measurement Facility (RMF) is started.
Dynamic I/O device information	OMEGAMON Subsystem is running.
Tape Drives workspace	OMEGAMON Subsystem is running.
4 Hour MSUs attribute in the System CPU Utilization workspace	A defined capacity is used as a basis for pricing and the z/OS system is *not* running as a guest on z/VM®.
User Response Time workspace	The End to End (ETE™) Response Time collector is started. **Note:** The ETE response time collector is shipped and installed with OMEGAMON XE on z/OS. It is configured as part of the configuration of the OMEGAMON II for MVS component and is started as a separate started task. If you start the collector after you start the Tivoli Enterprise Monitoring Server, you must restart the monitoring server.
Workspaces showing zIIP and zAAP (IFA) data	Either: • System z Application Assist Processors (zAAP) and Integrated Information Processors (zIIP) are configured on the systems , or: • the **PROJECTCPU** control in the SYS1.PARMLIB IEAOPT*xx* member is specified as YES. **Note:** The **PROJECTCPU** control replaces the **-Xifa:force** control for zAAP data. zIIP and **PROJECTCPU** are supported in z/OS Version 1.8 at the base level, or in z/OS versions 1.6 and 1.7 via maintenance.
LPAR cluster workspaces	The z/OS system is not running as a guest on z/VM.
z/OS UNIX System Services attributes	The address space where the OMEGAMON XE on z/OS product is running has SUPER USER authority. This level of authority is equivalent to root (UID=0).

Table 1. Prerequisites for data display (continued)

Data is available in	Only if
Coupling facility, cross-system coupling facility, and system lock data collected by the Resource Management Facility (RMF) Distributed Data Server	You have enabled use of RMF data collection as described in *Tivoli OMEGAMON XE on z/OS Planning and Configuration Guide*, and RMF has been started. **Note:** If you enable RMF data collection, there will be no data in the Paths Workspace for CF System.

Customizing workspaces

You can modify the predefined workspaces in a number of ways. You can:

- Add, delete, or modify views.
- Modify queries.
- Apply thresholds or filters.
- Change the appearance of tables and charts.
- Add links to other workspace, or make a workspace accessible using a URL.

To modify a predefined workspace, create a copy of the workspace, save it with a different name, and then modify the copy. If you retain the original name, your customized workspace will be overwritten the next time you apply updates to the product.

Note: You must have Modify Workspace authority to modify workspaces.

If you have OMEGAMON DE on z/OS, you can create workspaces that include both mainframe and distributed sites, applications, and business processes.

Using cross-product workspace links

Dynamic workspace linking allows you to easily navigate between workspaces that are provided by multiple products. This feature aids problem determination and improves integration across the monitoring products, allowing you to quickly determine the root cause of a problem. Predefined cross-product links provided by the OMEGAMON XE products allow you to obtain additional information about systems, subsystems, resources or network components that are being monitored by other monitoring agents.

When you right-click on a link, a list of links is displayed. This list may contain links to workspaces provided by one or more monitoring products. The product you are linking to must be installed and configured and your Tivoli Enterprise Portal user ID must be authorized to access the target product in order for the link to that product's workspace to be included in the list.

Choose a workspace from the list to navigate to that workspace. You will link to the target workspace in context, meaning that you will receive additional information that is related to the system, subsystem or resource you are currently viewing.

If you choose a workspace from the list and the target workspace is not available, you will receive message KFWITM081E. Refer to the *IBM Tivoli OMEGAMON XE on z/OS: Troubleshooting Guide* for more information.

Predefined links in the Address Space Overview workspace for each managed system allow you to link to the Tivoli OMEGAMON XE for CICS on z/OS Region Overview workspace and the Tivoli OMEGAMON XE for Mainframe Networks Applications Connections workspace to obtain further information related to a selected address space.

You can also link to the Tivoli OMEGAMON XE for CICS on z/OS Region Overview workspace from the Address Space CPU Utilization workspace.

Dynamic workspace links in a mixed environment

If you are staging an upgrade from OMEGAMON XE V3.1.0 or V4.1.0, to OMEGAMON XE V4.2.0 products, you may have a combination of V3.1.0, 4.1.0, and V4.2.0 monitoring agents installed in your environment. For example, you may have an OMEGAMON XE on z/OS V4.1 monitoring agent and an OMEGAMON XE on z/OS V3.1 monitoring agent running on the same z/OS system during the migration period.

In this migration scenario, dynamic workspace linking from an OMEGAMON XE V4.2.0 workspace to an OMEGAMON XE V3.1 product workspace will work as long as the target workspace exists in the V3.1 product. If the target workspace does not exist, you will receive message KFWITM081E.

In cases where the V4.1.0 or V4.2.0 version of the target workspace has been modified (for example to accept link parameters to limit the data displayed) you may notice different behavior when you upgrade the target product.

Finding more information about workspaces

For more information on predefined workspaces, the views they contain, and the attribute groups on which they are based, see Chapter 11, “Workspaces,” on page 219. For more information about customizing workspaces, see *IBM Tivoli Monitoring: Administrator’s Guide* and the Tivoli Enterprise Portal online Help.

Using the predefined situations

To help you begin monitoring quickly, OMEGAMON XE on z/OS provides a number of predefined situations. These situations monitor for conditions that are typically considered to be problematic or noteworthy and trigger ⊗ Critical or ⚠ Warning event indicators in the Navigator when those conditions occur.

You must distribute and start the predefined situations before they can begin monitoring. You should evaluate each situation careful and adjust its thresholds before you start monitoring.

Activating situations

To activate a situation, you use the Situation editor in the Tivoli Enterprise Portal to distribute (assign) the situation to one or more managed systems or managed system lists and then start the situation. Each situation is already associated with an appropriate Navigator item. After you distribute a situation, you will see its name listed under the name of its associated item in the Situation editor.

Some system-level situations are shipped with very high or very low values, which essentially disable them. Others have values that may be inconsistent with the policies, goals, or monitoring requirements of your site. Examine the predefined situations and customize them with values that are meaningful for your installation before you activate them.

Distributing situations

Distribute only the situations that you are going to set to autostart or plan to manually enable. If you distribute all the situations, they will be propagated to the agents when the Tivoli Enterprise Monitoring Server starts. This may simplify any subsequent activation procedures, but it extends startup time. Review the situations to determine which ones you plan to use and add distribution lists for only those situations. Once the situations are distributed, their alerts will appear on the Navigator items they are associated with.

You distribute situations using the Tivoli Enterprise Portal Situation editor. To distribute a situation:

1. Open the Situation editor.

You can access the Situation editor from the toolbar or by right-clicking an item in the Navigator and selecting Situations from the pop-up menu.

2. If necessary, use **Set Situation filter criteria** to view the situations available for distribution.

   Check **Eligible for Association** to see a list of all the situations which are written for this type of managed system (MVS Sysplex or MVS System, depending on where you access the Situation editor from; if you access the editor from the toolbar, you see situations for all types of managed systems).

   Any undistributed situations show their icon partially dimmed .
3. Select (click) the situation you want to distribute.

   The Situation editor displays the **Condition** tab for the situation.
4. Select the **Distribution** tab.

   The available managed systems and managed systems lists are displayed.
5. Select the systems and lists to which you want to distribute the situation, and then click the left arrow to assign the situations to the systems or system lists.
6. Click **Apply** to save and implement the change and continue editing; click **OK** to apply and save the change and close the Situation editor.

Starting situations

Some situations you might want to run for a limited time or only under specific conditions. These situations you should start and stop manually. Other situations you may want to run continuously. These situations you should set to run at Tivoli Enterprise Monitoring Server startup, so they will run across Tivoli Enterprise Monitoring Server restarts.

Initially, you might want to start situations manually to evaluate the impact of the monitoring and monitoring interval on system performance, adjust them accordingly, and then decide if you want the situation to run indefinitely, across Tivoli Enterprise Monitoring Server restarts.

To start a situation, right-click the situation name in the Situation editor tree and select Start from the pop-up menu.

To set a situation to start automatically when the Tivoli Enterprise Monitoring Server starts:

1. Select (click) the name of the situation in the Situation editor tree.
2. The settings for the situation are displayed in the right-hand frame of the editor.
3. On the **Conditions** tab, check **Run at startup**.
4. Click **Apply** to save and implement the change and continue editing; click **OK** to apply and save the change and close the Situation editor.

Modifying situations

Before activating any predefined situations you should examine the conditions and values they monitor and, if necessary, adjust them to ones better suited to your environment.

To modify a situation:

1. Open the Situation editor from the toolbar, or right-click a Navigator entry and select Situations from the pop-up menu.

Tip

If you open the Situation editor by right-clicking a Navigator item, the situation you create is automatically associated with that item. If you open the editor from the toolbar, you must manually associate the new situation with a Navigator item in order to see an alert indicator when the situation evaluates as true.

2. Use the **Set Situation filter criteria** to view the situations.

   If necessary, check **Associated with Monitored Application** to see all situations that were written for this type of agent, regardless of where they are distributed.
3. To create a copy, right-click the situation and select **Create Another . . .** from the popup menu.
4. Type a name for the new situation and click **OK**.
5. Modify the situation properties as required and click **OK** to save the new situation and close the Situation editor.

Migrating situations

Preexisting situations built for the attribute groups CF Clients and CF Policy will not trigger events when coupling facility data is being collected by RMF. To enable these situations function, change the threshold values in the predicates slightly using the Tivoli Enterprise Portal Situation editor and then save and restart the situation. This causes the SQL for the situation to be rebuilt and in the process it switches to using the new table name. You can then re-edit the situation and set the threshold back to what it was previously.

Existing OMEGAMON XE for UNIX System Services situations are migrated when OMEGAMON XE on z/OS V4.2.0 product-provided situations are installed into the hub monitoring server. You may continue to use the existing situations, especially if you have customized them, in which case you should not start the 4.2.0 versions. Alternatively, you can delete the original situations and activate the V4.2.0 versions. Some of these situations have been renamed, as shown in Table 2.

Table 2. Old and new UNIX System Services situation names

OMEGAMON XE for UNIX System Services V220	OMEGAMON XE on z/OS V4.2.0
Check_Missing_Mount_Point	Check_Missing_UNIX_Mount_Point
Excess_Kernel_CPU_Time	Excess_UNIX_Kernel_CPU_Time
Excess_Process_UNIX_Run_Time	Excess_Process_UNIX_Run_Time
Excess_UNIX_System_Time	Excess_UNIX_System_Time
Excess_UNIX_User_Time	Excess_UNIX_User_Time
ENQ_Contention_Critical	UNIX_ENQ_Contention_Critical
ENQ_Contention_Warning	UNIX_ENQ_Contention_Warning
File_System_Free_Space_Critical	UNIX_File_System_FreeSpace_Crit
File_System_Free_Space_Warning	UNIX_File_System_FreeSpace_Warn
Logged_On_User_Idle	UNIX_Logged_On_User_Idle
Missing_inetd_Process	Missing_UNIX_inetd_Process
Quiecsed_File_System	Quiesced_UNIX_File_System
Shortage_of_Processes_Critical	Shortage_of_UNIX_Processes_Crit
Shortage_of_Processes_Warning	Shortage_of_UNIX_Processes_Warn
Unwanted_inetd_Process	Unwanted_UNIX_inetd_Process

A number of situations included in earlier versions of OMEGAMON XE on z/OS were renamed in V4.1.0 and may require special handling.

- Real storage

  The real storage situations listed in Table 3 have been renamed to reflect the fact that OMEGAMON XE on z/OS reports only central storage. If you are running with a mixed environment during a staged upgrade, the older situations will continue to work with V3.1 agents and will be migrated to your V4.2.0 environment, but you can use only the new versions with V4.2.0 agents. You cannot distribute the new situations to V3.1 agents. When you have completed your upgrade, you should activate the new situations on the upgraded nodes and delete the older versions. (They will run, but they will never become true.)

Table 3. Old and new real storage situation names

Old name	New name
OS390_CentralAvailFrames_Crit	OS390_Available_Frames_Crit
OS390_CentralAvailFrames_Warn	OS390_Available_Frames_Warn
OS390_CentralOnlineFrames_Crit	OS390_Frames_Online_Crit
OS390_CentralOnlineFrames_Warn	OS390_Frames_Online_Warn

  The following real storage situations are no longer included. They will continue to run on V3.1 agents, in a mixed environment, but you should delete them when you have completed upgrading to V4.2.0.

  OS390_ExpandedOnlineFrames_Crit
  OS390_ExpandedOnlineFrames_Warn
  OS390_Real_Stor_Migrate_Age_Crit
  OS390_Real_Stor_Migrate_Age_Warn
  OS390_Migration_Rate_Crit
  OS390_Migration_Rate_Warn OS390
  ExpandedToCentralStor_Crit
  OS390_ExpandedToCentralStor_Warn
  OS390_CentraltoExpandedStor_Crit
  OS390_CentraltoExpandedStor_Warn

- Because of architectural changes, the names of the situations listed in Table 4 were shortened in V4.1.0. You should replace the old situations with the renamed versions and delete them. If you continue to use the older situations, when the situations evaluate to true, you will see correct data in the Initial Situation values column of the Situation Event Console, but not in the Current Situation column.

Table 4. Situations renamed for architectural reasons

Old name	New name
OS390_System_PageFault_Rate_Crit	OS390_System_PageFaultRate_Crit
OS390_System_PageFault_Rate_Warn	OS390_System_PageFaultRate_Warn
OS390_Channel_LPAR_Busy_Pct_Crit	OS390_Channel_LPAR_BusyPct_Crit
OS390_Channel_LPAR_Busy_Pct_Warn	OS390_Channel_LPAR_BusyPct_Warn
OS390_Cache_FastWrite_HitPt_Crit	OS390_Cache_FastWriteHitPt_Crit
OS390_Cache_FastWrite_HitPt_Warn	OS390_Cache_FastWriteHitPt_Warn
OS390_Common_PageDS_PctFull_Crit	OS390_Common_PageDSPctFull_Crit
OS390_Common_PageDS_PctFull_Warn	OS390_Common_PageDSPctFull_Warn
OS390_Tape_Permanent_Errors_Crit	OS390_Tape_Permanent_Error_Crit
OS390_Tape_Permanent_Errors_Warn	OS390_Tape_Permanent_Error_Warn

Finding more information about situations

For descriptions of all the predefined situations shipped with OMEGAMON XE on z/OS, including definitions and advice, see Chapter 12, "Situations," on page 279.

You can find more information on creating and modifying situations in *IBM Tivoli Monitoring: User's Guide* and in the Tivoli Enterprise Portal online Help.

Using historical data collection and reporting

In addition to monitoring and displaying real-time data, OMEGAMON XE on z/OS can log data to binary data sets so you can examine data for longer periods of time.

You can view the logged historical data in OMEGAMON XE on z/OS workspaces. Table and chart views for which historical data collection is enabled have a tool for setting a time span. You can see up to 24 hours of previously collected data. If you have configured data warehousing, you can view samples for longer periods of time.

In order for historical data to be available in workspaces, you must configure and start historical data collection for the appropriate attribute groups using the Tivoli Enterprise Portal (see "Configuring historical data collection"). In addition, to store short-term data historical data for agents running on z/OS or reporting to monitoring servers on z/OS, data sets must be allocated in the persistent data store and maintenance of the data store must be configured. To warehouse data, DB2 or Microsoft® SQL Server must be installed and your environment must be configured to include the Warehouse Proxy agent and Tivoli Data Warehouse. For information on setting up the persistent data store and configuring maintenance, see the *OMEGAMON XE and Tivoli Management Services on z/OS: Common Planning and Configuration Guide* and the *IBM Tivoli OMEGAMON XE on z/OS: Planning and Configuration Guide*. For information about installing and setting up the Tivoli Data Warehouse and the Warehouse Proxy agent, see *IBM Tivoli Monitoring: Installation and Setup Guide*.

You can also export the logged historical data to delimited flat files for use with third-party reporting tools to produce trend analysis reports and graphics. Data warehoused to the Tivoli Data Warehouse, a relational database, can be used to produce customized history reports.

Configuring historical data collection

You configure historical data collection using the History Collection Configuration dialog box in the Tivoli Enterprise Portal (Figure 3 on page 27).

Configuration is done on an attribute-group by attribute-group basis. This allows you to configure collection for different attribute groups at different intervals so important volatile data may be collected more often, while less dynamic data can be collected less frequently.

Not all attribute groups can collect historical data. This is because collecting history data for these attribute groups is not appropriate or would have a detrimental effect on performance. For example, collection might generate unmanageable amounts of data. Only those attribute groups for which data can be collected are listed in the Configuration dialog box.

Note that for a given attribute group, the same history collection options are applied to all Tivoli Enterprise Monitoring Servers for which collection for that attribute group is currently enabled. You cannot specify different intervals for the same attribute group for different monitoring servers.

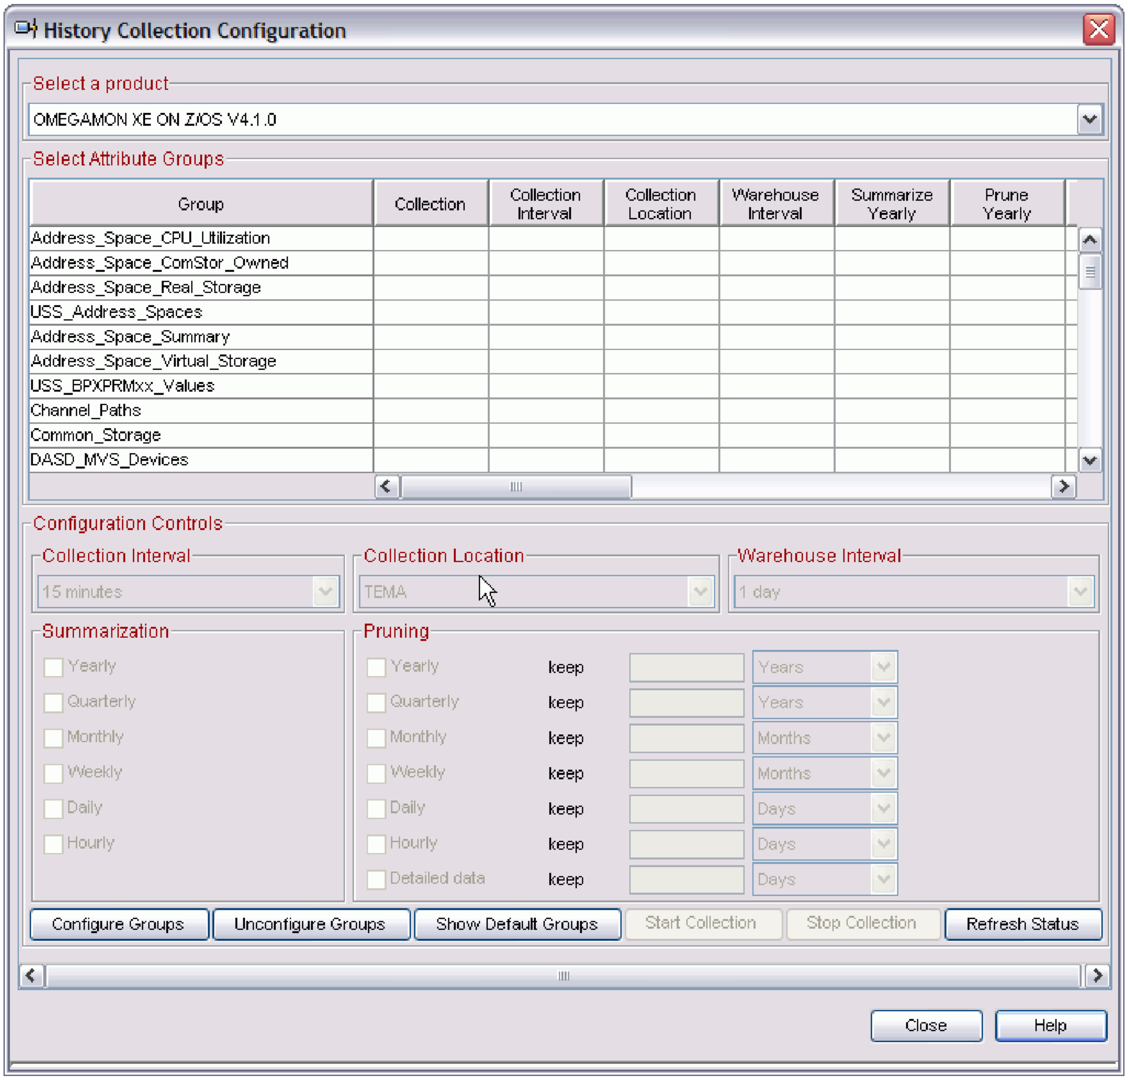

Select a product

OMEGAMON XE ON Z/OS V4.1.0

Select Attribute Groups

Group	Collection	Collection Interval	Collection Location	Warehouse Interval	Summarize Yearly	Prune Yearly
Address_Space_CPU_Utilization						
Address_Space_ComStor_Owned						
Address_Space_Real_Storage						
USS_Address_Spaces						
Address_Space_Summary						
Address_Space_Virtual_Storage						
USS_BPXPRMxx_Values						
Channel_Paths						
Common_Storage						
DASD_MVS_Devices						

Configuration Controls

Collection Interval: 15 minutes

Collection Location: TEMA

Warehouse Interval: 1 day

Summarization

- □ Yearly
- □ Quarterly
- □ Monthly
- □ Weekly
- □ Daily
- □ Hourly

Pruning

□ Yearly	keep		Years
□ Quarterly	keep		Years
□ Monthly	keep		Months
□ Weekly	keep		Months
□ Daily	keep		Days
□ Hourly	keep		Days
□ Detailed data	keep		Days

Figure 3. History Collection Configuration dialog box

Starting and stopping data collection

You start and stop historical data collection for individual attribute groups from the History Collection Configuration dialog box.

In the **Select Attribute Group** table, select the attribute group or groups for which you want to change collection status, then press the appropriate button. Collection continues until the agent or monitoring server is stopped or recycled.

Reducing the impact of requests from large tables

Requests for historical data from tables that collect a large amount of data have a negative impact on the performance of the product components involved. To reduce the performance impact on your system, set a longer collection interval for attribute groups that collect a large amount of data, in particular the Address Space groups, the DASD MVS Devices group, and the Enqueue group (for sites that are active with WebSphere®). You specify the collection interval from the Configuration tab of the History Collection Configuration dialog.

Note: No data is collected for sysplex-level shared DASD unless a DASD filter situation has been activated. This is to prevent high CPU and storage problems caused by the large amount of data that may be generated with large DASD volume counts. For more information, see Chapter 5, "Monitoring shared DASD," on page 55.

When you are viewing historical data, set the Time Span interval to the shortest time span setting sufficient to provide the information you need, especially for tables that collect a large amount of data. Selecting a long time span interval for the report time span increases the amount of data being processed, and may have a negative impact on performance. The program must dedicate more memory and processor cycles to process a large volume of report data.

If the amount of information requested is too large, the agent may take too long to process the request and the request may time out. However, the agent continues to process the report data to completion, and remains blocked, even though the report data is not viewable.

Using summarization and pruning can also help reduce the amount of historical data you store and request. For more information about using these features, see the *IBM Tivoli Monitoring: User's Guide* and the *IBM Tivoli Monitoring: Administrator's Guide*.

Tivoli Common Reporting reports

OMEGAMON XE on z/OS V4.2.0 includes reports that run under Tivoli Common Reporting , a reporting tool and strategy common across Tivoli products. Tivoli Common Reporting provides a consistent approach to viewing and administering reports. This reporting environment runs on Windows, Linux, and UNIX. For more information about Tivoli Common Reporting platforms, see the *Tivoli Common Reporting: User's Guide*. To learn more about how the OMEGAMON XE monitoring agents on zSeries use Tivoli Common Reporting, see the *OMEGAMON XE and Tivoli Management Services on z/OS: Reports for Tivoli Common Reporting*.

The reports provided with OMEGAMON XE on z/OS are historical reports run against data collected on a DB2–based Tivoli Data Warehouse V6.2 Fix Pack 1 or later running on Windows. Reports for OMEGAMON XE on z/OS are provided as *report packages*, compressed files containing reports, documentation, graphics, and dynamic link libraries. The OMEGAMON XE on z/OS report package is included as a .zip file on the application data DVD in the REPORTS directory, and the REPORTS directory is divided into subdirectories named with the three-character prefix that identifies the product (the prefix for OMEGAMON XE on z/OS is KM5). For example, on a Windows machine, if the DVD or CD drive is labelled D:, reports are in directories such as: D:\REPORTS\KM5. *These reports are not installed when application support is installed. You must obtain the reports from the media.*

Using the reports

Before you can use the reports in the report package, perform the following steps:

- Install Tivoli Common Reporting, using the information found in the *Tivoli Common Reporting: User's Guide*.
- Ensure that your environment meets the requirements described in the "Prerequisites" section of the *OMEGAMON XE and Tivoli Management Services on z/OS: Reports for Tivoli Common Reporting*.
- Perform the setup functions described in the *OMEGAMON XE and Tivoli Management Services on z/OS: Reports for Tivoli Common Reporting*
  - Ensuring that historical reporting is enabled
  - Importing the report package
  - Configuring the data source
  - Generating a sample report

  See the "Troubleshooting" section of the *OMEGAMON XE and Tivoli Management Services on z/OS: Reports for Tivoli Common Reporting* book if you have any problems performing these tasks.

To develop reports of your own or edit the reports provided with this product, you will need a Eclipse BIRT Report Designer, which is a noncharge, open-source tool. This tool is not included with Tivoli Common Reporting, but can be downloaded from http://www.eclipse.org/birt/phoenix/ or from the Tivoli Common Reporting page at IBM developerWorks® (http://www.ibm.com/developerworks/spaces/tcr. You will also need the *Tivoli Common Reporting: Development and Style Guide* on the IBM developerWorks Web site: http://www.ibm.com/developerworks/spaces/tcr.

Typical types of reports available with Tivoli Common Reporting

A report can either be run on demand or a snapshot can be created at any time for viewing later. An *on-demand* report is a formatted version of a report run with the currently available data and viewed immediately. When you run an on-demand report, you can specify the parameter values to use; the report is then formatted using the current data and displayed in the report viewer.

In addition to creating on-demand reports, you can also create, save, and access *snapshot* reports, saved versions of a report using data from a particular time. Snapshots are useful for generating unattended reports in advance. Snapshots are of particular value in avoiding the wait for a big report to run. Reports may take a long time to run when you request a huge amount of data, such as the past month of real time measurements. Refer to the *IBM Tivoli Common Reporting User's Guide* for information about this report type or look in the online help for Tivoli Common Reporting.

Because Tivoli Management Services supports the summarization and pruning of data, many OMEGAMON XE packaged reports can also generate *summarized* reports. If a packaged report supports summarized data and the Summarization Agent has been configured to collect data for the attribute group required (at the Tivoli Enterprise Portal console), then selected reports will provide an option to specify a summarization period (Hourly, Daily, Weekly, Monthly, Quarterly, or Yearly). The resulting reports reflect data summarized at that level. You can change the summarization period dynamically from a generated report without completing the parameters definition dialog again by clicking a different summarization period in the Available Summarization Time Periods area of a summarized report. For more information about summarization and pruning, see the *IBM Tivoli Monitoring: Administrator's Guide*.

Some OMEGAMON XE summarized reports in PDF and HTML format may also have embedded *drill-through* reports available. *Drill-through* refers to an embedded link from one report to another report that provides additional detail. You can open a drill-through report by clicking a point in a line graph or a bar in a bar graph. Table Group headings can also be selected for drill-through actions.

Note: If you are not seeing drill-through reports for your snapshot reports, first check the reports in this chapter to ensure that this report type is available. If it is available but is not displaying the correct data, check the setting on the **Maximum Levels Of Drill-through Snapshots** property. See the online help for information about the **Properties** dialog box.

OMEGAMON XE on z/OS includes the following reports:

- z/OS Status DASD Issues
- z/OS Status WLM Service Class Period
- z/OS Utilization Common Storage
- z/OS Utilization LPAR Summary
- z/OS Utilization LPAR Usage Summary
- z/OS Utilization Real Storage

Each of these reports is described in *OMEGAMON XE and Tivoli Management Services on z/OS: Reports for Tivoli Common Reporting* in tables that provide the following information:

- A description of the report
- The most logical users of this report
- The default attribute graphed in the report
- Other attributes from the same attribute group and workspace that you could graph instead

- Resource identifiers that form the matrix for the report and are not selectable
- Other resource identifiers you can specify that act as filters for the data shown in the report. The choices in this cell usually correspond to the choices in the Report Parameters dialog.
- The name of the workspace that displays the same data as this report
- The attribute group or table in Tivoli Data Warehouse that is the source of the attributes
- The types of reports available

Finding more information about historical data collection

For more information on configuring historical data collection and reporting in Tivoli Enterprise Portal, see the Tivoli Enterprise Portal online Help and *IBM Tivoli Monitoring: User's Guide*.

For more information on allocating data sets and configuring the persistent data store, see *IBM Tivoli OMEGAMON XE on z/OS: Planning and Configuration Guide*.

For information on maintaining the persistent data store, exporting historical data to flat files, and warehousing historical date, see *OMEGAMON XE and Tivoli Management Services on z/OS: Common Planning and Configuration Guide*and the *IBM Tivoli Monitoring: Administrator's Guide*.

For information on configuring the Tivoli Data Warehouse, the Warehouse Proxy Agent, and the Summarization and Pruning Agent, see *IBM Tivoli Monitoring: Installation and Setup Guide*.

Issuing UNIX commands from OMEGAMON XE on z/OS

You can use the Tivoli Enterprise Portal Take Action feature to enter a command or stop or start a process on any system in your network where one or more monitoring agents are installed, and you can add Take Action commands to any situations you create using OMEGAMON XE on z/OS attributes.

By default, any command issued on behalf of OMEGAMON XE on z/OS is issued as a z/OS command. However, prefixing a command with one of the following options causes the command to be issued as a UNIX command. Note that the colon (:) following the command is required.

- **OMVS:, Omvs:, or omvs:**
- **UNIX:, Unix:, or unix:**

Thus, for example, `D OMVS` is issued as a z/OS command. Alternatively, `omvs:ps -ef` is issued as a UNIX command. As with z/OS commands, the only result returned is whether or not the command appears to have started successfully.

Using one of the prefixes you can issue a UNIX program name as a command. You can also issue UNIX shell commands..

Important

The user ID of the address space where the OMEGAMON XE on z/OS product is defined (the Tivoli Enterprise Monitoring Server address space) must be defined to z/OS UNIX System Services and have superuser authority in order to collect UNIX data and relay UNIX commands. See "Authorizing users to issue UNIX commands" on page 32 and *Tivoli OMEGAMON XE on z/OS Planning and Configuration Guide* for more information on defining users to UNIX System Services. In addition, Tivoli Enterprise Portal user IDs must be authorized to issue UNIX commands. See "Authorizing users to issue UNIX commands" on page 32.

Environment for issuing UNIX commands

UNIX commands are issued in an environment having these characteristics:

- No terminal is available. Some commands, such as `ps` issued without options, will not work in this environment since they are designed to have a current terminal. Note that, in the case of `ps`, some options, for example, `-A`, enable `ps` to execute without requiring a terminal.
- The specified command may be any shell command or UNIX program, including any REXX™ program written for the shell, that does not require a terminal or a particular environmental condition (for example, a specific user ID).
- stdin is initially assigned to `/dev/null` (the equivalent of an empty file).
- stdout and stderr are initially assigned to `/dev/console` (the z/OS SYSLOG).
- stdin, stdout, and stderr can be redirected following standard UNIX redirection conventions.
- The shell is executed in a new process in a separate address space. This insulates OMEGAMON XE on z/OS from the effects of the command and vice versa.
- The shell has the same authorization as OMEGAMON XE on z/OS.
- The initial working directory is /.
- The HOME environment variable is set to /.
- No other environment variables are set before the shell is started. `shell/bin/sh` is used to issue the specified command.
- The shell performs normal login profile processing starting with `/etc/profile` before issuing the specified command. This can result in further environment variables being set before the specified command is issued.

  Note: If profile processing terminates the shell before the specified command has been issued, for example, by issuing the exit shell command, the specified command is not issued.
- The specified command will not terminate when OMEGAMON XE on z/OS terminates if OMEGAMON XE on z/OS terminates first.
- Termination, abnormal or otherwise, of the specified command will not cause OMEGAMON XE on z/OS to terminate.
- OMEGAMON XE on z/OS does not maintain or report any status information about the shell or specified command other than that collected as part of its normal system monitoring functions.

Redirecting UNIX commands output

By default, the output of both z/OS and UNIX commands are written to the z/OS system log. You cannot redirect the output of a z/OS command. However, you can redirect both the input stream and output of a UNIX command by following standard UNIX redirection conventions. For example, the command `omvs:ps -ef>/tmp/myoutput` sends the output of the `ps` command to a file called `/tmp/myoutput`. Redirect command output to a file for later examination and to avoid cluttering the z/OS log.

Running commands in the background

Each UNIX command is run as a process in a separate address space using the */bin/sh* shell. When OMEGAMON XE on z/OS is used to start a long-running UNIX command, you may notice an address space that persists until the command ends. This address space is in addition to the one running OMEGAMON XE on z/OS and the one running the command itself. You can avoid the extra address space by running the command in the background. To do so, end the command line with the UNIX shell symbol & (ampersand).

Testing commands

Before using commands as part of situations or policies, you should run some simple tests from the user interface to ensure that commands are working as expected. For example, issuing the `omvs:set>/tmp/cantest` Take Action command results in the output of the set command being placed in the `/tmp/cantest` file.

Authorizing users to issue UNIX commands

By default, only user IDs that have been defined to z/OS UNIX System Services and that have superuser, or root, authority are allowed to issue UNIX commands through the Tivoli Enterprise Portal. Users are defined to z/OS UNIX using RACF® commands. The z/OS UNIX attributes are kept in the OMVS segment of the RACF user's profile.

This means that to issue UNIX commands:

- The user's Tivoli Enterprise Portal user ID must be defined in RACF.
- The profile associated with the RACF user ID must contain an OMVS segment.
- In the OMVS segment, the z/OS UNIX user identifier (UID) must have a value of 0 (superuser).

You can override the default validation behavior by adding one of two parameters to the KDSENV member of *&shilev.&rtename.*RKANPARU on the system or LPAR on which the command is being executed:

- You can allow any RACF user ID defined to z/OS UNIX System Services to issue UNIX commands, regardless of level of authorization, by adding the variable KOE_ALLOW_ANY_UID=1 to *&shilev.&rtename*.RKANPARU(KDSENV) on the LPAR where the command is to be executed.
- You can allow any RACF user ID to issue UNIX commands, whether or not it has been defined to z/OS UNIX System Services, by adding the variable KOE_ALLOW_UNDEFINED=1 to *&shilev.&rtename*.RKANPARU(KDSENV) on the LPAR where the command is to be executed.

If you want any user with a Tivoli Enterprise Portal user ID to be able to issue UNIX commands, add both KOE_ALLOW_ANY_UID=1 and KOE_ALLOW_UNDEFINED=1 parameters.

Tip

For commands used in situations and policies, the user ID verified is the ID of the user who last saved the situation.

Launching the RMF Monitor III data portal

The Resource Measurement Facility (RMF) is an IBM-licensed program that measures selected areas of system activity and presents the data collected in the format of printed reports, System Management Facility (SMF) records, or display reports. RMF is used to evaluate system performance and identify reasons for performance problems. RMF Monitor III is a browser-based interactive monitor that collects data and reports contention for resources and their users. The data allows identification of system bottlenecks and determination of the reasons for possible system performance degradations.

OMEGAMON XE on z/OS provides a launch of the Monitor III data portal from the Sysplex Enterprise Overview workspace. You enable the launch by providing the location of the browser on the client host system and the URL of the Monitor III Web interface. You can configure the launch using the Create or Edit Launch Definitions window in the Tivoli Enterprise Portal, or from a command line using the **LaunchConfig** command.

Configuring the launch from Tivoli Enterprise Portal

To configure the launch using the Tivoli Enterprise Portal graphical user interface:

1. In the Tivoli Enterprise Portal Navigator, select **z/OS Systems**.

   The Sysplex Enterprise Overview workspace is displayed.
2. Right-click the **z/OS Systems**item and select **Launch** from the pop-up menu.

   The Create or Edit Launch Definitions window opens.
3. Type the name of the new link you are creating, for example `RMFIII`.
4. In the **Target** field, type the completely qualified path to the browser you want to use. For example:

   ```
   C:\Program Files\Internet Explorer\IEXPLORE.EXE
   ```

5. In the **Arguments** field, enter the URL of the Monitor III Data Portal. For example:

   ```
   http://sys1.itdept.anywhere.company.com:27703/
   ```

6. Press OK to close the window and save the configuration. Press Launch to launch a browser window and connect to the Monitor III data portal.

Configuring the launch from the command line

To configure the launch from the command line:

1. Access a command-line interface.
2. Run the **km5LaunchConfig** command to insert, change, or remove target points for the launch feature.

   Note: The command syntax must include the path for the **km5LaunchConfig** command, as in the following examples for a Windows installation of IBM Tivoli Monitoring in the default installation path. Alternatively, you can run the command from the directory where the code for the command is stored. (The default installation path on Windows is `C:\IBM\ITM\CNP\`.)

   - To add a new launch target use the following flags:

     ```
     C:\IBM\ITM\CNP\km5LaunchConfig -insert -loc=location
            -url=url
     ```

     where *location* is the fully qualified path to the browser application (for example, Microsoft Internet Explorer) on the client system and *url* is the URL of the Monitor III Data Portal. Both arguments must be enclosed in double quotes (" ").

     For example:

     ```
     C:\IBM\ITM\CNP>km5LaunchConfig -insert -loc="C:\Program Files\Internet Explorer\IEXPLORE.EXE"
          -url="http://sys1.itdept.anywhere.company.com:27703/"
     ```

   - To make a change or add a launch target use the following flags:

     ```
     C:\IBM\ITM\CNP\km5LaunchConfig -change -loc=LOCATION -url=URL
     ```

     where *LOCATION* is the path to the browser application (for example, Microsoft Internet Explorer) on the client system and *URL* is the URL of the Monitor III Data Portal.

     Enclose both arguments in double quotes (" ").

   - To remove all of the launch target points that were created by this command, use the following flags:

     ```
     C:\IBM\ITM\CNP\km5LaunchConfig -remove
     ```

Chapter 3. Using the Inspect function

This chapter discusses the Inspect function. It explains what Inspect does, how you invoke it, and how you specify the number of samples it collects and the sampling interval at which it collects them.

This chapter also describes the Inspect Address Space CPU Use workspace and how you can use it to identify inefficient code, or to determine where, within the code, a program might be looping.

About the Inspect function

Inspect is a diagnostic tool whose primary purpose is to help you understand where, within an address space, code is spending its time. You can then use that information either to optimize the code, or to identify where within the code a program might be looping. You might use the Inspect function, for example, when a workspace or situation event shows an address space with high processor usage.

You can specify the number of samples to be collected and the interval at which they are collected. The results are displayed in the Inspect Address Space CPU Use workspace (see “About the Inspect CPU Usage workspace” on page 38).

Accessing Inspect data

You access the Inspect data from the Address Space CPU Utilization workspace by linking from a row of the Address Space CPU Utilization table.

The Inspect Address Space CPU Use is displayed, populated with Inspect data for the selected address space. The parameters Inspect uses to collect the data are specified in the link definition.

OMEGAMON XE on z/OS provides two Inspect links:

- Inspect Address Space CPU Use

  This link uses the default parameters of 1000 samples at 5 millisecond intervals.
- Inspect with 5000 samples at 2ms interval

  This link uses the specified parameters (5000 samples at 2 millisecond intervals). It is intended to be used as a template to set your own sample count and sampling interval.

Modifying the template Inspect link

If you want to specify the sample count and sampling interval used to collect Inspect data, modify the template Inspect link (Inspect with 5000 samples at 2ms interval).

> **Tip**
>
> The Inspect Address Space CPU Use workspace is not populated with data until the Inspect agent completes on the target system. The time it takes to complete is a function of the number of samples and the sampling interval specified in the link definition. For example, taking 1000 samples at a 5-millisecond interval (the default settings) requires 5 seconds for the data collection process to complete. When you are selecting the values for number of samples and the sampling intervals, bear in mind that if the total time taken to execute the agent exceeds the client timeout value, the Tivoli Enterprise Portal will return no data, even if the agent subsequently completes normally.

To modify the template Inspect link:

1. Open Address Space Overview workspace for the target system.
2. In the **Address Space Counts** table, right-click the 🔗 link button and select **Address Space CPU Utilization** from the pop-up window.

3. After the Address Space CPU Utilization workspace is displayed, right-click the 🔗 link button beside a row in the **Address Space CPU Utilization** table and select **Link Wizard** from the pop-up menu.

   The Link Wizard editor is displayed.

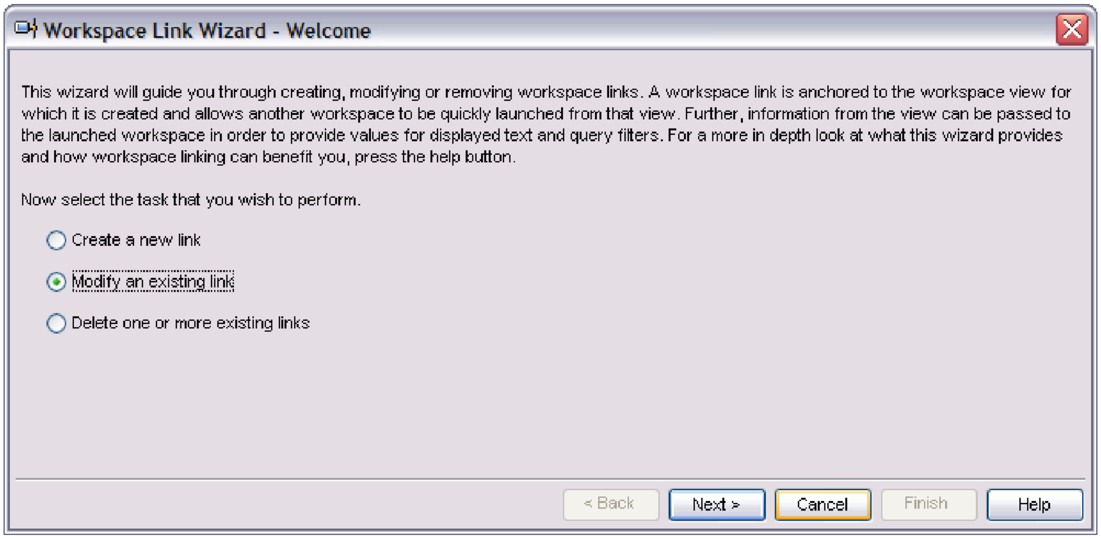

Workspace Link Wizard - Welcome

This wizard will guide you through creating, modifying or removing workspace links. A workspace link is anchored to the workspace view for which it is created and allows another workspace to be quickly launched from that view. Further, information from the view can be passed to the launched workspace in order to provide values for displayed text and query filters. For a more in depth look at what this wizard provides and how workspace linking can benefit you, press the help button.

Now select the task that you wish to perform.

○ Create a new link

● Modify an existing link

○ Delete one or more existing links

4. Select **Modify an existing link**, and then click **Next**.
5. In the Link to Modify window, select `Inspect with 5000 samples at 2ms interval`, and then click **Next**.

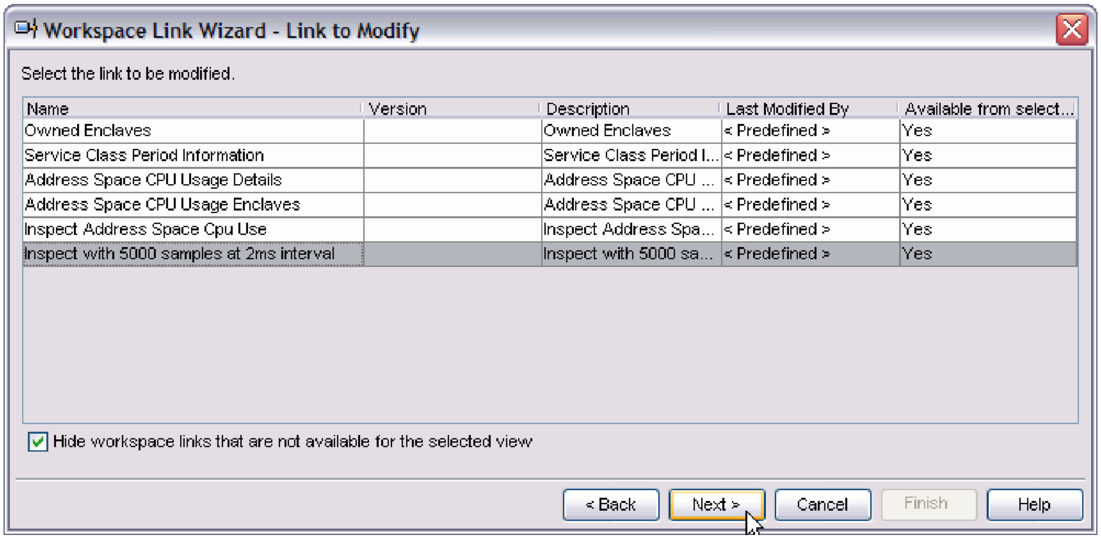

Workspace Link Wizard - Link to Modify

Select the link to be modified.

Name	Version	Description	Last Modified By	Available from select...
Owned Enclaves		Owned Enclaves	< Predefined >	Yes
Service Class Period Information		Service Class Period I...	< Predefined >	Yes
Address Space CPU Usage Details		Address Space CPU ...	< Predefined >	Yes
Address Space CPU Usage Enclaves		Address Space CPU ...	< Predefined >	Yes
Inspect Address Space Cpu Use		Inspect Address Spa...	< Predefined >	Yes
Inspect with 5000 samples at 2ms interval		Inspect with 5000 sa...	< Predefined >	Yes

☑ Hide workspace links that are not available for the selected view

6. In the **Link Name** window, type a new name for the link and description, if desired, and then click **Next**.

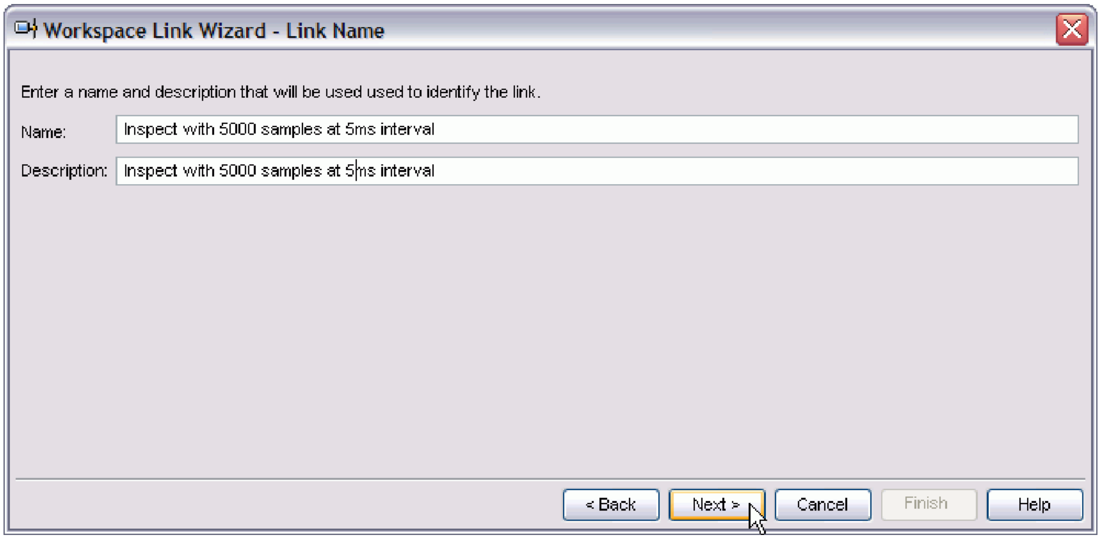

Workspace Link Wizard - Link Name

Enter a name and description that will be used used to identify the link.

Name: Inspect with 5000 samples at 5ms interval

Description: Inspect with 5000 samples at 5ms interval

7. In the Parameters window, select the v INTERVAL row, and then click **Modify Expression**.

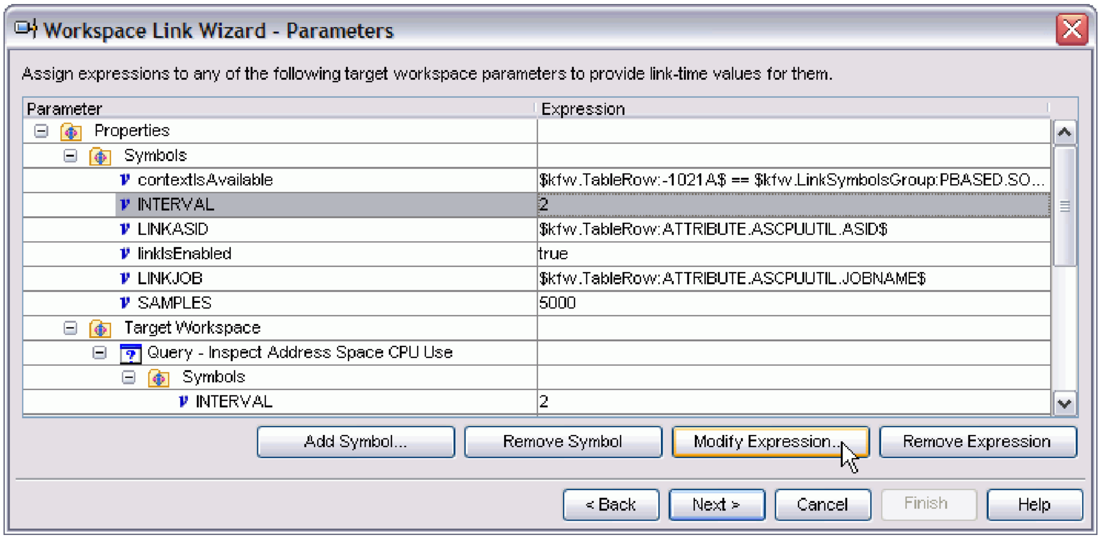

Workspace Link Wizard - Parameters

Assign expressions to any of the following target workspace parameters to provide link-time values for them.

Parameter	Expression
⊟ Properties	
⊟ Symbols	
v contextIsAvailable	$kfw.TableRow:-1021A$ == $kfw.LinkSymbolsGroup:PBASED.SO...
v INTERVAL	2
v LINKASID	$kfw.TableRow:ATTRIBUTE.ASCPUUTIL.ASID$
v linkIsEnabled	true
v LINKJOB	$kfw.TableRow:ATTRIBUTE.ASCPUUTIL.JOBNAME$
v SAMPLES	5000
⊟ Target Workspace	
⊟ Query - Inspect Address Space CPU Use	
⊟ Symbols	
v INTERVAL	2

8. In the Expression editor, type the interval, in milliseconds, at which you want Inspect to collect samples, and then click **OK**.

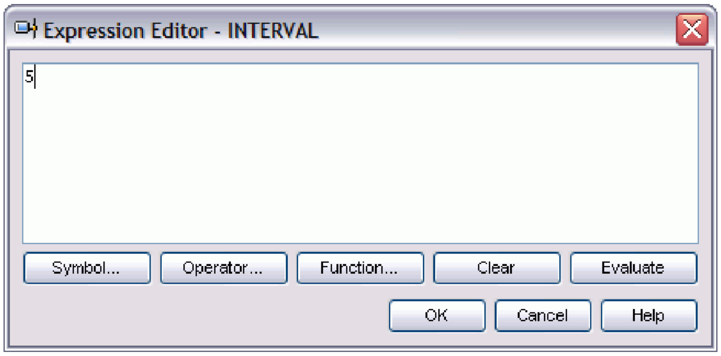

Expression Editor - INTERVAL

5

9. In the Parameter tree, select the SAMPLES row, and then click **Modify Expression**.

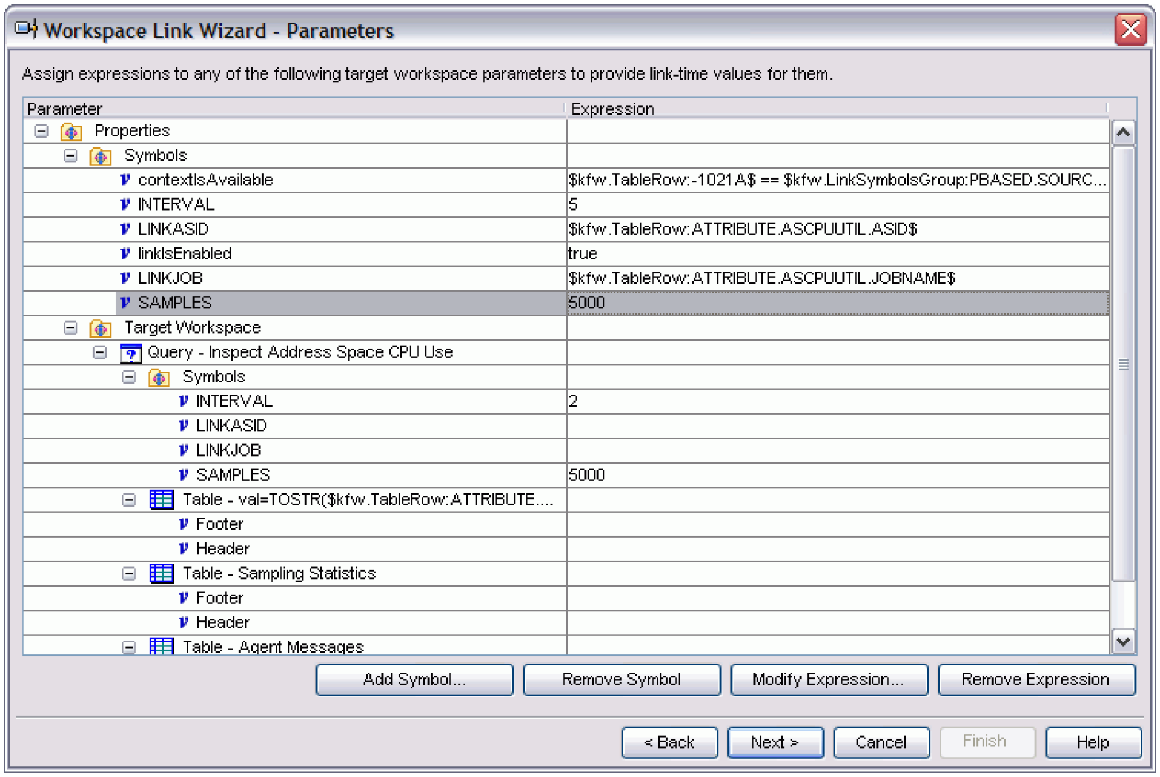

Assign expressions to any of the following target workspace parameters to provide link-time values for them.

Parameter	Expression
⊟ Properties	
⊟ Symbols	
contextIsAvailable	$kfw.TableRow:-1021A$ == $kfw.LinkSymbolsGroup:PBASED.SOURC...
INTERVAL	5
LINKASID	$kfw.TableRow:ATTRIBUTE.ASCPUUTIL.ASID$
linkIsEnabled	true
LINKJOB	$kfw.TableRow:ATTRIBUTE.ASCPUUTIL.JOBNAME$
SAMPLES	5000
⊟ Target Workspace	
⊟ Query - Inspect Address Space CPU Use	
⊟ Symbols	
INTERVAL	2
LINKASID	
LINKJOB	
SAMPLES	5000
⊟ Table - val=TOSTR($kfw.TableRow:ATTRIBUTE....	
Footer	
Header	
⊟ Table - Sampling Statistics	
Footer	
Header	
⊟ Table - Agent Messages	

10. In the Expression editor, type the number of samples you want Inspect to use in deriving the data, and then **OK**.

    The Parameters window is re-displayed with the changes.

11. Click **Next**.

    The wizard asks you to review the changes. If they are correct, click **Finish** to save the changes and close the wizard.

You see the name you assigned in the pop-up menu when you right-click a link button in the Address Space CPU Utilization table.

The new name, if any, and the parameters you set persist until you close the workspace.

About the Inspect CPU Usage workspace

The Inspect Address Space CPU Usage workspace contains three views:

- "Sampling Statistics view"
- "Agent Messages view" on page 39
- "Inspect Data view" on page 39

Sampling Statistics view

The four columns of this table view show the number of samples requested, the sampling interval in milliseconds, the number of samples collected, and the number of samples used.

Normally the number of samples collected will be same as the number requested. However, if the job being inspected ends before the Inspect agent has finished collecting data, the number of samples collected will be the number collected up to the point where Inspect detected that the target job had ended.

The number of samples used is the number of times that the Inspect agent saw CPU activity in the target address space and gives you some indication as to the statistical accuracy of the resultant inspect data. The number of samples used value does *not* represent the number of rows of Inspect data.

Agent Messages view

This view displays any error or informational messages returned by the Inspect agent. These messages help to explain the resultant data (or lack thereof) that you see in the other views. For example, if no CPU activity was seen by Inspect in the address space being inspected, the agent returns a message indicating that; the number of samples used column in the Sampling Statistics view would be zero; and no data would be displayed in the Inspect Data view.

Inspect Data view

The Inspect Data view contains the output from the inspection process. The Inspect agent returns data only for elements for which it saw CPU activity. The data is ordered in descending CPU activity order with the following hierarchy:

Task control block (TCB)
Load module within TCB
Control sector (CSECT) within load module
Block of code within CSECT
Task control block

For each TCB for which it sees processor activity, Inspect attempts to determine the executing load modules consuming the processor time (CPU). For each load module, Inspect then attempts to map the CSECT structure of the load module and assign the load module processor time to the appropriate CSECTs. This allows you to determine which load modules and CSECTs within the load module are consuming processor time.

Inspect maps the CSECT structure for each load module (with processor activity) by scanning the target address space for load libraries and attempting to read in the load module from each library in turn. It also scans SYS1.LINKLIB, SYS1.NUCLEUS and SYS1.LPALIB for load modules.

If Inspect cannot locate the load module, the CSECT name is unknown and the entire load module is considered to be one large CSECT.

Inspect then further breaks down the CPU time attributed to each CSECT into blocks of code, the size of which is calculated by Inspect once the data collection process has completed. In order to prevent Inspect from flooding the client workspace with rows of data, the Inspect agent attempts to calculate a granularity (block) size that will limit the number of rows of data returned to about 100, but where possible the Inspect agent will use a granularity size of 16 bytes (0X00000010). The granularity size used is displayed in the Agent Messages view of the Inspect CPU Use workspace. The granular data is shown in the rightmost two columns of the Inspect Data view. Again, these are displayed in descending processor use order with the most active blocks of code within each CSECT being at the top of the display rows for each CSECT.

For descriptions of the information in the columns of this table, see “Inspect Address Space CPU Use attributes” on page 166.

Using the Inspect data to understand a problem

Before examining the data, review the Samples Used field in the Sampling Statistics view. This field indicates the statistical validity of the sampled data. A low number of samples indicates that the Inspect data may not give a truly representative view of where the code in the target address space is spending its time.

Also review any messages in the Agent Messages view that may indicate that the data might be incomplete, or explain why there is no data. You can find descriptions of the Inspect messages in the online help.

Understanding the data

In the Inspect Data view, the data is organized in descending processor (CPU) by task control block (TCB) order, so the most active items are at the top of the display. In the example shown below, Inspect saw one TCB active. The program that was attached to create this TCB was called KLV. Because Inspect only saw one TCB active, 100% of the processor time that Inspect saw being used is attributed to this TCB. The TCB Ended column is blank, indicating that this TCB did not end while Inspect was running.

Inspect saw two load modules in use during its sampling activities, KM5AGENT and IEAVXPCA. Since KM5AGENT is at the top, this was the most active load module. However, you can also see that the second load module spent some of its time executing within the PCAUTH address space, as shown by the Load Module ASID Hex and Load Module Jobname columns.

TCB Address	Initial Program	TCB CPU % of job	TCB Ended	Load Module Name	Load Module ASID Hex	Load Module ASID Jobname
0X007F83D0	KLV	100.0		KM5AGENT		
				IEAVXPCA	0X0002	PCAUTH

The Load Module Address column shows where in storage in the target address space the load module is loaded. Following that are two columns which show how the CPU time used by this load module breaks down as a percentage of all the CPU time used by the address space (Load Module CPU % Job) and of this TCB (Load Module CPU % TCB). Because there is only one TCB in this instance, these numbers are the same for each load module. (If there were more TCBs active you see the CPU time used by each load module as a percentage of both the overall job and its owning TCB.)

Load Module Address	Load Module CPU % of TCB	Load Module CPU % of job
0x12E8F670	73.3	73.3
0X123020A0	26.6	26.6

To the right, the CSECT Name column shows that in this instance all the CPU time seen by Inspect was being consumed by one CSECT, in this case KM3PBM1.

The CSECT Offset in Load Module shows amount of offset within the load module of this CSECT, and the CSECT Address column shows the address of the CSECT within the target address space storage. You could use this information to locate the CSECT within a dump of the address space, for example to confirm an eye catcher that might contain a compile date or time or some other version information.

In the next columns, the CPU time attributed to the CSECT is displayed as a percentage of the total CPU time for the job, the load module, and the TCB. This can help you understand how much the CSECT is being used overall, by each load module, and by each task within the address space. This can help you to identify the CSECTs that are most heavily used.

CSECT Name	CSECT Offset in load module	CSECT Address	CSECT CPU % of Load Module	CSECT CPU % of TCB	CSECT CPU % of Job
KM3PBM1	0X000A9580	0X12F38BF0	100.0	73.3	73.3
IEAVXRFE	0X0005AB50	0X12307B50	100.0	26.6	26.6

In the final columns, the granular level data is shown. This breaks each CSECT down into blocks of code, based on the granularity size calculated by Inspect. The granularity size used is shown in the Agent Messages view. This allows you to further refine where within each CSECT the code is spending its time and identify areas that may be looping or that might benefit from optimization.

Offset in CSECT	CPU% of CSECT
0X00000AE0	36.3
0X00000BD0	33.0
0X00000CD0	30.6
0X00000AA0	100.0

To understand how the granular level data relates to the program source (section of code) you can then refer to a link editor and compile listings.

Chapter 4. Using dynamic links to OMEGAMON for MVS

A Tivoli Enterprise Portal terminal view enables you to connect to any TN3270 host system with TCP/IP from inside a Tivoli Enterprise Portal workspace. These terminal views also have scripting capability with record, playback, and authoring of entire scripts. If you associate a terminal view with a connection script and a query that returns appropriate values, you can configure a view that opens to a specific panel of a 3270 application. OMEGAMON XE on z/OS has taken advantage of this capability to create dynamic links from several existing workspaces to target workspaces that display related OMEGAMON for MVS screens in a terminal emulator view.

The data used to connect to OMEGAMON for MVS and navigate to the target screens is retrieved from variables specified during configuration of OMEGAMON XE on z/OS monitoring agents using the Configuration Tool, and attributes identified in the dynamic link.

Like the product-provided situations, you can use the predefined workspaces, queries, and scripts used to create these links as models or templates for creating additional links. You can also create your own links and target workspaces from scratch.

This chapter covers the following topics:

- "How dynamic linking works"
- "Product-provided links" on page 44
- "Creating new dynamic links" on page 45

How dynamic linking works

Dynamic linking involves the following components:

- A source workspace that provides the context (attribute values) for the link
- A target workspace that contains a terminal emulator view associated with a query and a suitable connection and navigation script
- A dynamic link from the source workspace to the target workspace that specifies the attribute values from the source workspace that are used for navigation to the appropriate screen.
- A query that retrieves additional required values from the OMEGAMON XE agent to be used by the script associated with the terminal adapter
- A script that uses the values from the query and the link to connect to OMEGAMON for MVS session and navigate to the appropriate screen

When a dynamic link on a row in the source workspace is selected, the link picks up the values of the attributes specified in the link definition from the row and passes them to the target workspace. The query associated with the terminal view in the target workspace is executed to obtain connection information. The query is processed by the monitoring agent and discovered or configured information is passed back to the terminal emulator.

When a Tivoli Enterprise Portal user attempts to connect to a session, a TN3270 logon window prompts for a user ID and password. These credentials can be saved for the duration of the Tivoli Enterprise Portal session so they will not have to be prompted for again. They are lost when the Tivoli Enterprise Portal client is terminated. If the script associated with the session is modified to include the user ID and password for the OMEGAMON session, or if no user ID or password is required for the logon, the window can be dismissed without them.

Note: The stored user credentials are associated with the specific OMEGAMON session being linked to, as defined by its host name, port number, LU Group and APPLID. If the same terminal emulator (and workspace) is used to connect to a different classic instance, the user is prompted again for another user ID and password.

After the user clicks **OK** to close the prompt, the startup script file is run. The script file retrieves all the query values, link values, and any values modified during logon and uses this information to drive the 3270 interface to connect to the APPLID returned in the query and to navigate to the appropriate display in OMEGAMON for MVS.

The connection values for the OMEGAMON for MVS session (host name, port number, Logical Unit (LU) Group and APPLID) are discovered by the monitoring agent. However, these values can be overridden using Configuration Tool. The host to which the TN3270 session connects must have an active TN3270 listener. By default, that host is assumed to be the LPAR on which the monitoring agent is located. If there is no active listener, the address of an LPAR that does have an active listener must be specified. The default port number for the Telnet listener is 23. This value can also be overridden. The Dynamic XE to 3270 (Classic) linking feature requires the VTAM® Unformatted System Services (USS) screen to accept a `LOGON APPLID() DATA()` command. If the default Telnet USS screen does not accept this command, the name of a Logical Unit (LU) group that does accept it must be provided. The TN3270 session will be joined to that LU group. The default values are overridden on the Add Runtime Environment (3 of 3) or Update Runtime Environment (3 of 3) panel.

The default values or the override values used for the session are displayed at the TN3270 logon and can be modified for an individual TN3270 session.

The terminal connection terminates after 5 minutes of inactivity or when you end this Tivoli Enterprise Portal work session.

Product-provided links

OMEGAMON XE on z/OS provides seven launch points to four target workspaces (Table 5). These links save you several time-consuming steps and the cutting and pasting of data required to navigate to desired OMEGAMON for MVS screens.

For example, suppose an alert is raised by OMEGAMON XE on z/OS in the Tivoli Enterprise Portal indicating there is a common storage shortage of CSA. Instead of logging on to an OMEGAMON for MVS TN3270 session and drilling down to the appropriate information, you can review the relevant information in the OMEGAMON XE on z/OS workspaces and then log on to OMEGAMON for MVS via TN3270 to display orphaned CSA storage and release it, if appropriate, freeing up this valuable resource and avoiding CSA shortage problems that might result in an IPL.

OMEGAMON for MVS screens that require level 3 authority or are high CPU users were avoided as targets of these links. Like the product-provided situations these links are meant to be useful examples that can be built on (see “Creating new dynamic links” on page 45).

Table 5. Product-provided dynamic links

From	To
Address Space Common Storage – Orphaned Elements workspace	OMEGAMON for MVS – CSA Analyzer workspace
LPAR Clusters workspace	OMEGAMON for MVS – LPAR PR/SM™ Processor Statistics workspace
System CPU Utilization workspace	OMEGAMON for MVS – License Manager MSU and WLM Capping
Address Space Overview workspace Address Space Bottlenecks Summary workspace Address Space CPU Utilization workspace Address Space Storage workspace	OMEGAMON for MVS – Job Details workspace

Creating new dynamic links

The process of creating a dynamic OMEGAMON XE to 3270 link can be summarized as follows:

- A target workspace is created with one or more terminal adapter (3270 emulator) views.
- A script is created and associated with the terminal emulator view.
- A query is associated with the terminal adapter to fetch connection values from the monitoring agent.
- The link is created from the source report to the target workspace that includes the values from the source report that will be needed for the script invoked by the terminal adapter.

This section contains instructions for creating new links by modifying product-provided workspaces and scripts:

1. "Create the target workspace."
2. "Modify the associated script" on page 46.
3. "Define the dynamic link" on page 48

If you want to create new workspaces and scripts, follow the instructions in the Tivoli Enterprise Portal online help or the *IBM Tivoli Monitoring: User's Guide*. Any links you create should use the Dynamic Link to 3270 query (see "Assigning the Dynamic Link to 3270 query" on page 51).

Tip

You must have Workspace Author Mode permission to create the target workspaces and the dynamic links. If you want to share the links and workspaces you create, you must have Workspace Administration Mode permission.

Create the target workspace

Complete the following steps to create a target workspace based on the OMEGAMON for MVS - Job Details workspace:

1. In the Navigator, open the Address Space Overview workspace.
2. In the Address Space CPU Utilization Summary table, right-click the 🔗 link in the first row and select OMEGAMON for MVS - Job Details.

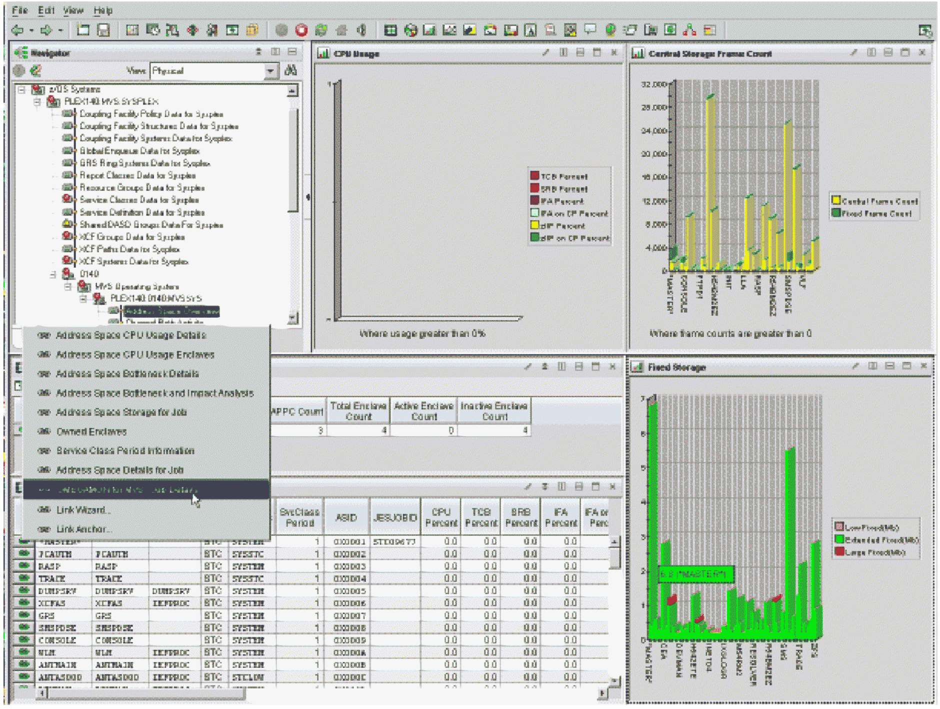
File Edit View Help
Navigator
CPU Usage
Central Storage Frame Count
Address Space CPU Usage Details
Address Space CPU Usage Enclaves
Address Space Bottleneck Details
Address Space Bottleneck and Impact Analysis
Address Space Storage for Job
Owned Enclaves
Service Class Period Information
Address Space Details for Job
Link Wizard...
Link Anchor...
Where usage greater than 0%
Where frame counts are greater than 0
Fixed Storage

The OMEGAMON for MVS - Job Details workspace opens, and a prompt asks for logon credentials for the TN3270 session.

3. Cancel the prompt, then select **Save as** from the **File** menu.
4. In the Save Workspace As window, provide a name and description for the workspace and check **Only selectable as the target of a Workspace Link**, then click **OK**.

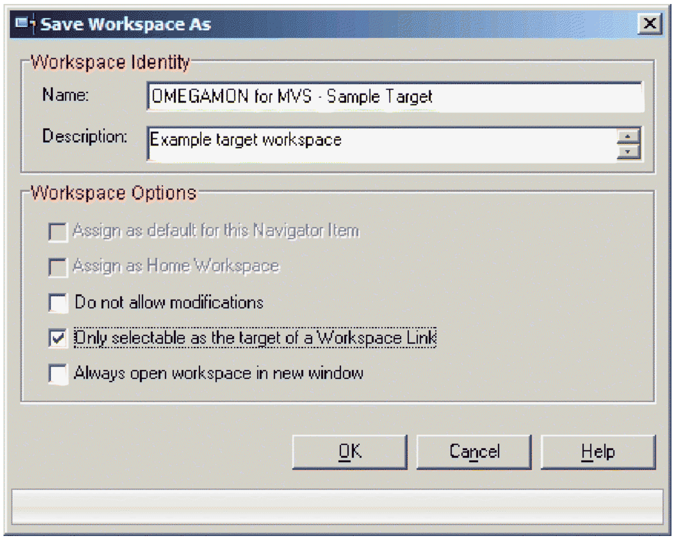

Save Workspace As

Workspace Identity

Name: OMEGAMON for MVS - Sample Target

Description: Example target workspace

Workspace Options

☐ Assign as default for this Navigator Item

☐ Assign as Home Workspace

☐ Do not allow modifications

☑ Only selectable as the target of a Workspace Link

☐ Always open workspace in new window

You have now created the target workspace. The terminal view in this workspace is already associated with Dynamic Link to 3270 query, so you do not need to assign it to the terminal view. However you must modify the script associated with the terminal view (see “Modify the associated script”). Leave the properties window open; otherwise you will have difficulty locating the new target workspace.

Modify the associated script

In this step you modify the script already associated with the new target workspace.

You can record or write your own scripts, using the instructions in the Tivoli Enterprise Portal online help or the *IBM Tivoli Monitoring: User's Guide*, but you will probably want to use the connection and signon code from an existing script.

Scripts can be global or local. You can only access local scripts through the Properties window of the terminal view with which they are associated.

The scripts associated with terminal views that are the target of dynamic links have essentially two parts: the first part handles the connection and logon, the second half handles the navigation to a specific OMEGAMON for MVS screen. This section of the script uses the symbol names assigned to the attributes to fetch the values it uses to navigate to the relevant screen. This is probably the only part of the script that you will need to modify.

To modify the script, take the following steps:

1. In the Properties window of the terminal view, select the **Configuration** tab (if necessary), then select the **Scripts** tab at the bottom.

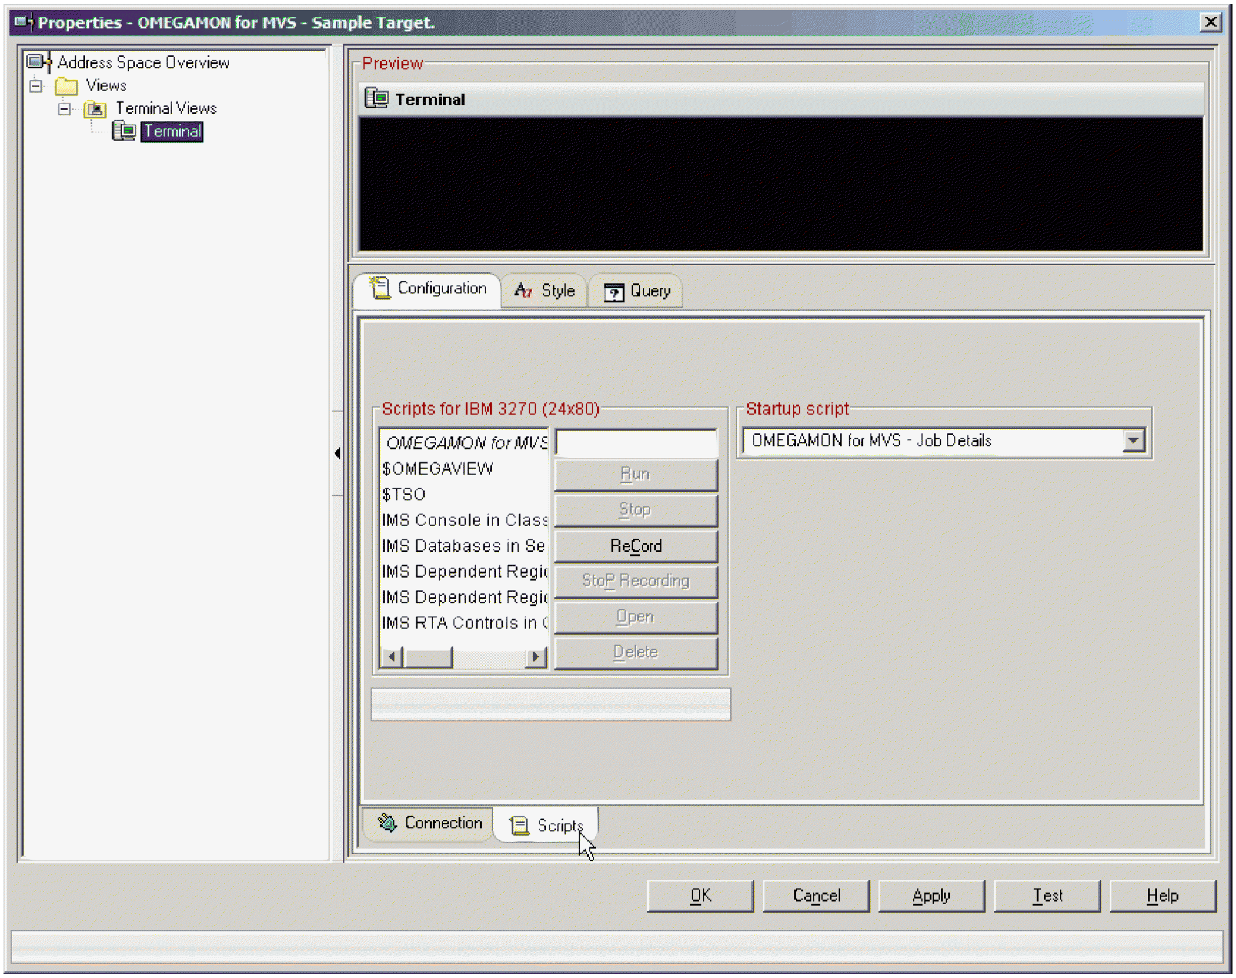

2. In the Scripts for IBM 3270 (24x80) list, select OMEGAMON for MVS, then select **Open**.

   The script is displayed in an editing window.

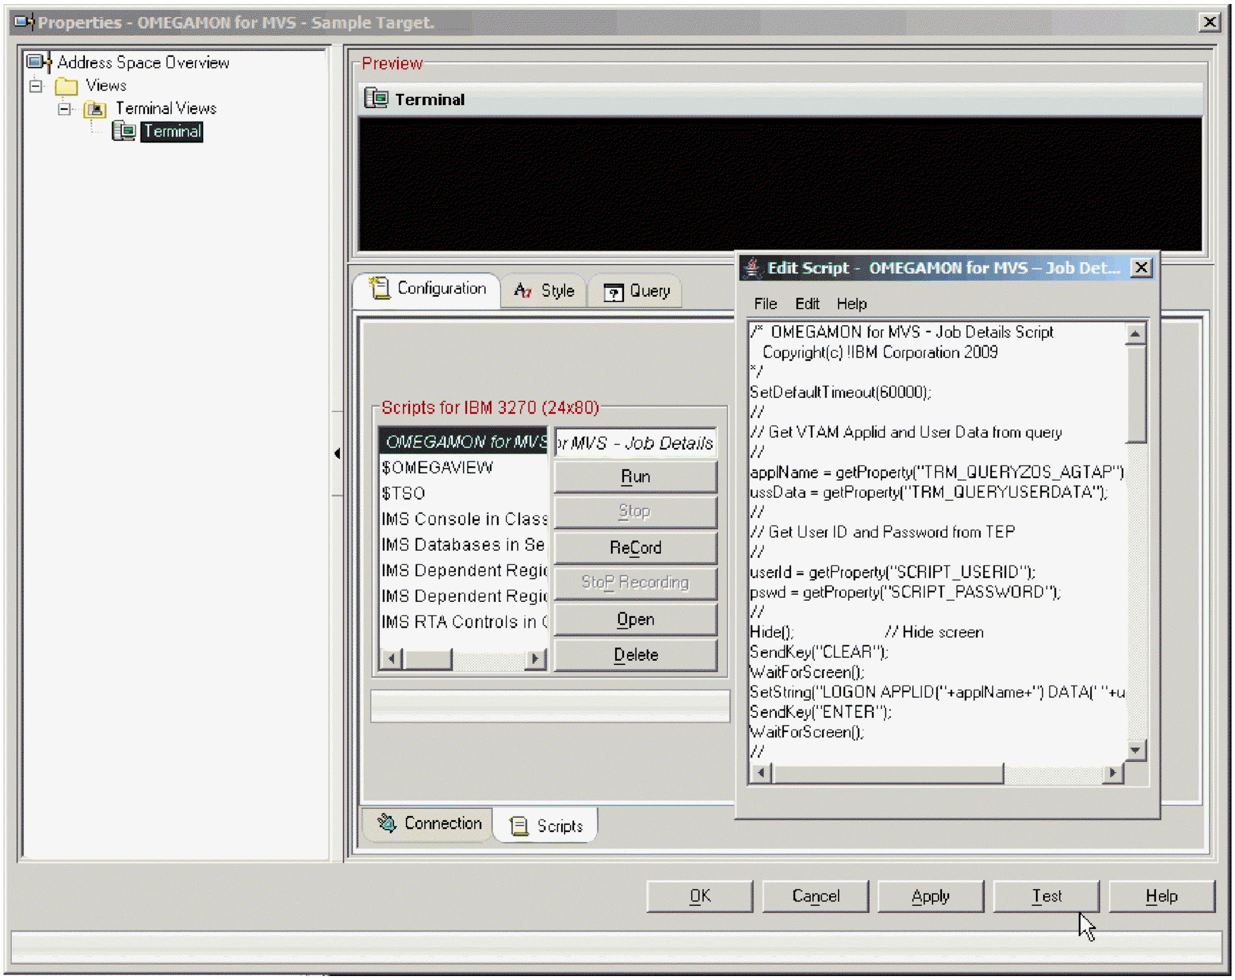

```
/* OMEGAMON for MVS - Job Details Script
  Copyright(c) !IBM Corporation 2009
*/
SetDefaultTimeout(60000);
//
// Get VTAM Applid and User Data from query
//
applName = getProperty("TRM_QUERYZOS_AGTAP")
ussData = getProperty("TRM_QUERYUSERDATA");
//
// Get User ID and Password from TEP
//
userId = getProperty("SCRIPT_USERID");
pswd = getProperty("SCRIPT_PASSWORD");
//
Hide();                  // Hide screen
SendKey("CLEAR");
WaitForScreen();
SetString("LOGON APPLID("+applName+") DATA(' "+u
SendKey("ENTER");
WaitForScreen();
//
```

3. Look for the comment:

   ```
   // We are at the OMEGAMON for MVS Main Menu
   ```

   Edit the script after this point as necessary to reflect the navigation path and symbol names that the script will need to retrieve from the link.
4. Select **Save As** from the **File** menu and name the modified script and specify a time-out value, then click **OK**.
5. Click **OK** to close the Properties editor.

Now you are ready to define the dynamic link.

Define the dynamic link

In this step you create a link to the target workspace. When you define the link, you identify the target workspace and specify the attributes whose values are passed to the script to use to navigate to the relevant OMEGAMON for MVS screen.

Follow these steps to create a link between a source OMEGAMON XE on z/OS workspace and a related OMEGAMON for MVS screen:

1. Open the workspace that you want to link from.
2. Right-click in a row in the table view from which you want to create the link and select **Link to > Link Wizard** (if there are no existing links in the row) or **Link Wizard** (if there are existing links) from the popup menu.

3. On the Link Wizard Welcome window, select **Create a new link** and click **Next**.
4. On the Link Name window, type a **Name** and **Description** to identify the link and click **Next**.
5. On the Link Type window, select **Dynamic** as the link type and click **Next**.
6. In the Navigator pane of the Target Workspace window, expand the branches and select the Address Space Overview node. ("Cloned" workspaces are listed under the node associated with the workspace they are cloned from; in this case, the OMEGAMON for MVS - Job Details workspace, which is associated with the Address Space Overview node.)

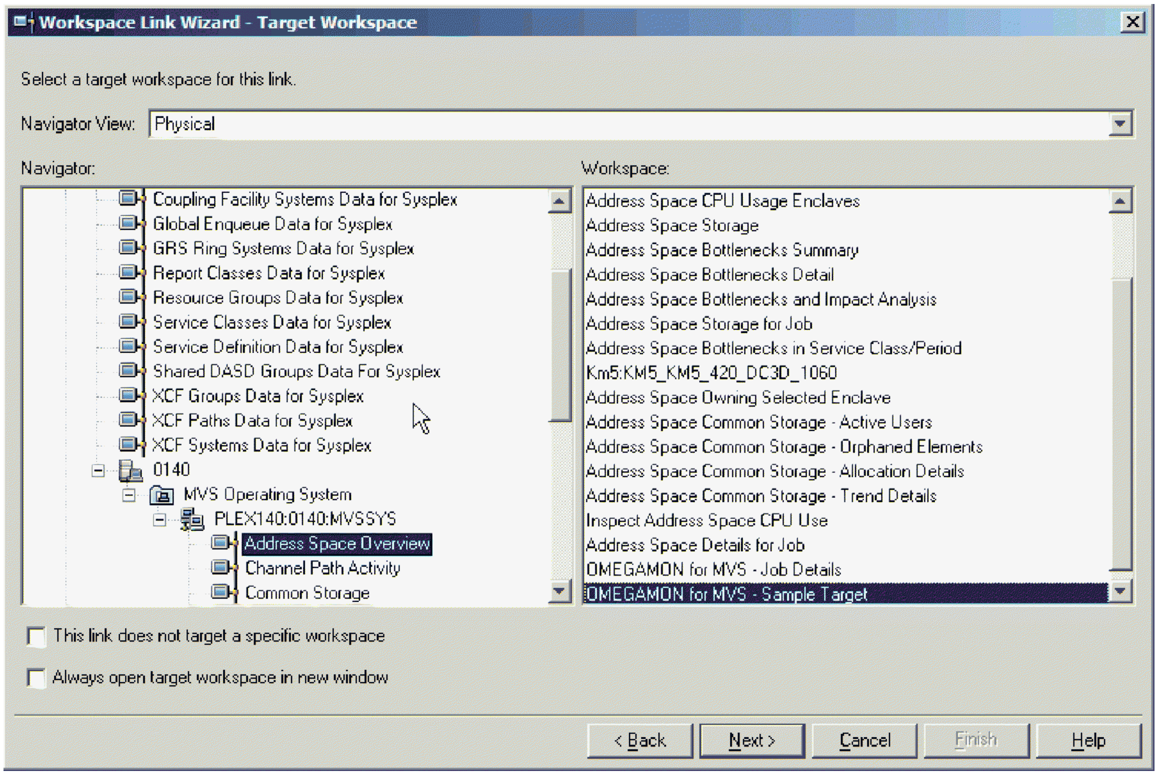

Workspace Link Wizard - Target Workspace

Select a target workspace for this link.

Navigator View: Physical

Navigator:

- Coupling Facility Systems Data for Sysplex
- Global Enqueue Data for Sysplex
- GRS Ring Systems Data for Sysplex
- Report Classes Data for Sysplex
- Resource Groups Data for Sysplex
- Service Classes Data for Sysplex
- Service Definition Data for Sysplex
- Shared DASD Groups Data For Sysplex
- XCF Groups Data for Sysplex
- XCF Paths Data for Sysplex
- XCF Systems Data for Sysplex
- 0140
  - MVS Operating System
    - PLEX140:0140:MVSSYS
      - Address Space Overview
      - Channel Path Activity
      - Common Storage

Workspace:

- Address Space CPU Usage Enclaves
- Address Space Storage
- Address Space Bottlenecks Summary
- Address Space Bottlenecks Detail
- Address Space Bottlenecks and Impact Analysis
- Address Space Storage for Job
- Address Space Bottlenecks in Service Class/Period
- Km5:KM5_KM5_420_DC3D_1060
- Address Space Owning Selected Enclave
- Address Space Common Storage - Active Users
- Address Space Common Storage - Orphaned Elements
- Address Space Common Storage - Allocation Details
- Address Space Common Storage - Trend Details
- Inspect Address Space CPU Use
- Address Space Details for Job
- OMEGAMON for MVS - Job Details
- OMEGAMON for MVS - Sample Target

□ This link does not target a specific workspace

□ Always open target workspace in new window

7. In the Workspace pane, select the workspace to target and click **Next**.
8. In the Target Filters window, select Managed System Name and click **Modify Expression** to open the Expression Editor.
9. In the Expression Editor, click **Symbol** to open the Symbols list.
10. Expand the Agent node and select Node.

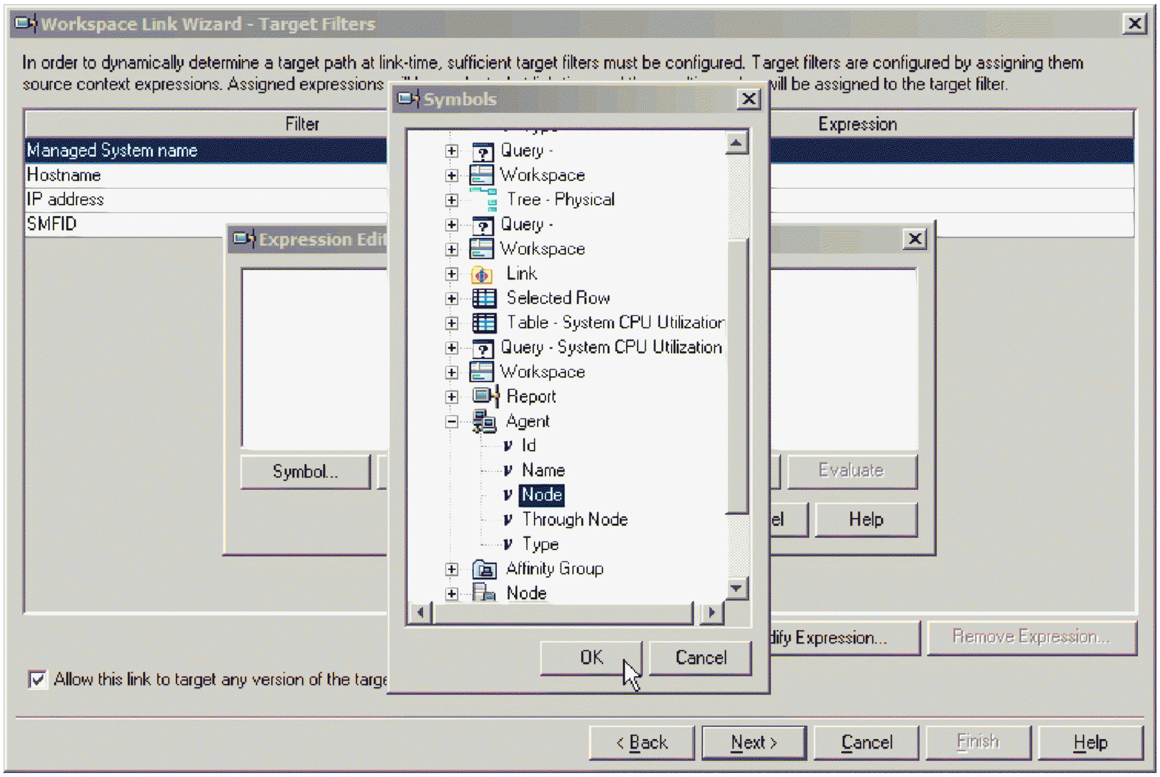

In order to dynamically determine a target path at link-time, sufficient target filters must be configured. Target filters are configured by assigning them source context expressions. Assigned expressions will be assigned to the target filter.

11. Click **OK** to close the Symbols list and **OK** again to close the Expression editor.

    $AGENT:NODE$ appears as the expression for Managed System Name.

12. Click **Next** to open the Parameters window.

13. In the Parameters window, click **Add Symbol** and provide a symbol name for the first attribute whose value you want to pass to the navigation script, then click **OK** to add the symbol to the Parameter panel.

14. Select the symbol you just added, then click **Modify Expression**.

15. In the Expression Editor, click **Symbol** and select the name of the attribute whose value you want to associate with the symbol you just named.

    The list is a reverse hierarchy of available v symbols starting from the source context and the current Navigator item and ending at the root Navigator item.

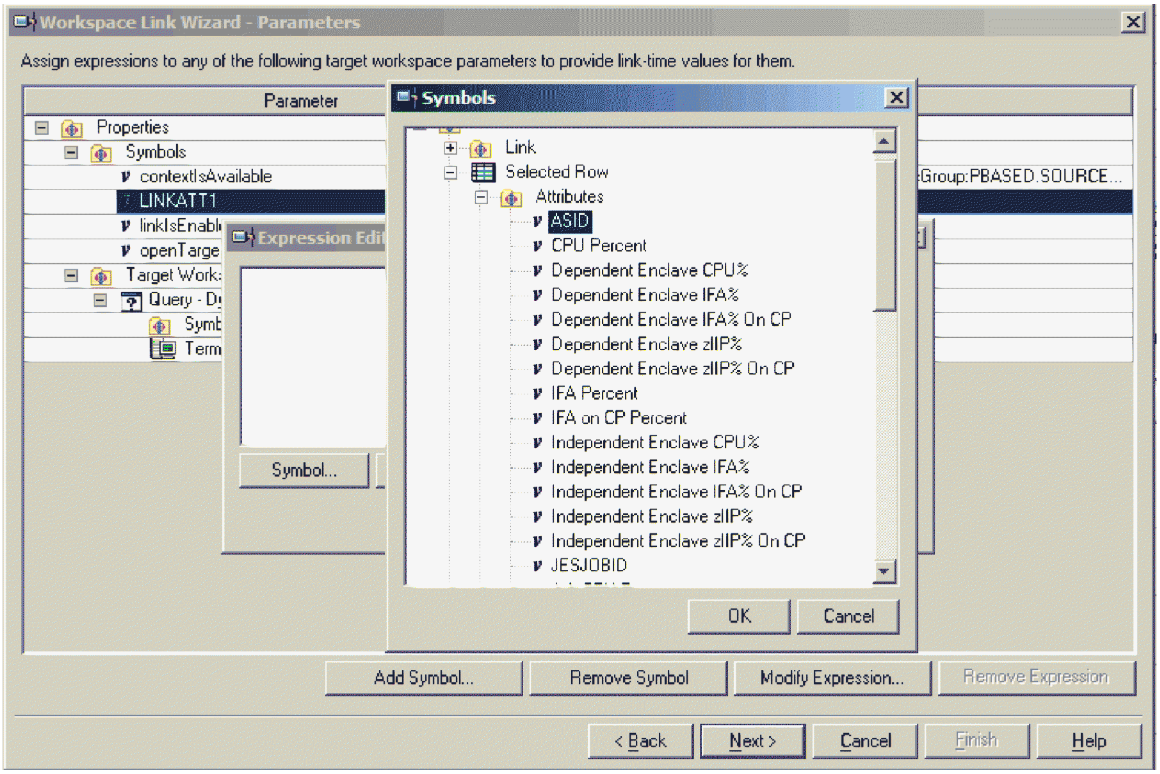

Assign expressions to any of the following target workspace parameters to provide link-time values for them.

16. Click **OK** to close the Symbol window, then click **OK** again to close the Expression editor.

    The attribute name is added in the expression window.

17. Repeat steps 12–15 until you have added all the attributes whose values you want passed to the script, then click **Next**.

    The Summary window is displayed.

18. Review the description of the link that is going to be created. If you want to change something about the link, click **Back**. If the description is correct, click **Finish**.

    The Wizard closes, and the link is added to the view.

Assigning the Dynamic Link to 3270 query

The values returned by the query associated with the terminal view provide the information needed to connect to the correct host, log onto the appropriate OMEGAMON for MVS session. The values passed on the link are used to navigate to a specific OMEGAMON for MVS screen. All terminal views that are the target of OMXE on z/OS dynamic links should use the Dynamic Link to 3270 query.

Note: If you based the target workspace on an existing workspace that is the target of an OMEGAMON XE on z/OS dynamic link, you do not need to assign the query. It is already assigned.

Take the follow steps to assign the Dynamic Link to 3270 query to the terminal view:

1. Right-click in the terminal view and select **Properties**.

   The Properties editor opens and the **Configuration** tab for the view is displayed.

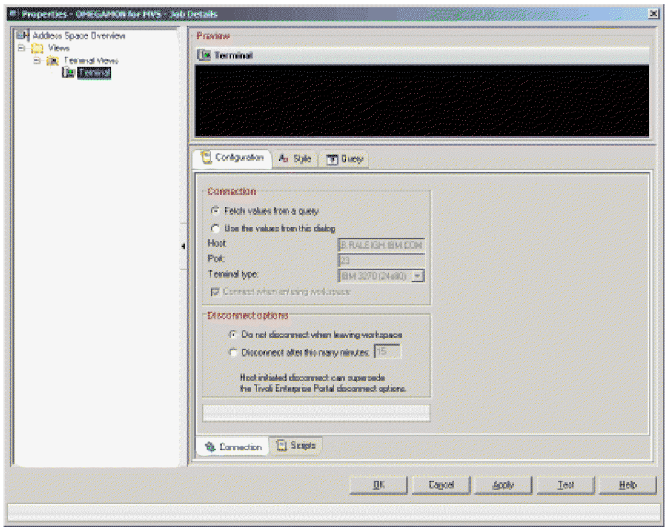

● Fetch values from a query
○ Use the values from this dialog

Host: B.RALEIGH.IBM.COM

Port: 23

Terminal type: IBM 3270 (24x80)

☑ Connect when entering workspace

● Do not disconnect when leaving workspace
○ Disconnect after this many minutes: 15

Host initiated disconnect can supersede the Tivoli Enterprise Portal disconnect options.

2. Select the **Query** tab, then click "Click here to assign a query." The Query editor opens.
3. In the left frame, expand the MVS System entry, and then the DWL to 3270 entry and select Dynamic Link to 3270.

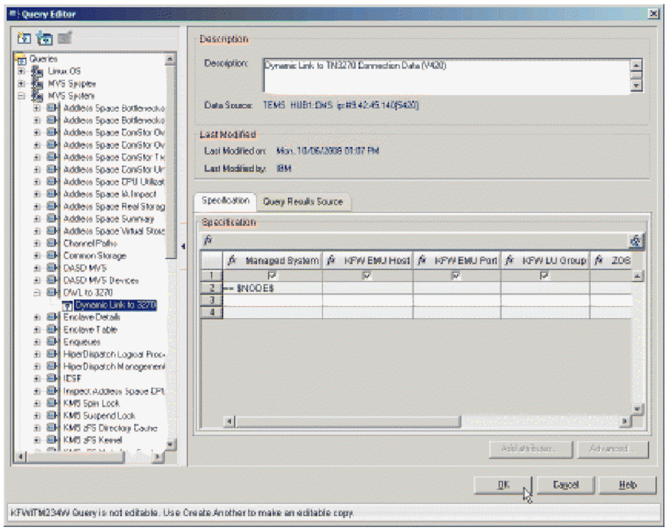

Description: Dynamic Link to TN3270 Connection Data (V420)

Data Source: TEMS HUB1:DNS ip:#9.42.45.140[5420]

Last Modified on: Mon, 10/06/2008 01:07 PM

Last Modified by: IBM

	Managed System	KFW EMU Host	KFW EMU Port	KFW LU Group	ZOS
1	☑	☑	☑	☑	
2	== $NODE$				
3					
4					

KFWITM234W Query is not editable. Use Create Another to make an editable copy.

4. Click **OK** to assign the query and return to the Properties editor.

Part 2. Monitoring scenarios

This section contains scenarios that illustrate how you can use the workspaces, attributes, and situations provided with OMEGAMON XE on z/OS to monitor your z/OS systems and sysplexes.

Chapter 5, “Monitoring shared DASD,” on page 55, provides examples of how you can use situations to monitor the performance of shared DASD in your sysplex. More importantly, this chapter includes instructions for creating situations for filtering the collection of data for DASD devices to reduce processing overhead in environments with large numbers of devices.

Chapter 6, “Monitoring virtual storage and missing jobs,” on page 61, contains two scenarios. The first illustrates how you can use OMEGAMON XE on z/OS to monitor storage usage. The second shows you how to create a monitoring situation to alert you to the failure of critical tasks.

Chapter 7, “Monitoring service class goals,” on page 65, shows you how to monitor service class goals by creating a situation notifies the appropriate parties when a service class missing its goals. It also shows you how to create a threshold in a workspace table view that matches the situation parameters to help pinpoint the problem service classes when you are doing problem analysis.

Chapter 8, “Monitoring cryptographic coprocessors,” on page 71, contains three scenarios that illustrate how you can use how the data collected and presented by OMEGAMON XE on z/OS to monitor and improve your cryptographic services.

Chapter 9, “Detecting CPU looping address spaces,” on page 75 describes the CPU Loop Index metric and how it can be used to detect CPU loops.

Chapter 5. Monitoring shared DASD

Because of the large DASD volume counts that have become common in recent years, monitoring DASD devices without a filter that eliminates some of the devices can lead to high CPU or storage problems and may even cause the monitoring server to fail. Because of these potential costs, although OMEGAMON XE on z/OS provides the capability to monitor sysplex shared DASD, it does not collect DASD device data unless a DASD filter situation is active. An auto-started warning situation (KM5_No_Sysplex_DASD_Filter_Warn) notifies you if no filtering situation is in place and no devices are being monitored.

You can turn DASD data collection on by running a DASD filter situation. OMEGAMON XE on z/OS includes a model filter situation (KM5_Model_Sysplex_DASD_Filter), which uses the DASD Device Collection Filtering attributes Average Response Time and I/O Rate. By customizing this filter to exclude well behaved devices, you can enable monitoring of devices of particular interest and avoid being overwhelmed with unwanted data. A third product-provided situation, KM5_Weak_Plex_DASD_Filter_Warn, alerts you when too many devices are being monitored and filter criteria should be strengthened.

This chapter provides instructions for creating situations for filtering the collection of data for DASD devices to reduce processing overhead in environments with large numbers of devices. This chapter also illustrates how you can use Tivoli Enterprise Portal and OMEGAMON XE on z/OS to monitor and manage the performance of shared DASD in your sysplex.

The section on identifying causes of I/O delays should be of interest to all users. The section on DASD device collection filtering should be of interest to those with administrative authority who are responsible for configuring data collection.

Filtering collection of DASD device data

With OMEGAMON XE on z/OS, the best way to reduce processing overhead is to control the amount of DASD information being sent to the sysplex proxy for sort merge processing. Six thousand unit addresses on each of nine LPARs in a sysplex, for example, requires the proxy to sort merge a considerable amount of data before the data can be evaluated. However, by creating a DASD filter situation to reduce the number of rows of data sent to the proxy and limiting data collection to DASD devices that are performing poorly or experiencing contention, you can dramatically reduce overhead.

In most cases, you create a situation to alert you to problems in the monitored system. When you create a situation for DASD device collection filtering, you are filtering the devices that are being monitored and identifying the devices that need further monitoring. You use the attributes in the DASD Device Collection Filtering attribute group to create the filter situations.

When you create a situation for DASD device collection filtering, OMEGAMON XE on z/OS builds a list of DASD devices based on the situation. This list is rebuilt on the DASD filter situation interval. The lower the interval, the more overhead is incurred as Resource Measurement Facility (RMF) data is collected on all the devices to determine if they qualify. The higher the interval, the more likely that spikes in activity on previously inactive volumes will go unnoticed, as they were not in the monitored volume list.

If a DASD device meets the requirements in the situation, all data for the device is forwarded to the sysplex proxy for sort merge. If the device does not meet the requirements in the situation, the device data is not forwarded to the sysplex proxy. Every monitored LPAR is checked to see if the device meets the filter criteria on that LPAR. If it does meet the criteria, the device is included for monitoring on every LPAR so its activity can be combined over the whole sysplex.

After you have enabled the situations, the number of devices exceeding the situation thresholds should be no more than 500 or whatever number seems viable for your site.

Requirements and restrictions for filter situations

Situations for DASD collection filtering have the following requirements and restrictions.

- You must create the situation using the attributes in the DASD Device Collection Filtering group.
- For each sysplex, you can have only one situation for DASD collection filtering. (For this reason, the one situation must contain all the conditions for monitoring the devices.)
- You can create one situation for DASD collection filtering and distribute the same situation to more than one sysplex.
- The collection interval can be as small as 5 minutes, but should probably be somewhere between 15 and 30 minutes. You may want to use the same interval as you use for Resource Measurement Facility (RMF).

  Overhead for the situation is dependent on both DASD farm size and refresh interval. Big farms or short intervals increase overhead. In establishing the collection interval, weigh the probability of missing an important performance event. If you feel that the usage pattern of your DASD farm changes dramatically every 10 minutes, for example, make the interval 10 minutes.

Creating a filtering situation

OMEGAMON XE on z/OS provides a model filter situation, KM5_Model_Sysplex_DASD_Filter, which uses the DASD Device Collection Filtering attributes Average Response Time and I/O Rate. You can quickly create an effective filter situation by customizing this situation to suit your site requirements.

Complete the following steps to create a filter situation based on KM5_Model_Sysplex_DASD_Filter:

1. In the Tivoli Enterprise Portal Navigator, open the Situation editor.
2. To display the KM5_Model_Sysplex_DASD_Filter situation under MVS Sysplex, click Set Situation Criteria Filter and check **Associated with Monitored Application**. Click **OK**.
3. Select the KM5_Model_Sysplex_DASD_Filter and click the Create Another Situation button.

   The **Create Situation** dialog box is displayed.

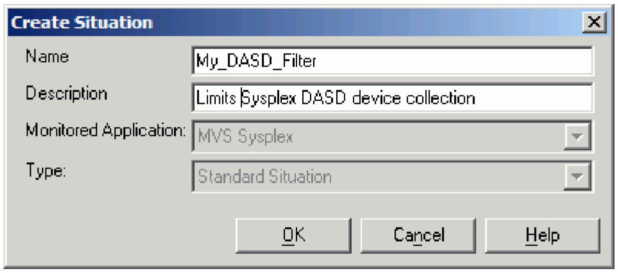

Create Situation

Name: My_DASD_Filter

Description: Limits Sysplex DASD device collection

Monitored Application: MVS Sysplex

Type: Standard Situation

Figure 4. Create Situations dialog

4. Specify a name and description for the situation and click **OK**.

   The **Formula** dialog for the new situation is displayed.
5. Set the threshold values for the Average Response Time and I/O Rate attributes by clicking in the row beneath each attribute name and typing in the values.

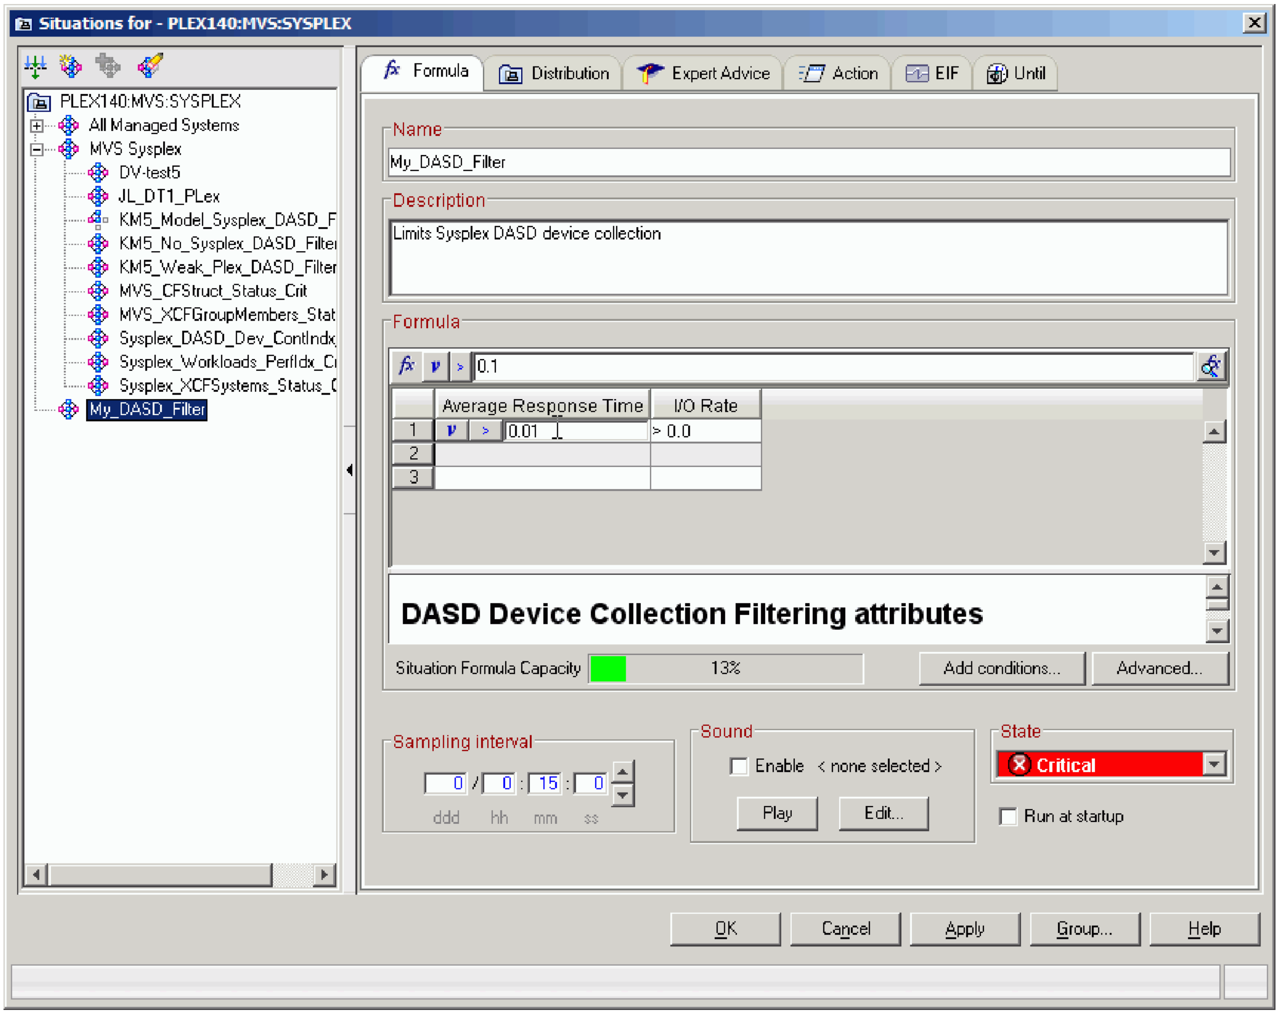

Name: My_DASD_Filter

Description: Limits Sysplex DASD device collection

Formula: > 0.1

	Average Response Time	I/O Rate
1	> 0.01	> 0.0
2		
3		

DASD Device Collection Filtering attributes

Situation Formula Capacity 🟩 13%

Sampling interval: 0 / 0 : 15 : 0 (ddd hh mm ss)

Sound: □ Enable < none selected >

State: Critical

□ Run at startup

Figure 5. The Formula tab for a filter situation

Note: Filter criteria can be specified for other columns, but these two are the only attributes that will have a noticeable effect on refresh time and CPU consumption.

6. Use the other options of this tab to customize the situation: add additional attribute to further refine the filter criteria, change the monitoring interval, and so on.
7. Click Apply to save the new situation.

After you have customized the situation, you must distribute and start it before filtering can take effect.

Distributing the situation

After you have created a situation for DASD device collection filtering, you must distribute the situation to the systems.

Take the following steps to distribute the situation:

1. With the situation selected in the Situation editor, click the **Distribution** tab.
2. Make one of the following assignments:
   - To distribute the situation to a single sysplex, in the **Available Managed Systems** box, click the name of the sysplex to which you want to assign the situation and click the ⇦ left arrow. Tivoli Enterprise Portal moves the name of the sysplex to the **Assigned** box.

   - To distribute the situation to all sysplexes, select *MVS_SYSPLEX in the **Available Managed Systems List** box, and click the ⇦ left arrow to add the managed system list to the Assigned box.
3. Click **OK** to distribute the situation.

Starting and stopping the situation

You can start and stop the situation using the **Situations** dialog box:

1. In the left hand frame of the Situation editor, right-click the name of the situation.
2. Select the appropriate option from the pop-up menu:
   - To start the situation, click **Start**.
   - To stop the situation, click **Stop**.

If you want the situation to run continuously across Tivoli Enterprise Monitoring Server restarts, check **Run at startup** on the **Formula** tab.

Displaying messages for the situations you create

OMEGAMON XE on z/OS provides messages for DASD device collection filtering. You can display these messages in the RKLVLOG on the Tivoli Enterprise Monitoring Server on the host. The messages for sysplex situations have the prefix KOS.

Example situations

The following are examples of situations that filter DASD device collection.

Filtering for devices that are busy:

The following situation limits data collection to devices that have

- Some activity
- Long average response times

```
Average response time GT 70.0 AND
I/O Rate GT 2.0
```

Filtering for performance problems on critical volumes:

The following situation limits the data collection to critical database devices that have

- A volume serial name that begins with *DB2* or contains *IMS*™
- Some activity
- Slow response times

```
(Average response time GT 70.0 AND I/O Rate GT 2.0 AND Volume Serial Number EQ DB2*) OR
(Average response time GT 70.0 AND I/O Rate GT 2.0 AND Volume Serial Number EQ *IMS*)
```

Using shared DASD data collection to identify the cause of I/O delays

OMEGAMON XE on z/OS provides three sysplex level DASD related workspaces:

- The Shared DASD Groups Data for Sysplex workspace displays information on device contention and usage for all the groups in a sysplex. This information can help determine how equitably a device is serving all systems in the sysplex.
- The Shared DASD Devices workspace displays statistics for the individual devices in a selected group. The Shared DASD Devices workspace displays information about the activity of the shared devices for a group, averaged over all systems in the sysplex. This information can help determine how equitably a device is serving all systems in the sysplex.

- The Shared DASD Systems workspace displays information about the systems that share a device. This information helps you measure the performance and exceptions from the perspective of each system.

This scenario illustrates how you can use these workspaces, in conjunction with a monitoring situation designed to alert you to device contention, to identify the devices and data sets responsible for significant I/O time or delays for important workloads in your sysplexes.

Getting situation event notification

You are looking at the Tivoli Enterprise Portal and see that the **z/OS Systems** Navigator icon is overlaid by a warning event indicator.

You move your cursor over the icon to see a flyover list of situations that are currently true for your mainframe systems. The Sysplex_DASD_Dev_ContIdx_Warn situation, which you activated to monitor DASD device contention in your sysplexes, is listed in the flyover.

Glancing down the expanded **z/OS Systems** tree, you notice warning indicators for the **Service Classes Data for Sysplex** and the **Shared DASD Group Data for Sysplex** items in the **SYSPLEX1** tree, which leads you to suspect the problem is in the distribution of I/O activity among your DASD volumes.

Analyzing I/O distribution

The Sysplex_DASD_Dev_ContIdx_Warn situation has alerted you to the fact that DASD device contention index has reached a level that warrants attention. Now you want to find information about the device or devices causing the contention.

In the Navigator, you select **Shared DASD Groups Data for Sysplex** and the default workspace is displayed. To determine if I/O activity is unevenly distributed among the devices, you examine:

- The average true busy percentage of the group. If, for example, the value ranges between 1.1 at the lower end and 60.2 at its highest, that means that on average, devices in the group spend 1.1% of their time doing work for all systems; at its busiest, one device is spending 60.2% of its time doing work for all systems.
- The average device contention for the group. For example, the value of the average contention index might be 0.75, and the highest device contention index of the group 1.5. A high number (more than 1.0) means that I/O requests for a device are substantially delayed because of contention for the device or path generated by other systems. You would like to lower the contention index for a device in contention.

Examining this information allows you to determine which group or groups of devices is experiencing uneven I/O distribution, but you need information about the specific devices involved.

Isolating the problem

To get information about the devices in a group you have identified, you click the 🔗 icon beside its name in the Shared DASD Groups table view to link to the Shared DASD Devices workspace for that group.

The Shared DASD Devices workspace displays statistics for each device in the group. With the information in this workspace, you can determine if a device is not serving all systems in the sysplex equitably. You determine this by examining:

- True percent busy. If this value is too high for one device, while other devices have a lower value, the workload should be balanced between the devices.
- Contention index. If this value is too high, this is an indication that I/Os for the device are substantially delayed due to contention for the device or path generated by other systems.

- System response time. A high value in this column indicates an inordinate amount of time is required for this device to process I/O activity.
- Cumulated I/O rate. This indicates the device is not processing I/Os at an acceptable rate.

To view additional information for a specific device to help you determine if it is causing I/0 delays, you click the 🔗 link button beside the volume serial number in the **Shared DASD Devices** table. The link takes you to the Shared DASD Systems workspace, which shows:

- The systems that share the device and indicates which systems have a high response time and high I/O rate.
- Performance measures and exceptions presented for each system and indicates how the device is performing for each system.

Based on this information, you can identify which devices are causing excessive I/O delays.

Taking action to resolve a shared DASD problem

From the information you have gathered, you decide to implement one of the following solutions.

- Redistribute the work among devices to eliminate the contention for resources.
- Reschedule the work for a single device to a time when device contention is usually low.

Chapter 6. Monitoring virtual storage and missing jobs

This chapter contains two scenarios. “Monitoring paging and virtual storage” illustrates how you can use OMEGAMON XE on z/OS to monitor storage usage; “Monitoring critical started tasks” on page 62 illustrates how to create a monitoring situation to alert you when critical tasks fail.

Monitoring paging and virtual storage

The maximum amount of system virtual storage is bounded by the amount of real storage plus paging space limits, so paging space performance and availability become a vital factor in the effective execution of applications. With the introduction of 64-bit or large storage objects, the memory requirements of paging became higher and storage can become exhausted more readily. When usage approaches 30%, paging efficiency begins to decline, and blocked paging disappears at about 35% occupancy. Severe problems can occur if page space used is greater than 85%. Because the percentage of paging space increases as well as paging rate, it is a good indictor that a problem may be ensuing.

The following scenario suggests how you can use the resources provided by OMEGAMON XE on z/OS to detect and analyze paging and storage problems.

Activating the OS390_Local_PageDS_PctFull_Crit situation

To alert you to impending problems and to simplify analysis when problems occur, you can activate the predefined situation OS390_Local_PageDS_PctFull_Crit. This situation monitors to determine whether the percentage of slots in use on a local page dataset is greater than or equal to 35% and issues a Critical alert if this condition is found to be true. (See “Activating situations” on page 22 for instructions for activating the situation.)

You can also modify this situation to include page rate as an additional indicator that the paging system performance is becoming impacted.

Using the Page Dataset Activity workspace to gather information

When you see an event indicator alerting you that the OS390_Local_PageDS_PctFull_Crit condition is true, you might start investigating the problem using the Page Dataset Activity workspace for the affected system.

This workspace provides information about availability and response time for a specific page dataset. Page data sets are auxiliary storage datasets that back up all frames of virtual storage. They must be large enough to contain all common and private virtual storage. Page data sets are used when an address space references data that is not in either real or expanded storage. The process of bringing in data is called a page-in and is coordinated by the auxiliary storage manager (ASM). If swap data sets are not defined, page data sets also contain the swapped out part of an address space.

Because the process of paging is very slow when compared to referencing data from real or expanded storage, it is important that page dataset devices be isolated from contention with other kinds of work. This is especially true if there is contention for real and expanded storage, and the page fault rate is high. The Percent Full and Response Time bar charts in this workspace provide visual representations of the availability of space in the various types of page datasets and the response times for those data sets.

Even if page rates are low, data sets over 35% percent full can indicate a performance issue is developing and some action may be required.

Your next step is to decide what action to take to resolve any problems. Rather than simply adding another dataset, you can use the Address Space Storage workspace to evaluate if there are any jobs which can be trimmed or moved to a different system or time slot to balance out system resources.

Examining storage usage

You can link to the Address Space Storage workspace from the Address Space Counts table of the Address Space Overview workspace.

In the workspace, the **Fixed Storage** bar graph shows who the heavy users of fixed storage are. The Virtual Memory table at the bottom of the workspace provides data on fixed and virtual low, extended, and large storage use. This information can be used as an application tuning tool as well as system performance tuning tool.

Check for applications using large percentages of the storage memory limits. This information can be used in deciding how to manage those applications later.

Evaluating your options

Consider the following options to address any paging problems you have detected:

- Increase paging to align with potential virtual storage demand

  This would mean increasing the size of paging space by adding more local page data sets or increasing the size of existing ones.
- Redistribute larger applications to other LPARs

  You can redistribute large applications or applications using a lot of storage to other systems where the workload is lighter or paging demand lower. Alternately or in addition, you could rebalance large jobs at different times so that they do not run concurrently and compete for virtual storage or paging space.
- Decrease the applications memlimit to reduce storage demand on system

  The viability of this option depends on your application service requirements, as such a change could directly impact performance in the application. This could also depend on the behavioral effects on the application, as some applications may not be able to effectively function with lowered memory limits.

Monitoring critical started tasks

In most environments, there is a set of started tasks, such as CICS®® tasks or WebSphere® tasks, that should always be running. This scenario shows you how to define a situation that will alert you if one or more of these task fails. The situation is based on the Job Name attribute of the Address Space CPU Utilization attribute group and uses the **Check for Missing Items** function.

Creating the MissingTaskAlert situation

Note: This scenario assumes that you are already familiar with the basic steps for creating a situation. If you are not sure of the steps, see "Creating the zOS_Critical_SvcClass_Missed_Goal situation" on page 65.

For this scenario, you are creating a situation to monitor two tasks, CICS1 and WEB5, running on SYSTEM1 in SYSPLEX1.

1. Navigate to the Address Space Overview workspace for SYSTEM1.

   Enterprise
   - ⊟ z/OS Systems
     - ⊟ SYSPLEX1:MVS:SYSPLEX
       - Coupling Facility Policy Data for Sysplex
       - Coupling Facility Structures Data for Sysplex
       - ...
       - ⊟ SYSTEM1
         - ⊟ MVS Operating System

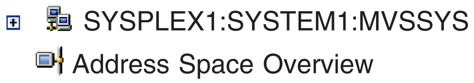

⊞ SYSPLEX1:SYSTEM1:MVSSYS

Address Space Overview

2. With the workspace displayed, access the Situation editor by right-clicking **Address Space Overview** in the Navigator and selecting Situations from the pop-up menu.
3. In the Situation editor, create a new situation.
4. In the Create Situation dialog box, type a name and description for the situation, for example, `MissingTaskAlert`.

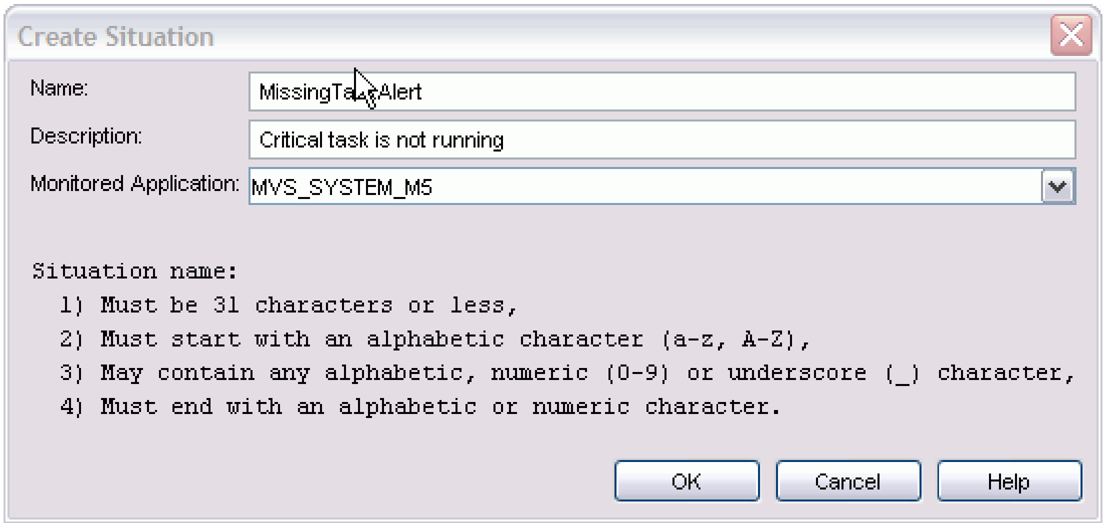

Name: MissingTaskAlert

Description: Critical task is not running

Monitored Application: MVS_SYSTEM_M5

```
Situation name:
  1) Must be 31 characters or less,
  2) Must start with an alphabetic character (a-z, A-Z),
  3) May contain any alphabetic, numeric (0-9) or underscore (_) character,
  4) Must end with an alphabetic or numeric character.
```

5. In the Select Attribute dialog box, select the `Address Space CPU Utilization` attribute group and the `Job Name` attribute.

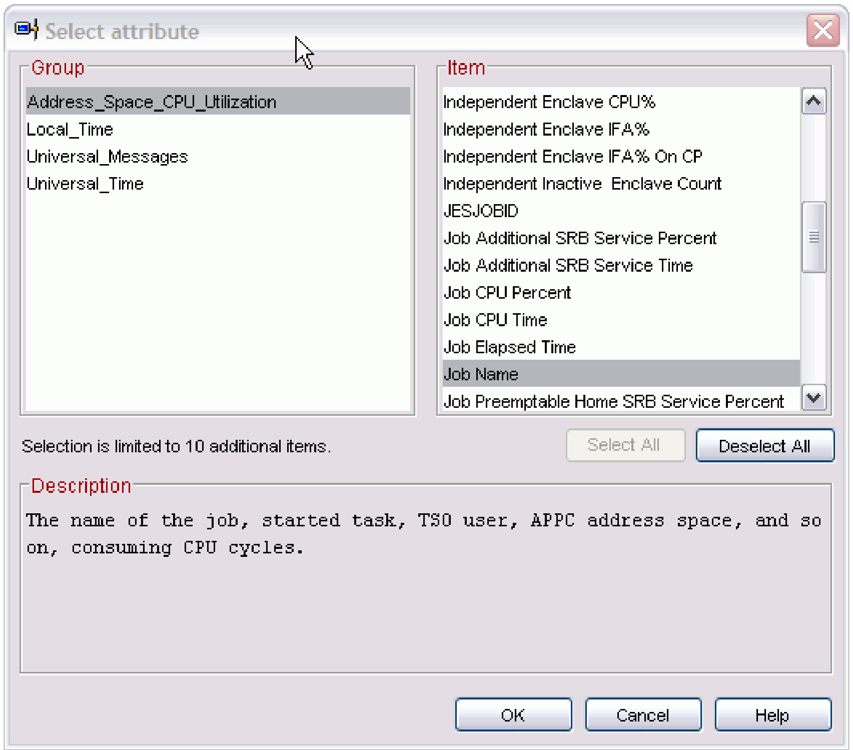

Group

- Address_Space_CPU_Utilization
- Local_Time
- Universal_Messages
- Universal_Time

Item

- Independent Enclave CPU%
- Independent Enclave IFA%
- Independent Enclave IFA% On CP
- Independent Inactive Enclave Count
- JESJOBID
- Job Additional SRB Service Percent
- Job Additional SRB Service Time
- Job CPU Percent
- Job CPU Time
- Job Elapsed Time
- Job Name
- Job Preemptable Home SRB Service Percent

Selection is limited to 10 additional items.

Description

The name of the job, started task, TSO user, APPC address space, and so on, consuming CPU cycles.

6. In the expression editor, select the **Check for Missing Items** function:

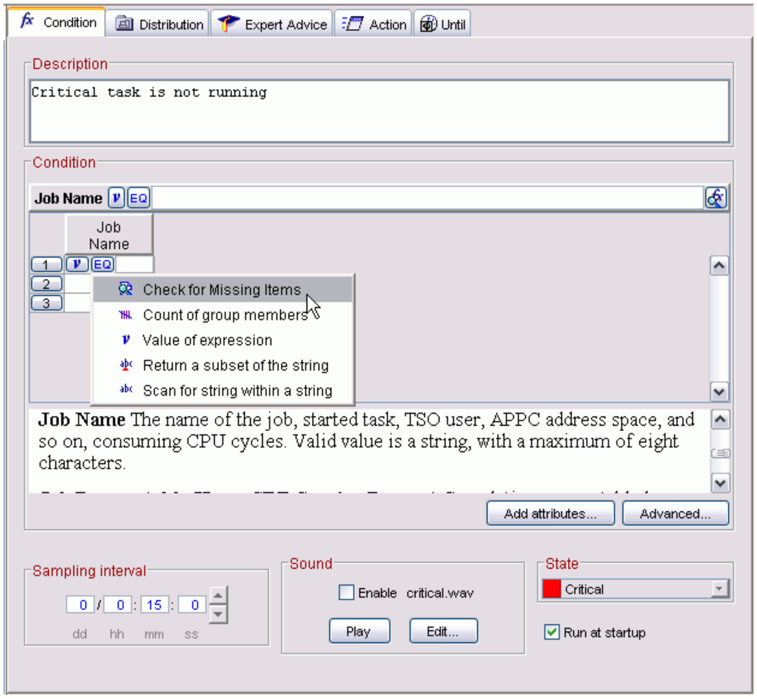

Description: Critical task is not running

Job Name The name of the job, started task, TSO user, APPC address space, and so on, consuming CPU cycles. Valid value is a string, with a maximum of eight characters.

Sampling interval: 0 / 0 : 15 : 0

Sound: □ Enable critical.wav

State: 🟥 Critical

☑ Run at startup

7. In the Missing Item popup, enter your list of critical tasks, and then click OK.
8. If you want to make the situation trigger independently for each job name:
   a. Click the Advanced . . . button.

      The Advanced Situation Options dialog box appears.
   b. Select the **Display Item** tab, and then select `Job Name` as the display item.

      This is especially helpful if you want to attach an action to the situation, such as a start command.
   c. Click OK to close the Advanced Situation Options dialog.
9. Complete the situation, for example by selecting a different sampling interval, adding any advice or instructions you want to provide, adding a Start command to restart the task, or distributing it to other systems you want it to run on.

   Remember to stop the situation if you migrate the task it is monitoring.

Chapter 7. Monitoring service class goals

This chapter presents a scenario which illustrates how you can use OMEGAMON XE on z/OS resources to monitor service class goals. This scenario also provides step-by-step instructions for creating a situation, defining a reflex action using a TSO command, and setting a threshold in a table view.

Setting the scene

Your DB2 and IMS transaction servers service applications that are important to your business operation. You run these address spaces on your SYSTEM1 z/OS system under the STCONLN Service Class.

You want to use OMEGAMON XE on z/OS to monitor this service class and notify the appropriate parties when the service class is missing its goals. You also want to use OMEGAMON XE on z/OS to determine why the service class is missing its goals.

To set up monitoring, alerting and analysis, you perform the following tasks using the Tivoli Enterprise Portal:

- Create a situation that will raise a critical alert when a service class is not meeting its goals.
- Define an action for the situation that will cause a system command to be executed at the z/OS system associated with the raised situation
- Modify the WLM Service Class Resources include to provide a column threshold that matches the situation parameters. This will help pinpoint the problem service classes when doing problem analysis.

Creating the zOS_Critical_SvcClass_Missed_Goal situation

Your first step is to create a situation that raises an event and notifies the appropriate personnel when a service class misses its goal.

Using the Situation editor, you create a new situation named zOS_Critical_SvcClass_Missed_Goal. This situation raises a ⊗ Critical event indicator and sends a TSO message to a designated system administrator when the Performance Index for a service class period is above 1.5. You set the situation to start whenever the Tivoli Enterprise Monitoring Server starts.

Because you access the Situation editor from the **WLM Service Class Resources** Navigator item for SYSTEM1, this situation is automatically associated with that item and a critical event indicator will appear on the item when the situation is true.

To create the zOS_Critical_SvcClass_Missed_Goal situation:

1. Navigate to system SYSTEM1 in the Navigator and ⊞ expand the item, if necessary.
2. Right-click the **WLM Service Class Resources** item and select **Situation**s from the pop-up menu.

   The Situation editor opens. The WLM Service Class Resources node and any associated situations are displayed in the left-hand frame.
3. Select the Create New Situation icon.

   The Create Situation dialog box appears.
4. Provide a name and description for the situation, and then click OK:

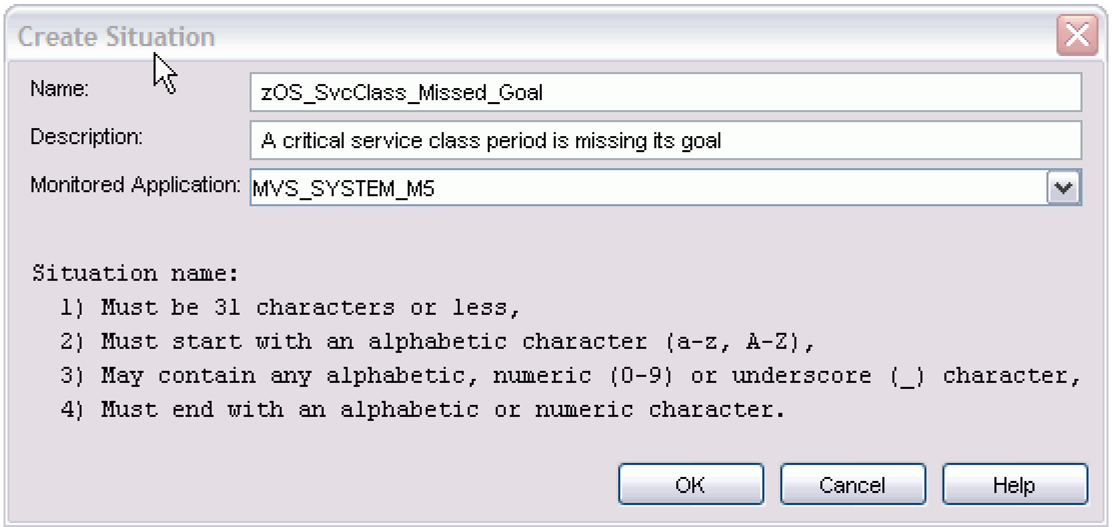

Create Situation

Name: zOS_SvcClass_Missed_Goal

Description: A critical service class period is missing its goal

Monitored Application: MVS_SYSTEM_M5

```
Situation name:
  1) Must be 31 characters or less,
  2) Must start with an alphabetic character (a-z, A-Z),
  3) May contain any alphabetic, numeric (0-9) or underscore (_) character,
  4) Must end with an alphabetic or numeric character.
```

The Select attribute dialog box appears, with the WLM_Service_Class_Resources group selected.

5. Scroll down the item list and select the `Performance Index` attribute, and then click OK.

   The **Condition** tab for your new situation is displayed with the Performance Index attribute added to the expression editor.

6. Create the situation expression:

   a. Click in the first row of the **Performance Index** column.

   b. Click the EQ relational operator button and select **Greater Than (GT)** from the pop-up menu.

   c. Type enter a value of 1.5.

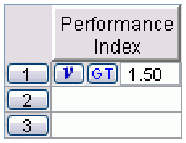

	Performance Index
1	GT 1.50
2	
3	

7. Accept the default sampling interval of every 15 minutes, or set the sampling interval that suits your monitoring requirements.

8. By default, **Run at startup** is selected. This means that after the situation has been distributed, it will start whenever the Tivoli Enterprise Monitoring Server is started. If you want to start and stop the situation manually, for example if you want to test the effects of the selected sampling interval, deselect Run at startup.

9. Click Apply to save the situation properties you have defined so far.

10. To generate a separate action message for each service class name, make Service Class a display item, using the Advanced. . . button:

    a. In the **Condition** tab, click the Advanced... button.

       The Advanced Situation Options dialog box appears.

    b. Select the **Display Item** tab, and then use the drop-down menu for the **Item** field and select `Service_Class`.

    c. Click OK to close the Advanced Situation Options dialog box.

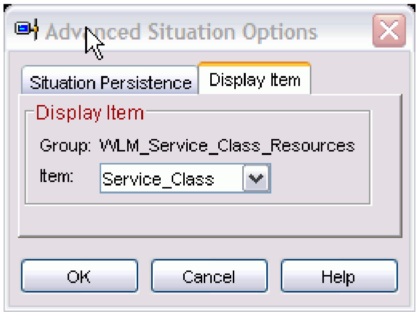

Display Item

Group: WLM_Service_Class_Resources

Item: Service_Class

Defining the Take Action command

You want to use a TSO command to send a message to the system administrator whenever a service class misses its goal. You want to include in the message the name of the service class and the value of the Performance Index attribute at the time the situation became true.

To define the Take Action command:

1. With the zOS_Critical_SvcClass_Missed_Goal Properties window open, select the **Action** tab.
2. In the `System Command` text field, type the following command, using attribute substitution for *service_class_name* and *performance_index* value.

   ```
   SEND 'service class service_class_name is missing goal. Performance Index is
   performance_index',USER=(userid)
   ```
3. If you want separate notification for every monitored item for which the situation is true, click **Take action** on each item.
4. Set the command to run at the Tivoli Enterprise Monitoring Server.

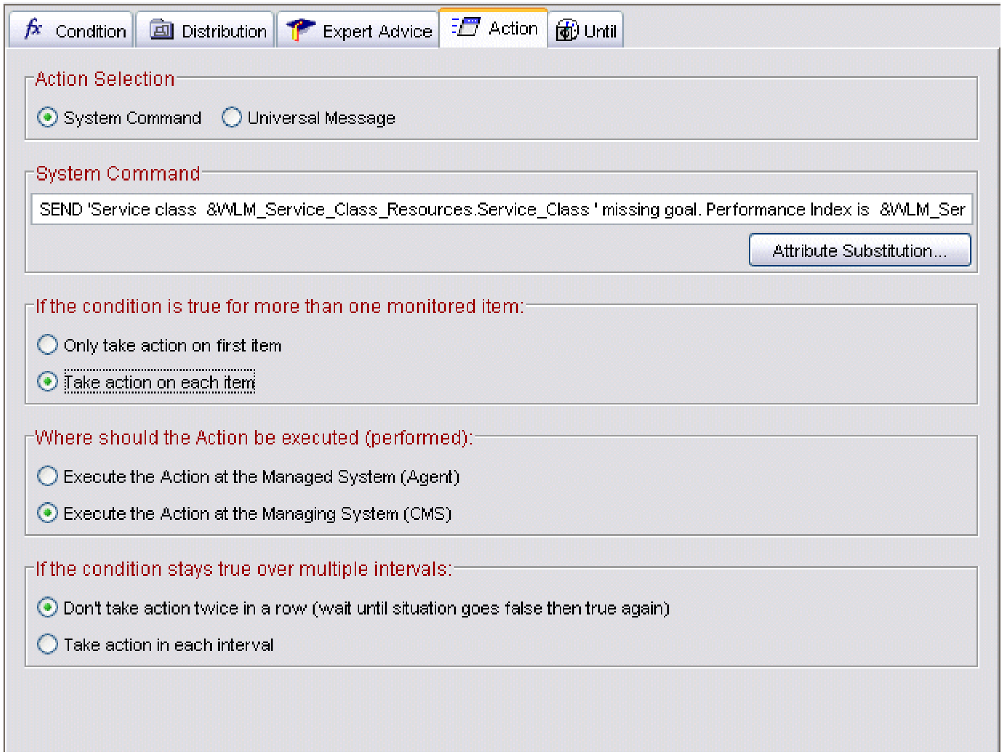

Action Selection

● System Command ○ Universal Message

System Command

SEND 'Service class &WLM_Service_Class_Resources.Service_Class ' missing goal. Performance Index is &WLM_Ser

If the condition is true for more than one monitored item:

○ Only take action on first item
● Take action on each item

Where should the Action be executed (performed):

○ Execute the Action at the Managed System (Agent)
● Execute the Action at the Managing System (CMS)

If the condition stays true over multiple intervals:

● Don't take action twice in a row (wait until situation goes false then true again)
○ Take action in each interval

5. Click Apply to save the command.

Setting thresholds in the WLM Service Class Resources workspace

Now you want to set thresholds in the WLM Service Class Resources workspace that will mirror the condition you defined in the situation. This will allow you to analyze problems more easily. You decide that you also want to create a threshold that will produce a warning indicator when a service class is nearing the critical threshold.

You use the Properties editor for the **WLM Service Class Resources** table to add these thresholds to the **Performance Index** column.

Setting the thresholds

To set the thresholds:

1. Select the **WLM Service Class Resources** item for SYSTEM1 in the Navigator.

   The default workspace is displayed.
2. In the workspace, right-click in the **WLM Service Class Resources** table view and select **Properties** from the pop-up menu.
3. In the Properties editor, select the **Thresholds** tab.
4. To create the critical threshold:
   a. Scroll in the Thresholds editor until you see the **Performance Index** attribute.
   b. Click in the first row beneath the column heading. The critical (red) indicator is already selected.
   c. Select a relational operator of **Greater Than (GT)** and type a value of `1.5`.

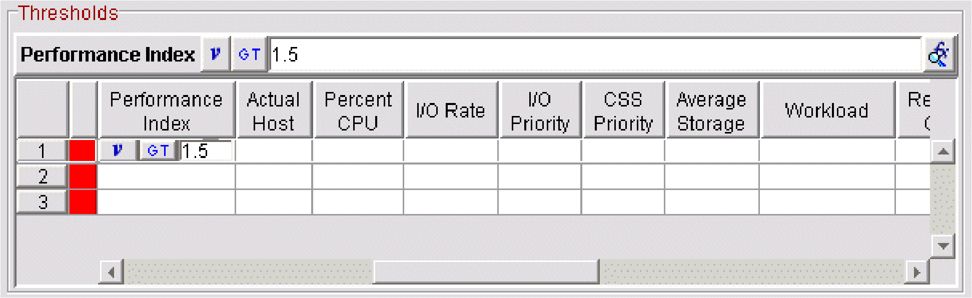

Thresholds

Performance Index GT 1.5

		Performance Index	Actual Host	Percent CPU	I/O Rate	I/O Priority	CSS Priority	Average Storage	Workload	Re (
1	🟥	GT 1.5								
2	🟥									
3	🟥									

5. To create a warning threshold:
   a. In the second row click the alert indicator selector next to the row number and select the warning (yellow) indicator.
   b. Click in the second row below the **Performance Index** attribute and select the Greater Than (GT) relational operator and type a value of `1.0`.

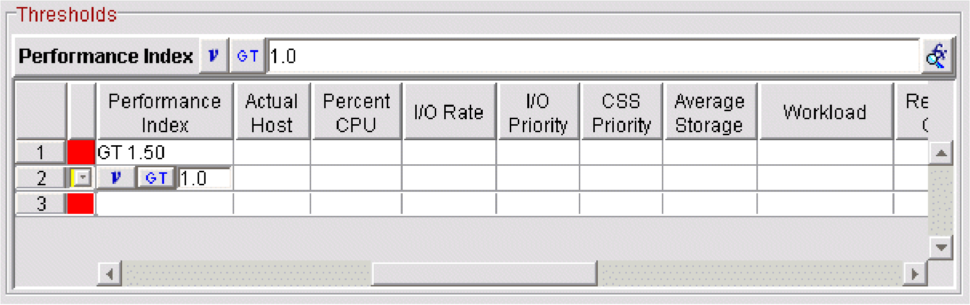

Thresholds

Performance Index GT 1.0

		Performance Index	Actual Host	Percent CPU	I/O Rate	I/O Priority	CSS Priority	Average Storage	Workload	Re (
1	🟥	GT 1.50								
2	🟨	GT 1.0								
3	🟥									

   c. Click OK to save the thresholds and close the editor.

Analyzing the problem

When the situation evaluates to true, you receive notification that a service class is not meeting its goal. In Tivoli Enterprise Portal, you move your mouse pointer over the event indicator on the Enterprise icon in the Navigator to see the Event flyover. In the flyover, you right-click the zOS_Critical_SvcClass_Missed_Goal situation and select Acknowledge from the pop-up menu to create an acknowledgment to let the operators monitoring the situation in the data center know you are working on the situation.

In the flyover, you click the 🔗 link icon next to the service class situation to open its event workspace. You compare the initial situation values and the current situation values to see if the high performance index value is persisting. Since it has taken you several minutes to respond to the situation notification, and the value is staying high, you decide to analyze the problem further to see if you can prevent this problem from arising.

Using the WLM Service Class Resources workspace

You navigate to the WLM Service Class Resources workspace for the affected system to examine the performance information. The thresholds you previously set in the Performance Index column help you pinpoint the problem service class periods, in this scenario the STCONLN service class. You note the goal importance for the service class in case an adjustment is required to address the problem. In addition, you examine the overall performance of service classes with less important goals to determine if they can be adjusted to allow more resources to be available to the more important service classes.

Using the Address Spaces Workspace for Service Class workspace

To get more information about the problem, you link to the Address Spaces Workspace for Service Class from the sysplex level Service Classes Data for Sysplex workspace. In this workspace you can examine data for all the address spaces in the STCONLN service class. For service classes like STCONLN with a goal type of Velocio, CPU will be the primary resource required for meeting this goal. You note the address spaces with the highest CPU utilization, as these may require further investigation. Sorting the table by the CPU Percent column helps you identify the address spaces with highest CPU usage. You also examine the address space list to verify that there are no unexpected address spaces.

Using the Service Class Workflow Analysis workspace

Stepping back ⇦ to the Service Classes Data for Sysplex workspace, you link to the Workflow Analysis Workspace for Service Class to determine the greatest resource impactor for the STCONLN service class. In this case, you determine the major bottleneck for the service class or its associated address spaces is Waiting on CPU, so you want to examine the performance information for the LPAR and central processing complex where the workload is running.

Using the LPAR Clusters workspace

You use the LPAR Clusters workspace to examine the performance information for the LPAR and CPC where the workload is running. You examine the performance parameters listed below to determine a path for problem resolution:

- *CPU % Index:* A value of 1 or greater indicates that the actual LPAR physical processor utilization meets or exceeds the configured targets. A value less than 1 indicates that the LPAR is not able to obtain the resources it is targeted to obtain (based on its defined weight).
- *Effective Weight Index*: A value of 1 or greater indicates that the ability of the LPAR to obtain logical processor resource meets or exceeds the defined targets. A value less than 1 indicates that the LPAR is not able to obtain the resources it is targeted to demand (based on its defined weight).
- Current, Initial, Minimum and Maximum Weights for the LPAR: In a shared physical processor configuration, the LPAR WEIGHT determines the relative importance of the LPAR for the allocation of processor resources. In an IRD configuration, the weights will be adjusted within the maximum and minimum bounds.

- *CPU % Ready*: Indicates the percent of time that the LPAR had “ready” work and was not dispatched (for example, because no processors are available).
- *LPAR Capping Status*: Indicates if “capping” is defined for the LPAR. LPAR will prevent an LPAR from obtaining processor resource even when other LPARs are not using available resources.

In a sysplex configuration, the performance of a given service class should be examined in all the LPARs where the service class workload is running. This will help balance any performance adjustments that may be implemented to resolve the problem.

Chapter 8. Monitoring cryptographic coprocessors

This chapter contains three scenarios that illustrate how you can use how the data collected and presented by OMEGAMON XE on z/OS to monitor and improve your cryptographic services:

- "Validating your cryptography configuration"
- "Monitoring and improving cryptography performance" on page 72
- "Monitoring and improving cross-system ICSF performance" on page 73

Validating your cryptography configuration

Many cryptography problems are the result of configuration errors such as failure to assign coprocessors to the z/OS system, offline or unavailable coprocessors, disabled public keys, or invalid master keys. This scenario illustrates how you can use OMEGAMON XE on z/OS to check your cryptography configuration and correct any errors you discover.

To check the configuration of your coprocessors:

1. In the Navigator, expand the item for a system in a sysplex and scan the tree for the Cryptographic Coprocessor entry:

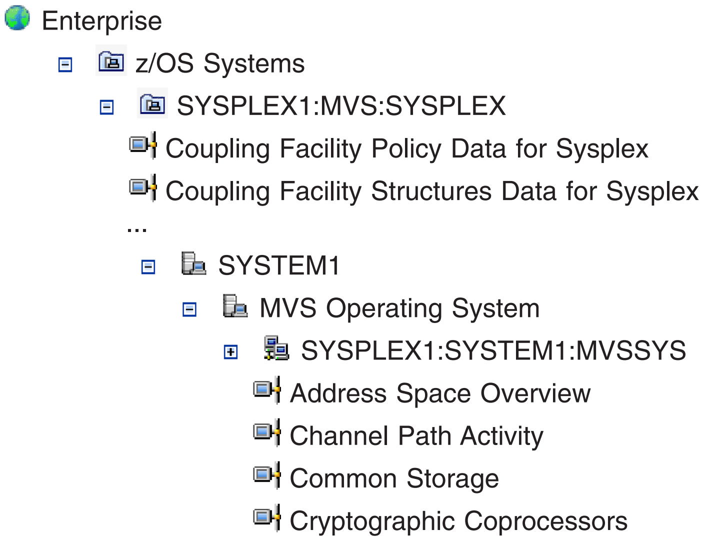

```
Enterprise
    ⊟ z/OS Systems
        ⊟ SYSPLEX1:MVS:SYSPLEX
            Coupling Facility Policy Data for Sysplex
            Coupling Facility Structures Data for Sysplex
            ...
            ⊟ SYSTEM1
                ⊟ MVS Operating System
                    ⊞ SYSPLEX1:SYSTEM1:MVSSYS
                    Address Space Overview
                    Channel Path Activity
                    Common Storage
                    Cryptographic Coprocessors
```

2. Select (click) **Cryptographic Coprocessors** to display the default Cryptographic Services workspace.
3. Check the data in the **ICSF Subsystem Status** view to make sure that
   - The ICSF subsystem is configured correctly.
   - Master keys are loaded and set correctly.
   - Coprocessors are online and active.
   - Cryptography services are operational.
4. If several ICSF subsystems are installed on images that share coprocessors, and a monitoring agent is installed on each subsystem, inspect the values for each subsystem. Also, be sure to check cross-system ICSF performance (see "Monitoring and improving cross-system ICSF performance" on page 73).

   Recheck the Cryptographic Coprocessor workspaces after any changes or adjustments to the cryptography configuration.

Monitoring and improving cryptography performance

The cryptographic coprocessor data collected by OMEGAMON XE on z/OS helps you make load-balancing decisions to improve cryptography performance.

Cryptography performance monitoring on each Integrated Cryptographic Service Facility (ICSF) subsystem has two main components.

- Service call performance monitoring, which involves gathering data such as arrival rate of service requests, time to complete each service call, and queue lengths.
- Top user performance monitoring, which involves determining which job names are the heaviest users of cryptography services.

Checking service call performance

Use the Service Call Performance workspace to evaluate how well service requests are being handled. You can link to this workspace from the **Service Call Performance** table of the Cryptographic Services workspace for a particular system (see steps 1 and 2 of “Validating your cryptography configuration” on page 71), or from the **Service Call Performance by System** table of the Cross-System Cryptographic Coprocessors Overview workspace. Starting from the Cross-System workspace allows you to gain an overview of performance and to quickly check performance details for several systems in succession.

1. Select (click) the **z/OS Systems** item in the Navigator.

   The Sysplex Enterprise Overview workspace is displayed.
2. Right-click the **z/OS Systems** Navigator item and select **Workspace > Cross-System Cryptographic Coprocessors Overview** from the pop-up menu.

   The Cross-System Cryptographic Coprocessors Overview workspace is displayed.
3. In the **Service Call Performance by System** table view of the workspace, click the 🔗 link icon.
4. In the **Select Target** dialog box, select the Cryptographic Coprocessors node for the system for which you want data.

   The Service Call Performance workspace is displayed.
5. In the Service Call Performance workspace:
   - Check the Average Arrivals per Minute bar chart to see which services are being called most frequently.
   - Check the Average Service Time per Call bar chart to see which services are taking the longest time to complete.

     Service calls that are taking the longest time to complete, but are not arriving frequently, probably do not pose a performance problem. Performance problems tend to occur when a particular service call both arrives frequently and takes a long time to complete.
   - Check the Average Pending per Call bar chart for queue length.

     A high number of pending requests for a particular service call would indicate a performance problem.
   - Use the Average Bytes per Service Call bar chart to see whether the service calls with the largest number of bytes also have the highest service times.

     Relatively high services are to be expected for calls that have the largest number of bytes.

     Note: Byte counts are available for some but not all service calls. Therefore, the data shown in the Average Bytes per Service Call bar chart are correct but incomplete. For a list of the service calls for which byte counts are available, see the online help.
6. If any of the statistics displayed in the bar charts do not seem to make sense or if you need more information about a service call’s performance, examine detailed data in the Service Call Performance table. For explanations of all the attributes in this table, see the online help.

7. To check service call performance on another system, use the Backward button (on the Tivoli Enterprise Portal toolbar in desktop mode) or the Back button of your browser to return to the Cross-System Cryptographic Coprocessors Overview workspace, and then repeat steps 3–6.

Checking top user performance

Use the Top User Performance workspace to learn which jobs are the heaviest users of cryptography services. You can link to this workspace from the Cryptographic Services workspace for a particular system (see steps 1 and 2 of “Validating your cryptography configuration” on page 71) or from the **Top Users by System** table in the Cross-System Cryptographic Coprocessors Overview workspace (see steps 1 and 2 of “Checking service call performance” on page 72).

Check the bar charts in the Top User workspace to see which jobs

- Are requesting cryptography services most frequently.
- Have the highest average service time.
- Are waiting longest for their requests to move to the top of the queue.
- Are requesting services with the highest byte counts.

Note: Byte counts are not available for all service calls. Therefore, the data shown in the **Top 10 Average Bytes per Call** bar chart are correct but incomplete. For a list of the service calls for which byte counts are available, see the online help.

If you need more information about top user performance, you can find detailed data in the **Performance by Top Users** table. One particularly useful piece of information in this table is the **LastSvcDesc** column, which shows the service requested most recently by each of the top users. Click the refresh button several times and see whether the same service call keeps showing up. If so, that service call is being used heavily by the top users and may be implicated in any performance problems that arise.

For explanations of all the columns, see the online help, or the *Tivoli OMEGAMON XE on z/OS Interface Reference Guide*.

Improving performance

The following are suggestions for improving performance.

- If service times are unacceptably high, you might consider decreasing the strength of cryptography by reducing the length of the key. Conversely, if you need to increase the strength of cryptography, you can observe and weigh the performance risk.
- If any of the top job names are relatively unimportant, you might want to reduce their priority on the system.
- For the most frequently called services and the services with the highest normal service times, you might want to create your own situations. In each situation, specify combinations of arrival rate and service time that you consider worrisome (warning) or truly unacceptable (critical). Whenever a service call reaches a specified threshold, an event alert will be posted on Tivoli Enterprise Portal. You can then take immediate action to correct the problem.

For instructions on creating situations, see the Tivoli Enterprise Portal online help.

Monitoring and improving cross-system ICSF performance

If several ICSF subsystems are installed on z/OS images that share coprocessors in a sysplex or Processor Resource/System Manager (PR/SM) complex, the workloads on each subsystem can affect the performance of the other subsystems. Use the Cross-System Cryptographic Coprocessor Overview workspace to compare the subsystems’ cryptography performance and to troubleshoot performance problems.

Checking and improving cross-system cryptography performance

Access the Cross-System Cryptographic Coprocessor Overview workspace as described in steps 1 and 2 of “Checking service call performance” on page 72.

To check cross-system performance:

- Check the **ICSF Subsystems by System** table to make sure the coprocessors are online and cryptography services are active on all systems.

  The predefined situation Crypto_No_PCI_Coprocessors defines the lack of an online PCI coprocessor as a warning condition. A matching threshold has been set for the 1 PC1 column. So when this condition occurs, warning event indicators appear in the table and in the Navigator.

  If you find configuration or availability problems, correct them immediately.
- Check the **Service Call Performance by System** table to determine the average request arrival rate, service time in milliseconds, queue length, and byte count per service call for each system.

  If all service calls have been processed on a single system, but service time is well under a millisecond, and the Pending column shows no request queue, there is no performance problem. However, if the same system continues to be used exclusively and the arrival rate increases, service time and queue length would also increase, and a serious performance problem might develop. In such a case, you might consider rebalancing workloads among systems to correct the problem.
- Check the **Top Users by System** table to see whether one or two job names seem to be monopolizing cryptographic services.

  If any of the top job names are relatively unimportant, you might want to reduce their priority on the system.

  Look at the **LastSvcDesc** column, which shows the service requested most recently by each of the top users. Refresh the workspace several times over various intervals and see whether the same service call keeps showing up. If so, that service call is being used heavily by the top users and may be implicated in performance problems. See “Improving performance” on page 73 for suggestions for improving service call performance.

Chapter 9. Detecting CPU looping address spaces

Address spaces can occasionally fall into a CPU loop where a task executes instructions endlessly. Looping is usually an unproductive event, using CPU resources that could be better spent on other workloads. Looping may be a symptom of a storage corruption within the application or perhaps a design flaw that did not anticipate some rare set of environmental circumstances. Whatever the cause of a loop, it has proven difficult in the past to detect that an address space has begun to loop.

On the surface, detecting a CPU loop should be simple: just look for an address space that is using 100% CPU. On z/OS, this is not a good strategy. Most logical partitions (LPARs) are defined with several logical processors (LPs) that can each run instructions for different dispatchable units such as tasks, enclaves, and service request blocks (SRBs). Under these circumstances, does 100% CPU refer to 100% of a single LP or 100% of all the logical processors in the LPAR? Typically, a looping address space is consuming a single logical processor.

Another confounding factor is the z/OS dispatching algorithms, including Workload Manager (WLM), which actively try to distribute processor resources appropriately among all the competing dispatchable units. These algorithms interrupt a looping job to dispatch other work that is not getting CPU resource which it is due under the service policy defined. Eventually, the looping job gets redispatched and squanders more CPU. But because of the interrupts, the job's measured CPU% will drop. Often looping jobs will be parked by the system policy so much that their measured CPU percentage will be too low to be detected by simple threshold settings.

Another popular strategy for detecting CPU loops is to look for address spaces with high CPU usage and low or no I/O activity. As already noted, it is not easy to define what "high CPU usage" means; can it be said with confidence that a job using little or no I/O is clearly misbehaving? What about an application that has much of its working data cached in memory so that it can be more responsive to transactions? This is a good performance strategy, but it means that the application will also present a profile of low or no I/O activity.

OMEGAMON XE on z/OS offers a metric, CPU Loop Index, which is designed to overcome these issues and make detecting CPU loops an easier task. The purpose of this metric is to characterize the intent of an address space to use the CPU. Looping jobs will show an unrelenting intent to use CPU to the exclusion of any other resource. Even when they are parked by WLM or other z/OS policy actions, their intent to use the CPU can be detected.

Determining the intent of a job

The OMEGAMON XE on z/OS Bottleneck Analysis feature samples the execution state of every address space in z/OS every few seconds. An address space typically occupies one execution states at a time. These execution states are things like Using I/O, Waiting for I/O, Waiting for Enqueue, Waiting for HSM Recall, Swapped MPL Delay, Using CPU, and Waiting for CPU. OMEGAMON XE on z/OS defines more than 60 execution states.

The execution states that indicate intent to use CPU are Using CPU, CPU Wait, Using IFA, IFA Wait, Using zIIP, and zIIP Wait.

A note on terminology

z/OS systems have three processor types that can be used by normal applications. The first type is the standard general processor, which is commonly referred to as a CPU in OMEGAMON XE on z/OS. The second processor type is the System z Application Assist Processor, which is usually referred to as a zAAP processor. In its early days this processor was called the Integrated Facility for Applications, or IFA. OMEGAMON XE on z/OS still uses this terminology. zAAP processors are used to execute Java Virtual Machine processes. The third processor type is the System z Integrated Information Processor, or zIIP. zIIPs are most often used to execute DB2 processes, but can be used by other specialized SRB processes.

On every sample taken by Bottleneck Analysis, each address space is classified into one of these execution states and a count is added to the bucket for this execution state. Over time, the profile of execution states builds up for every address space. If you divide the count in any bucket, for example the Using I/O bucket, by the total count summed over all buckets, you get the percentage of samples where the address space was found using I/O. If you sum the counts in all the “intent to use CPU” buckets and divide by the total count, you get the CPU Loop Index. Address spaces in a CPU loop have counts almost exclusively in the “intent to use CPU” buckets. This means the index will likely be at or very near 100%.

The KM5_CPU_Loop_Warn situation

OMEGAMON XE on z/OS includes a model situation, KM5_CPU_Loop_Warn, which can be distributed to all monitored LPARs and started automatically. This situation looks for jobs with a CPU Loop Index above 95%. The situation takes a sample every 5 minutes looking for this condition and requires that 2 consecutive samples show the condition to be true before it raises an alert.

The requirement that two consecutive samples meet the condition before an alert is raised is intended to minimize false positive results. However, it is still possible for false positives to be generated when there are two or more consecutive samples in which different address spaces happen to have high CPU Loop Indexes. While the address spaces were exhibiting looping like symptoms for a short time, they are not really looping. This behavior could be the result of deliberately using a low priority service class designed to allow less important work to soak up whatever CPU is left over in the LPAR. This tactic is usually used to accommodate batch work that has no deadline. Work in this service class would likely have long periods where it is waiting for CPU. Its bottleneck profile would show some CPU using and lots of CPU waiting. When the using and waiting values are combined, the total occasionally exceeds the 95% threshold in the situation.

You can adjust the KM5_CPU_Loop_Warn situation to avoid false positives like these by adding conditions specifically for the low priority work. For example, say that service class LOWBATCH is designed to soak up remaining CPU and the KM5_CPU_Loop_Warn situation is giving false positives for work in that service class. You can adjust the KM5_CPU_Loop_Warn to exclude the service class by taking the following steps.

1. Open the Situation editor from the toolbar.
2. Expand the MVS System node in the navigation tree and scroll down until you can select the KM5_CPU_Loop_Warn situation.

   If necessary, check **Associated with Monitored Application** to see all situations that were written for this type of agent, regardless of where they are distributed.
3. To create a copy, right-click the situation and select **Create Another . . .** from the popup menu. (If you modify the original situation instead of creating your own, your changes will be overwritten the next time OMEGAMON XE on z/OS is updated.)
4. Type a name for the new situation (for example, CPULoop_Warn and click **OK**.

5. Use the **Add conditions** button to add the Service Class attribute from the Address Space Bottlenecks group:

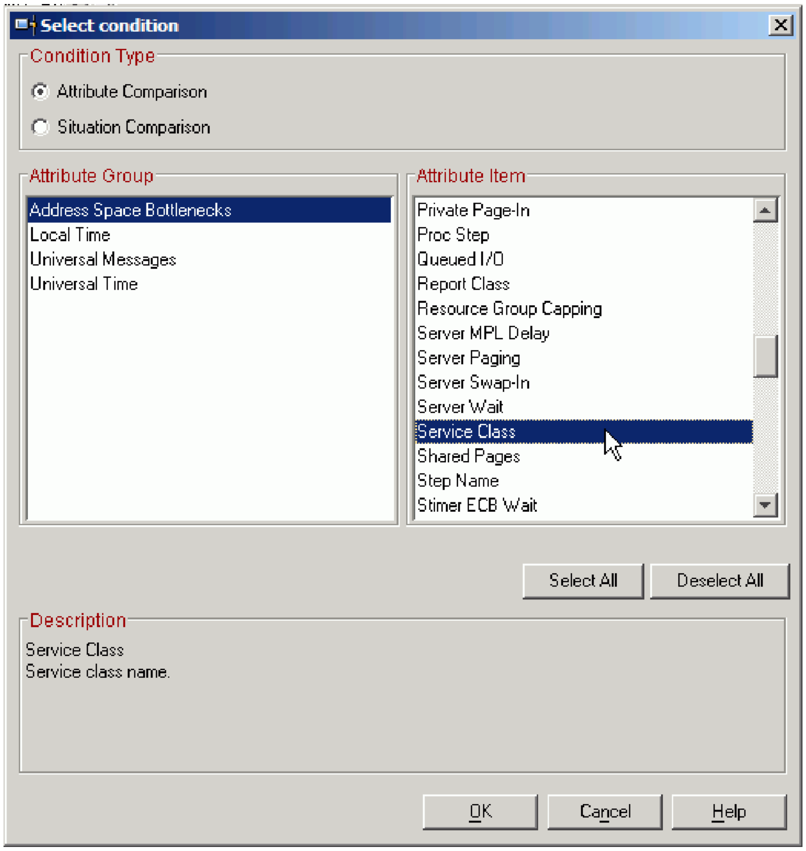

Figure 6. Adding the Service Class attribute

6. Click in the cell beneath the Service Class heading and select != (Not equal) from the dropdown operator menu:

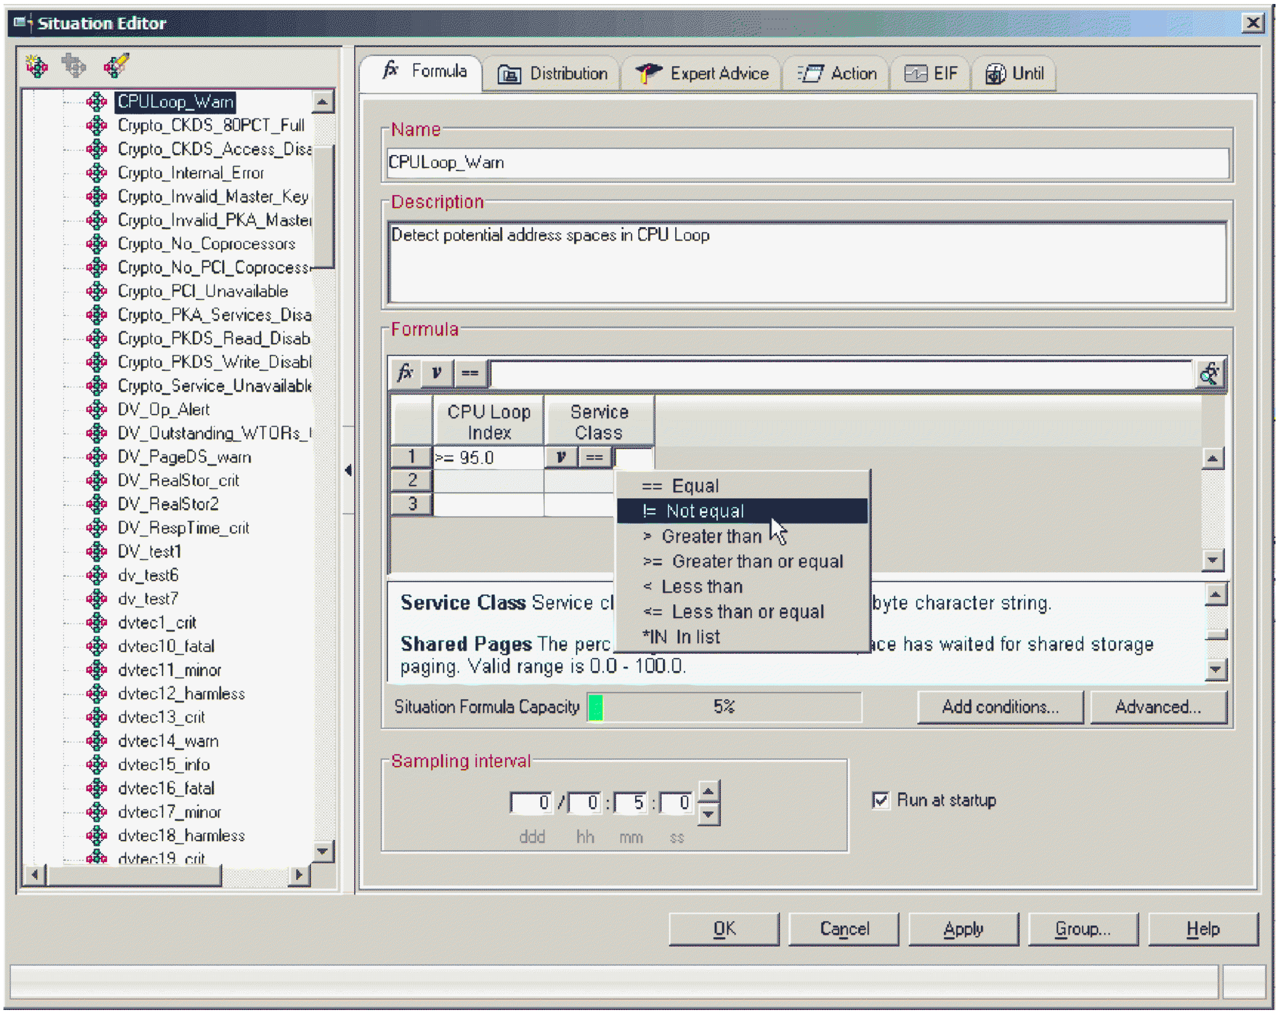

Name: CPULoop_Warn

Description: Detect potential address spaces in CPU Loop

Formula

	CPU Loop Index	Service Class
1	>= 95.0	v ==
2		
3		

Service Class Service cl… …byte character string.

Shared Pages The perc… …ace has waited for shared storage paging. Valid range is 0.0 - 100.0.

Situation Formula Capacity 5%

Sampling interval: 0 / 0 : 5 : 0 (ddd hh mm ss)

☑ Run at startup

Figure 7. Situation editor with Service Class added to formula and != Not equal selected in the Operator popup menu

7. Type the name of the service class to be excluded in single quotation marks (for example, 'LOWBATCH') in the text field beside the operator:

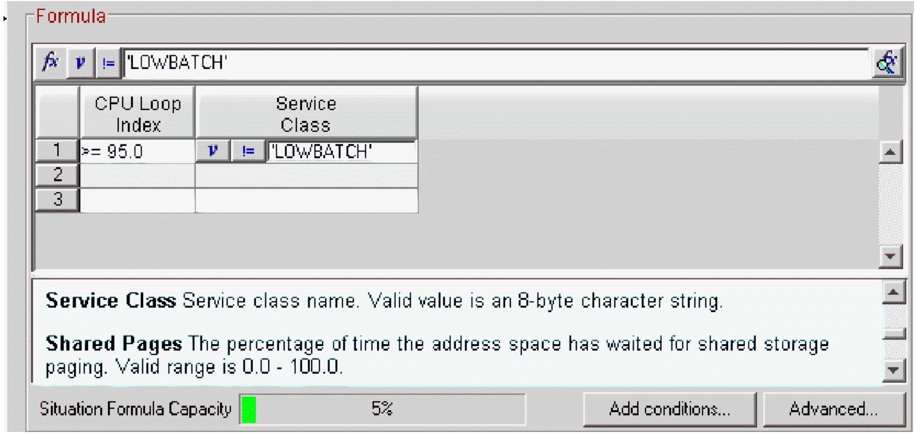

Formula

fx v != 'LOWBATCH'

	CPU Loop Index	Service Class
1	>= 95.0	v != 'LOWBATCH'
2		
3		

Service Class Service class name. Valid value is an 8-byte character string.

Shared Pages The percentage of time the address space has waited for shared storage paging. Valid range is 0.0 - 100.0.

Situation Formula Capacity 5%

Figure 8. Formula section of the Situation editor window, with 'LOWBATCH' added as the value for the expression.

8. Distribute the situation to the managed systems or managed system lists on which you want it to run.

9. Click **OK** to save the new situation and close the Situation editor.

With this modification, work running outside of service class LOWBATCH will continue to trigger the situation when the CPU Loop Index exceeds 95%. You can then create a separate situation for the LOWBATCH work where the threshold has been raised to 99%.

Investigating and identifying looping jobs

The following scenario involves three jobs that show a CPU loop: BKEALCP1, BKEALCP2, and BKEALCP3. This scenario also has three other jobs that run high CPU for a while, but then run I/O activity for a while and then return to the CPU: BKEALIO, BKEALIO2, and BKEALIO3. These jobs represent high CPU usage jobs that are not in a CPU loop.

Using normal ISPF SDSF screens, sorted for CPU% ascending, you can see that these jobs are the most CPU intensive work in the LPAR (Figure 9). Each of the sample looping jobs is using 15.81% of the LPAR's CPU. The I/O jobs are using about 4.33% each.

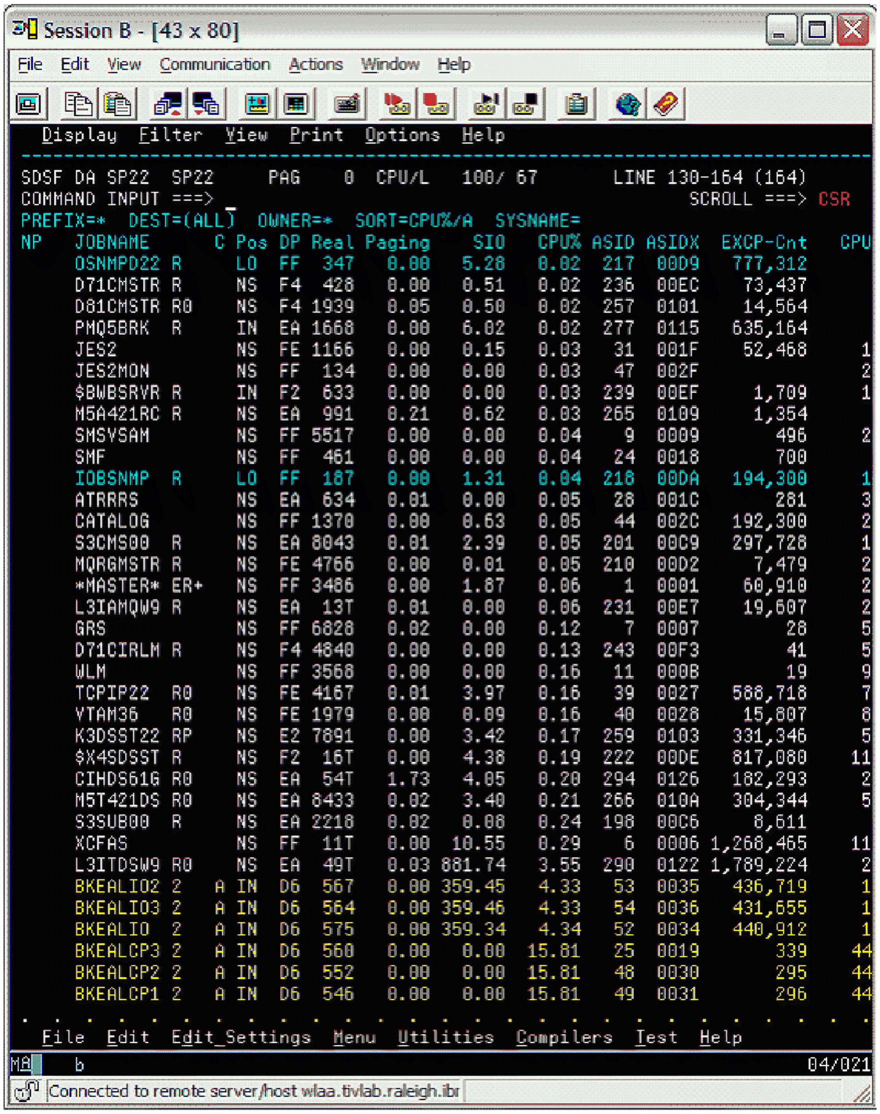

```
  Display  Filter  View  Print  Options  Help
 -------------------------------------------------------------------------------
 SDSF DA SP22  SP22      PAG    0  CPU/L   100/ 67       LINE 130-164 (164)
 COMMAND INPUT ===>                                          SCROLL ===> CSR
 PREFIX=*  DEST=(ALL)  OWNER=*  SORT=CPU%/A  SYSNAME=
 NP   JOBNAME     C Pos DP Real Paging    SIO   CPU% ASID ASIDX  EXCP-Cnt    CPU
      OSNMPD22 R    LO  FF  347   0.00   5.28   0.02  217  00D9   777,312
      D71CMSTR R    NS  F4  428   0.00   0.51   0.02  236  00EC    73,437
      D81CMSTR RO   NS  F4 1939   0.05   0.50   0.02  257  0101    14,564
      PMQ5BRK  R    IN  EA 1668   0.00   6.02   0.02  277  0115   635,164
      JES2          NS  FE 1166   0.00   0.15   0.03   31  001F    52,468      1
      JES2MON       NS  FF  134   0.00   0.00   0.03   47  002F                2
      $BWBSRVR R    IN  F2  633   0.00   0.00   0.03  239  00EF     1,709      1
      M5A421RC R    NS  EA  991   0.21   0.62   0.03  265  0109     1,354
      SMSVSAM       NS  FF 5517   0.00   0.00   0.04    9  0009       496      2
      SMF           NS  FF  461   0.00   0.00   0.04   24  0018       700
      IOBSNMP  R    LO  FF  187   0.00   1.31   0.04  218  00DA   194,300      1
      ATRRRS        NS  EA  634   0.01   0.00   0.05   28  001C       281      3
      CATALOG       NS  FF 1370   0.00   0.63   0.05   44  002C   192,300      2
      S3CMS00  R    NS  EA 8043   0.01   2.39   0.05  201  00C9   297,728      1
      MQRGMSTR R    NS  FE 4766   0.00   0.01   0.05  210  00D2     7,479      2
      *MASTER* ER+  NS  FF 3486   0.00   1.87   0.06    1  0001    60,910      2
      L3IAMQW9 R    NS  EA  13T   0.01   0.00   0.06  231  00E7    19,607      2
      GRS           NS  FF 6828   0.02   0.00   0.12    7  0007        28      5
      D71CIRLM R    NS  F4 4840   0.00   0.00   0.13  243  00F3        41      5
      WLM           NS  FF 3568   0.00   0.00   0.16   11  000B        19      9
      TCPIP22  RO   NS  FE 4167   0.01   3.97   0.16   39  0027   588,718      7
      VTAM36   RO   NS  FE 1979   0.00   0.09   0.16   40  0028    15,807      8
      K3DSST22 RP   NS  E2 7891   0.00   3.42   0.17  259  0103   331,346      5
      $X4SDSST R    NS  F2  16T   0.00   4.38   0.19  222  00DE   817,080     11
      CIHDS61G RO   NS  EA  54T   1.73   4.05   0.20  294  0126   182,293      2
      M5T421DS RO   NS  EA 8433   0.02   3.40   0.21  266  010A   304,344      5
      S3SUB00  R    NS  EA 2218   0.02   0.08   0.24  198  00C6     8,611
      XCFAS         NS  FF  11T   0.00  10.55   0.29    6  0006 1,268,465     11
      L3ITDSW9 RO   NS  EA  49T   0.03 881.74   3.55  290  0122 1,789,224      2
      BKEALIO2 2  A IN  D6  567   0.00 359.45   4.33   53  0035   436,719      1
      BKEALIO3 2  A IN  D6  564   0.00 359.46   4.33   54  0036   431,655      1
      BKEALIO  2  A IN  D6  575   0.00 359.34   4.34   52  0034   440,912      1
      BKEALCP3 2  A IN  D6  560   0.00   0.00  15.81   25  0019       339     44
      BKEALCP2 2  A IN  D6  552   0.00   0.00  15.81   48  0030       295     44
      BKEALCP1 2  A IN  D6  546   0.00   0.00  15.81   49  0031       296     44
 .  .  .  .  .  .  .  .  .  .  .  .  .  .  .  .  .  .  .  .  .  .  .  .  .  .  .
  File  Edit  Edit_Settings  Menu  Utilities  Compilers  Test  Help
MA      b                                                                 04/021
```

Figure 9. Sample jobs

It is not likely that a performance analyst looking at this report would expect that any of these jobs was in a CPU Loop. The KM5_CPU_Loop_Warn situation, however, does see BKEALCP1, BKEALCP2, and BKEALCP3 as suspect jobs and an event is raised in the Tivoli Enterprise Portal (Figure 10).

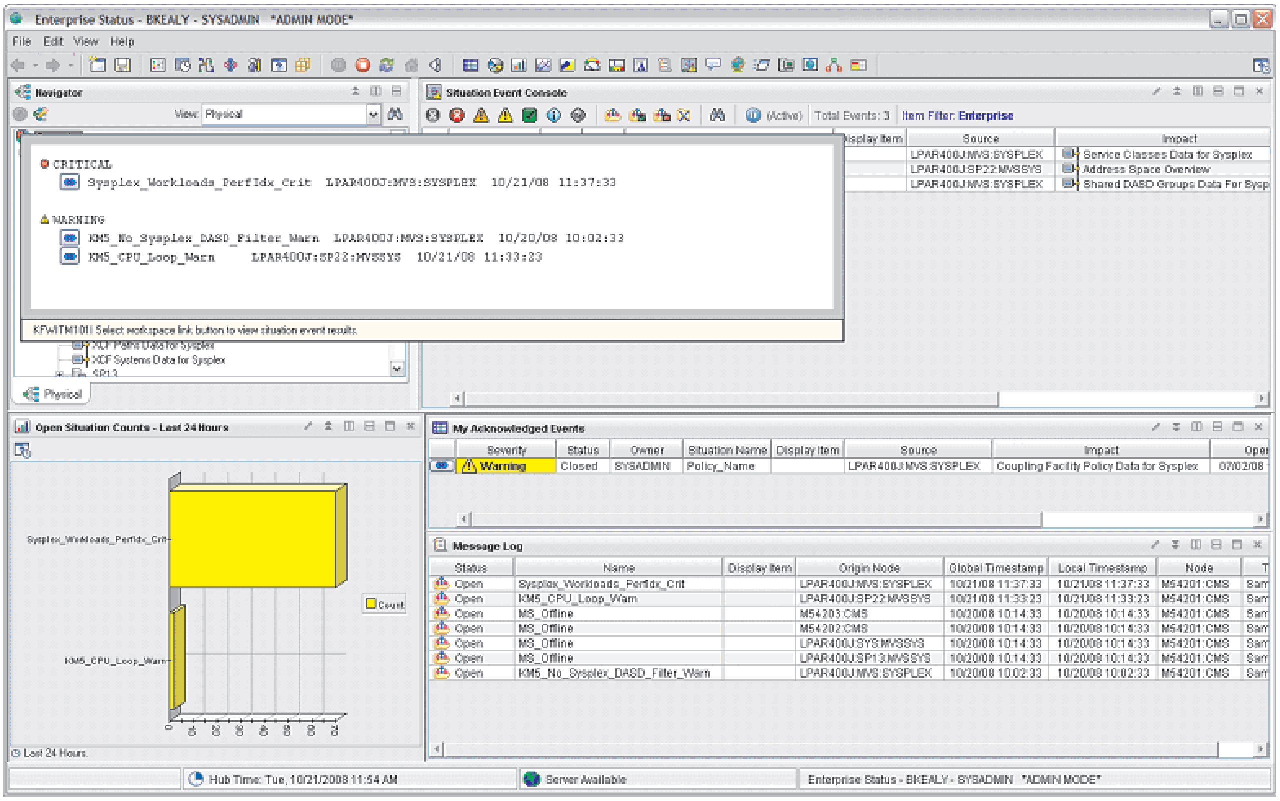
Enterprise Status - BKEALY - SYSADMIN *ADMIN MODE*
File Edit View Help
Navigator
View: Physical
Situation Event Console
(Active)
Total Events: 3
Item Filter: Enterprise
XCF Paths Data for Sysplex
XCF Systems Data for Sysplex
SP13
Physical
Open Situation Counts - Last 24 Hours
Sysplex_Workloads_Perfidx_Crit
KM5_CPU_Loop_Warn
Count
Last 24 Hours
My Acknowledged Events
Message Log
Hub Time: Tue, 10/21/2008 11:54 AM
Server Available
Enterprise Status - BKEALY - SYSADMIN *ADMIN MODE*

```
CRITICAL
   Sysplex_Workloads_PerfIdx_Crit  LPAR400J:MVS:SYSPLEX  10/21/08 11:37:33

WARNING
   KM5_No_Sysplex_DASD_Filter_Warn  LPAR400J:MVS:SYSPLEX  10/20/08 10:02:33
   KM5_CPU_Loop_Warn      LPAR400J:SP22:MVSSYS  10/21/08 11:33:23
```

KFWITM101I Select workspace link button to view situation event results.

Display Item	Source	Impact
	LPAR400J:MVS:SYSPLEX	Service Classes Data for Sysplex
	LPAR400J:SP22:MVSSYS	Address Space Overview
	LPAR400J:MVS:SYSPLEX	Shared DASD Groups Data For Sysp

	Severity	Status	Owner	Situation Name	Display Item	Source	Impact	Oper
	Warning	Closed	SYSADMIN	Policy_Name		LPAR400J:MVS:SYSPLEX	Coupling Facility Policy Data for Sysplex	07/02/08

Status	Name	Display Item	Origin Node	Global Timestamp	Local Timestamp	Node	T
Open	Sysplex_Workloads_Perfidx_Crit		LPAR400J:MVS:SYSPLEX	10/21/08 11:37:33	10/21/08 11:37:33	M54201:CMS	Sam
Open	KM5_CPU_Loop_Warn		LPAR400J:SP22:MVSSYS	10/21/08 11:33:23	10/21/08 11:33:23	M54201:CMS	Sam
Open	MS_Offline		M54203:CMS	10/20/08 10:14:33	10/20/08 10:14:33	M54201:CMS	Sam
Open	MS_Offline		M54202:CMS	10/20/08 10:14:33	10/20/08 10:14:33	M54201:CMS	Sam
Open	MS_Offline		LPAR400J:SYS:MVSSYS	10/20/08 10:14:33	10/20/08 10:14:33	M54201:CMS	Sam
Open	MS_Offline		LPAR400J:SP13:MVSSYS	10/20/08 10:14:33	10/20/08 10:14:33	M54201:CMS	Sam
Open	KM5_No_Sysplex_DASD_Filter_Warn		LPAR400J:MVS:SYSPLEX	10/20/08 10:02:33	10/20/08 10:02:33	M54201:CMS	Sam

Figure 10. Event flyover for KM5_CPU_Loop Warn

The situation event workspace for KM5_CPU_Loop_Warn (Figure 11 on page 81) indicates that BKEALCP1, BKEALCP2, and BKEALCP3 have very high CPU Loop Index values. The accompanying expert advice notes that a high CPU Loop Index is not a guarantee that the job is looping and it is possible that a well behaved job is in a normal period of intensive CPU activity. It is up to the site analyst to recognize jobs that normally run long periods of CPU instructions. Some scientific workloads may fit this profile, but it unlikely that a normal business workload would use this much CPU for this long (5 to 10 minutes at minimum). Users are advised to examine the address spaces.

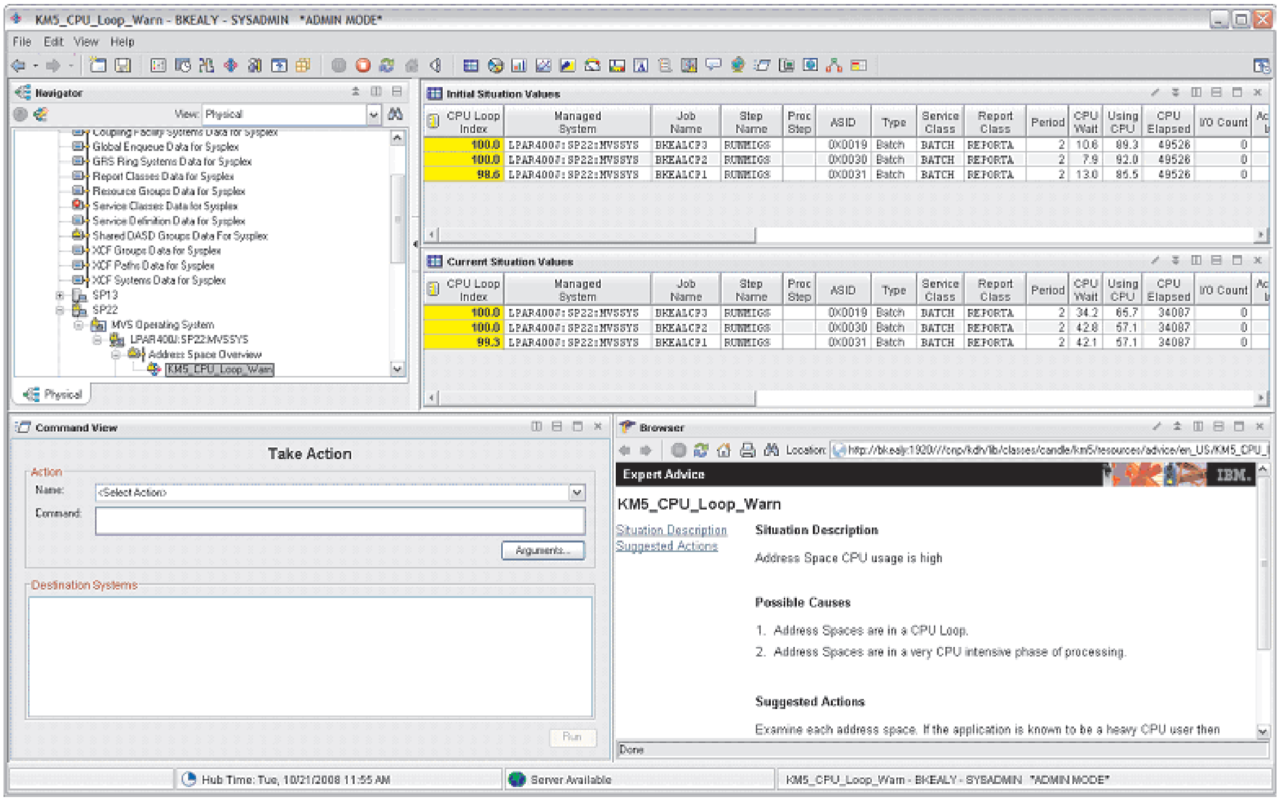

KM5_CPU_Loop_Warn - BKEALY - SYSADMIN *ADMIN MODE*
File Edit View Help
Navigator
View: Physical
Coupling Facility Systems Data for Sysplex
Global Enqueue Data for Sysplex
GRS Ring Systems Data for Sysplex
Report Classes Data for Sysplex
Resource Groups Data for Sysplex
Service Classes Data for Sysplex
Service Definition Data for Sysplex
Shared DASD Groups Data For Sysplex
XCF Groups Data for Sysplex
XCF Paths Data for Sysplex
XCF Systems Data for Sysplex
SP13
SP22
MVS Operating System
LPAR400J:SP22:MVSSYS
Address Space Overview
KM5_CPU_Loop_Warn
Physical
Command View
Take Action
Action
Name:
<Select Action>
Command:
Arguments...
Destination Systems
Run
Browser
Location: http://bkealy:1920///cnp/kdh/lib/classes/candle/km5/resources/advice/en_US/KM5_CPU_
Expert Advice
IBM.
Done
Hub Time: Tue, 10/21/2008 11:55 AM
Server Available
KM5_CPU_Loop_Warn - BKEALY - SYSADMIN *ADMIN MODE*

Initial Situation Values

CPU Loop Index	Managed System	Job Name	Step Name	Proc Step	ASID	Type	Service Class	Report Class	Period	CPU Wait	Using CPU	CPU Elapsed	I/O Count	Ac
100.0	LPAR400J:SP22:MVSSYS	BKEALCP3	RUNMIGS		0X0019	Batch	BATCH	REPORTA	2	10.6	89.3	49526	0	
100.0	LPAR400J:SP22:MVSSYS	BKEALCP2	RUNMIGS		0X0030	Batch	BATCH	REPORTA	2	7.9	92.0	49526	0	
98.6	LPAR400J:SP22:MVSSYS	BKEALCP1	RUNMIGS		0X0031	Batch	BATCH	REPORTA	2	13.0	85.5	49526	0	

Current Situation Values

CPU Loop Index	Managed System	Job Name	Step Name	Proc Step	ASID	Type	Service Class	Report Class	Period	CPU Wait	Using CPU	CPU Elapsed	I/O Count	Ac
100.0	LPAR400J:SP22:MVSSYS	BKEALCP3	RUNMIGS		0X0019	Batch	BATCH	REPORTA	2	34.2	65.7	34087	0	
100.0	LPAR400J:SP22:MVSSYS	BKEALCP2	RUNMIGS		0X0030	Batch	BATCH	REPORTA	2	42.8	57.1	34087	0	
99.3	LPAR400J:SP22:MVSSYS	BKEALCP1	RUNMIGS		0X0031	Batch	BATCH	REPORTA	2	42.1	57.1	34087	0	

KM5_CPU_Loop_Warn

Situation Description
Suggested Actions

Situation Description

Address Space CPU usage is high

Possible Causes

1. Address Spaces are in a CPU Loop.
2. Address Spaces are in a very CPU intensive phase of processing.

Suggested Actions

Examine each address space. If the application is known to be a heavy CPU user then

Figure 11. Situation event workspace for KM5_CPU_Loop_Warn

A number of OMEGAMON XE on z/OS workspaces can be used to further examine these jobs, starting at the Address Space Overview workspace (Figure 12 on page 82).

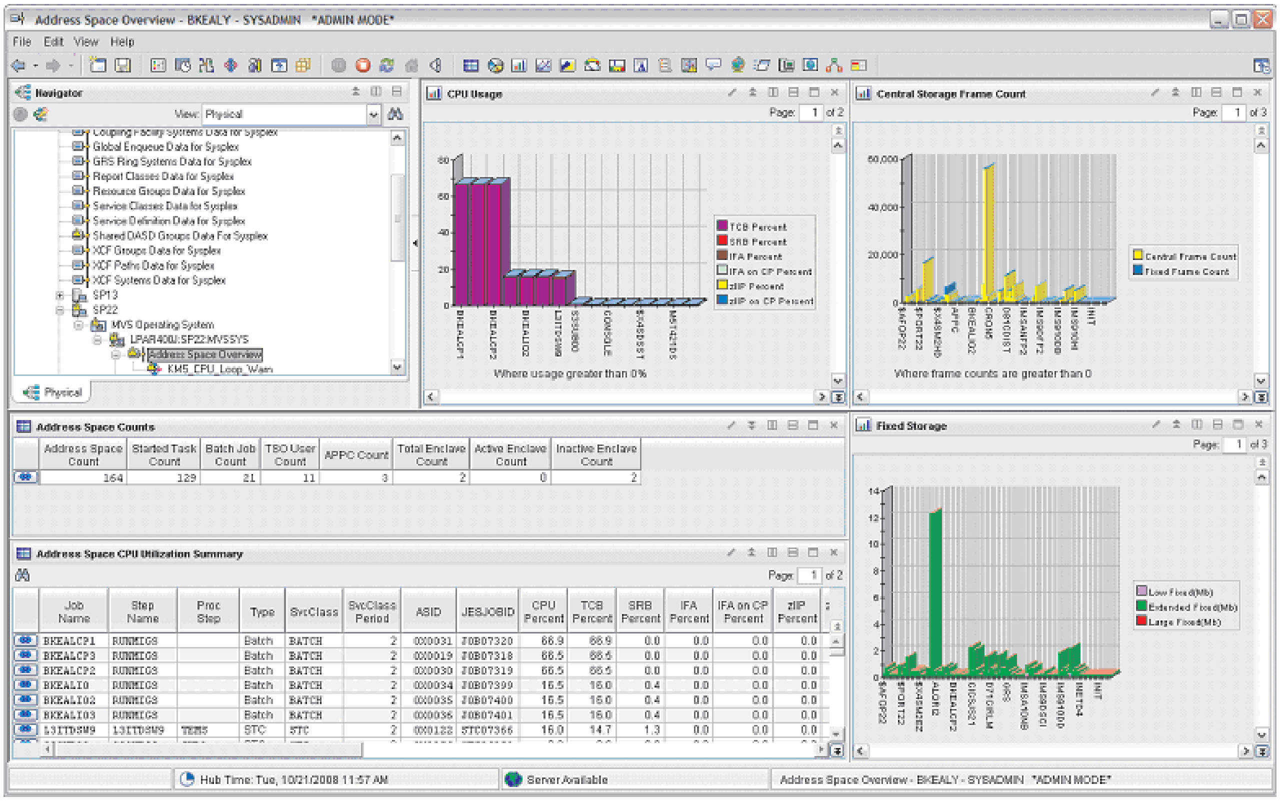

Address Space Counts

	Address Space Count	Started Task Count	Batch Job Count	TSO User Count	APPC Count	Total Enclave Count	Active Enclave Count	Inactive Enclave Count
👁	164	129	21	11	3	2	0	2

Address Space CPU Utilization Summary

Page: 1 of 2

	Job Name	Step Name	Proc Step	Type	SvcClass	SvcClass Period	ASID	JESJOBID	CPU Percent	TCB Percent	SRB Percent	IFA Percent	IFA on CP Percent	zIIP Percent
👁	BKEALCP1	RUNMIGS		Batch	BATCH	2	0X0031	JOB07320	66.9	66.9	0.0	0.0	0.0	0.0
👁	BKEALCP3	RUNMIGS		Batch	BATCH	2	0X0019	JOB07318	66.5	66.5	0.0	0.0	0.0	0.0
👁	BKEALCP2	RUNMIGS		Batch	BATCH	2	0X0030	JOB07319	66.5	66.5	0.0	0.0	0.0	0.0
👁	BKEALIO	RUNMIGS		Batch	BATCH	2	0X0034	JOB07399	16.5	16.0	0.4	0.0	0.0	0.0
👁	BKEALIO2	RUNMIGS		Batch	BATCH	2	0X0035	JOB07400	16.5	16.0	0.4	0.0	0.0	0.0
👁	BKEALIO3	RUNMIGS		Batch	BATCH	2	0X0036	JOB07401	16.5	16.0	0.4	0.0	0.0	0.0
👁	L3ITDSW9	L3ITDSW9	TEMS	STC	STC	2	0X0122	STC07366	16.0	14.7	1.3	0.0	0.0	0.0

Figure 12. Address Space Overview workspace

The six jobs are displayed at the top of the Address Space CPU Utilization Summary view. In this view, the rows are sorted in descending order of CPU Percent so that no higher CPU users are missed on subsequent view pages. To look at the bottleneck analysis for these jobs to see which execution states they are spending their time in, you right-click the 👁 icon in the Address Space Counts view and select the Address Space Bottleneck Summary link.

The Address Space Bottleneck Summary workspace (Figure 13 on page 83) displays the CPU Loop Index value for the three suspect looping jobs. You notice that the jobs that are using heavy CPU but are not looping, BKEALIO, BKEALIO2, and BKEALIO3 all have significantly lower indexes. When you look at the other bottleneck data you see that the Using CPU values by themselves are not unequivocal indicators of a loop. It is only when you combine the Using CPU and CPU Wait numbers that you see the dramatic effect.

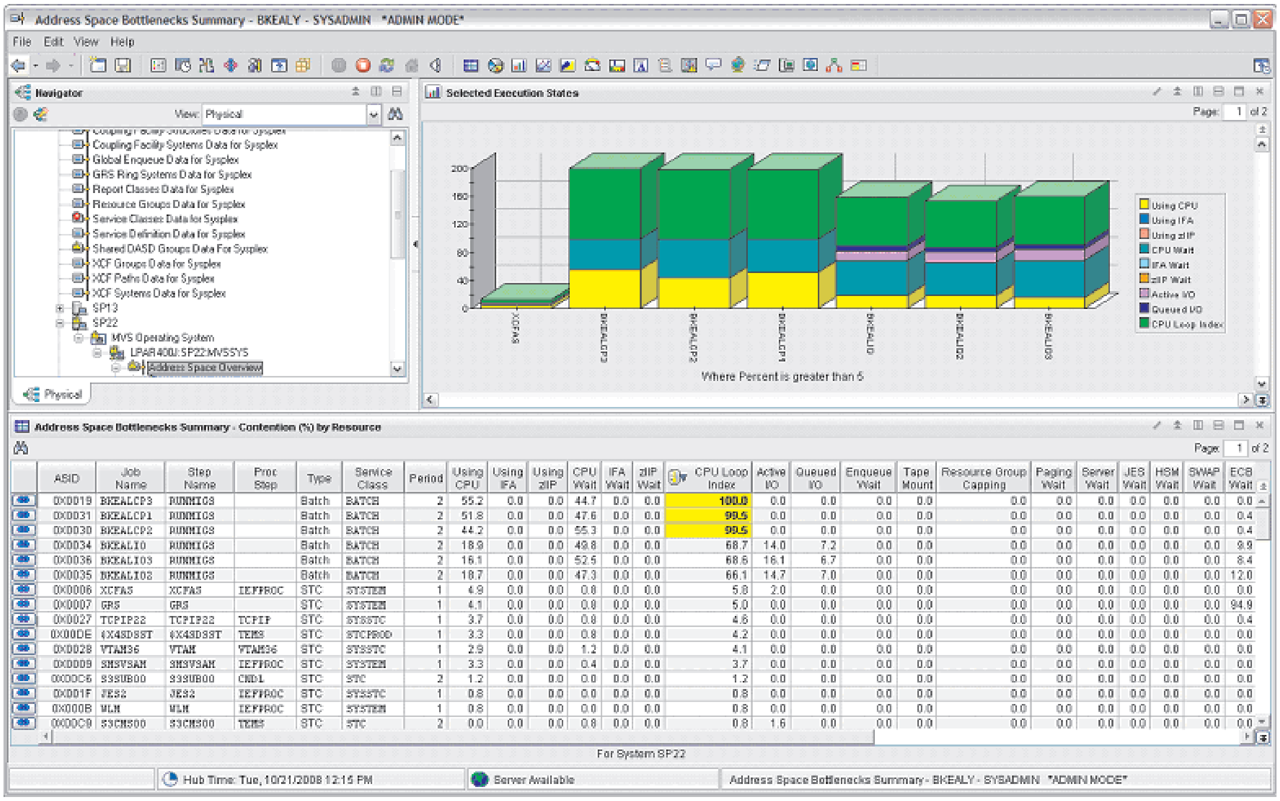
Address Space Bottlenecks Summary - BKEALY - SYSADMIN *ADMIN MODE*
File Edit View Help
Navigator
View: Physical
Coupling Facility Structures Data for Sysplex
Coupling Facility Systems Data for Sysplex
Global Enqueue Data for Sysplex
GRS Ring Systems Data for Sysplex
Report Classes Data for Sysplex
Resource Groups Data for Sysplex
Service Classes Data for Sysplex
Service Definition Data for Sysplex
Shared DASD Groups Data For Sysplex
XCF Groups Data for Sysplex
XCF Paths Data for Sysplex
XCF Systems Data for Sysplex
SP13
SP22
MVS Operating System
LPAR400J:SP22:MVSSYS
Address Space Overview
Physical
Selected Execution States
Page: 1 of 2
200
160
120
80
40
0
XCFAS
BKEALCP3
BKEALCP2
BKEALCP1
BKEALIO
BKEALIO2
BKEALIO3
Where Percent is greater than 5
Using CPU
Using IFA
Using zIIP
CPU Wait
IFA Wait
zIIP Wait
Active I/O
Queued I/O
CPU Loop Index
Address Space Bottlenecks Summary - Contention (%) by Resource
Page: 1 of 2

	ASID	Job Name	Step Name	Proc Step	Type	Service Class	Period	Using CPU	Using IFA	Using zIIP	CPU Wait	IFA Wait	zIIP Wait	CPU Loop Index	Active I/O	Queued I/O	Enqueue Wait	Tape Mount	Resource Group Capping	Paging Wait	Server Wait	JES Wait	HSM Wait	SWAP Wait	ECB Wait
	0X0019	BKEALCP3	RUNMIGS		Batch	BATCH	2	55.2	0.0	0.0	44.7	0.0	0.0	100.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	0X0031	BKEALCP1	RUNMIGS		Batch	BATCH	2	51.8	0.0	0.0	47.6	0.0	0.0	99.5	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4
	0X0030	BKEALCP2	RUNMIGS		Batch	BATCH	2	44.2	0.0	0.0	55.3	0.0	0.0	99.5	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4
	0X0034	BKEALIO	RUNMIGS		Batch	BATCH	2	18.9	0.0	0.0	49.8	0.0	0.0	68.7	14.0	7.2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	9.9
	0X0036	BKEALIO3	RUNMIGS		Batch	BATCH	2	16.1	0.0	0.0	52.5	0.0	0.0	68.6	16.1	6.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	8.4
	0X0035	BKEALIO2	RUNMIGS		Batch	BATCH	2	18.7	0.0	0.0	47.3	0.0	0.0	66.1	14.7	7.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	12.0
	0X0006	XCFAS	XCFAS	IEFPROC	STC	SYSTEM	1	4.9	0.0	0.0	0.8	0.0	0.0	5.8	2.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	0X0007	GRS	GRS		STC	SYSTEM	1	4.1	0.0	0.0	0.8	0.0	0.0	5.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	94.9
	0X0027	TCPIP22	TCPIP22	TCPIP	STC	SYSSTC	1	3.7	0.0	0.0	0.8	0.0	0.0	4.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.4
	0X00DE	4X43DSST	4X43DSST	TEMS	STC	STCPROD	1	3.3	0.0	0.0	0.8	0.0	0.0	4.2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	0X0028	VTAM36	VTAM	VTAM36	STC	SYSSTC	1	2.9	0.0	0.0	1.2	0.0	0.0	4.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	0X0009	SMSVSAM	SMSVSAM	IEFPROC	STC	SYSTEM	1	3.3	0.0	0.0	0.4	0.0	0.0	3.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	0X00C5	S3SUB00	S3SUB00	CNDL	STC	STC	2	1.2	0.0	0.0	0.0	0.0	0.0	1.2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	0X001F	JES2	JES2	IEFPROC	STC	SYSSTC	1	0.8	0.0	0.0	0.0	0.0	0.0	0.8	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	0X000B	WLM	WLM	IEFPROC	STC	SYSTEM	1	0.8	0.0	0.0	0.0	0.0	0.0	0.8	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	0X00C9	S3CMS00	S3CMS00	TEMS	STC	STC	2	0.0	0.0	0.0	0.8	0.0	0.0	0.8	1.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

For System SP22

Hub Time: Tue, 10/21/2008 12:15 PM

Server Available

Address Space Bottlenecks Summary - BKEALY - SYSADMIN *ADMIN MODE*

Figure 13. Address Space Bottleneck Summary workspace

To explore job BKEALCP3 in more detail, you use a dynamic link to drill down to OMEGAMON for MVS. You right-click the ⛓ navigation link in the row for job BKEALCP3 and select the OMEGAMON for MVS – Job Details link. The OMEGAMON for MVS – Job Details workspace opens and a log-on window is displayed. You enter a user ID and password for the terminal session (if required), and click **OK**. The Examine Details for Job BKEALCP3 screen space in OMEGAMON for MVS opens in the terminal view (see Figure 14 on page 84).

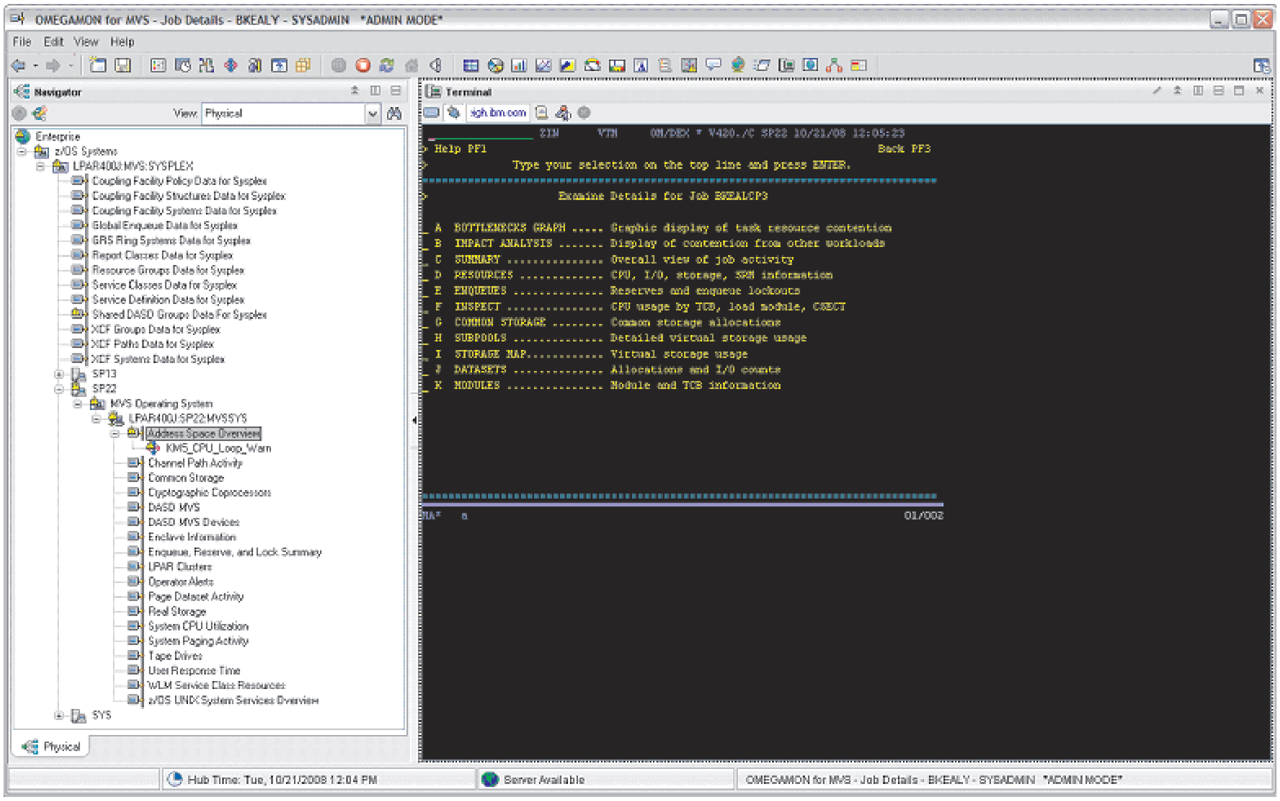
OMEGAMON for MVS - Job Details - BKEALY - SYSADMIN *ADMIN MODE*
File Edit View Help
Navigator
View: Physical
Enterprise
z/OS Systems
LPAR400J:MVS:SYSPLEX
Coupling Facility Policy Data for Sysplex
Coupling Facility Structures Data for Sysplex
Coupling Facility Systems Data for Sysplex
Global Enqueue Data for Sysplex
GRS Ring Systems Data for Sysplex
Report Classes Data for Sysplex
Resource Groups Data for Sysplex
Service Classes Data for Sysplex
Service Definition Data for Sysplex
Shared DASD Groups Data For Sysplex
XCF Groups Data for Sysplex
XCF Paths Data for Sysplex
XCF Systems Data for Sysplex
SP13
SP22
MVS Operating System
LPAR400J:SP22:MVSSYS
Address Space Overview
KM5_CPU_Loop_Warn
Channel Path Activity
Common Storage
Cryptographic Coprocessors
DASD MVS
DASD MVS Devices
Enclave Information
Enqueue, Reserve, and Lock Summary
LPAR Clusters
Operator Alerts
Page Dataset Activity
Real Storage
System CPU Utilization
System Paging Activity
Tape Drives
User Response Time
WLM Service Class Resources
z/OS UNIX System Services Overview
SYS
Physical
Terminal
sgh.ibm.com
Hub Time: Tue, 10/21/2008 12:04 PM
Server Available
OMEGAMON for MVS - Job Details - BKEALY - SYSADMIN *ADMIN MODE*

```
________________     ZIN       VTM      OM/DEX * V420./C SP22 10/21/08 12:05:23
> Help PF1                                                           Back PF3
>             Type your selection on the top line and press ENTER.
===============================================================================
>                    Examine Details for Job BKEALCP3

_ A  BOTTLENECKS GRAPH ..... Graphic display of task resource contention
_ B  IMPACT ANALYSIS ....... Display of contention from other workloads
_ C  SUMMARY ............... Overall view of job activity
_ D  RESOURCES ............. CPU, I/O, storage, SRM information
_ E  ENQUEUES .............. Reserves and enqueue lockouts
_ F  INSPECT ............... CPU usage by TCB, load module, CSECT
_ G  COMMON STORAGE ........ Common storage allocations
_ H  SUBPOOLS .............. Detailed virtual storage usage
_ I  STORAGE MAP............ Virtual storage usage
_ J  DATASETS .............. Allocations and I/O counts
_ K  MODULES ............... Module and TCB information








===============================================================================
MA*    a                                                                 01/002
```

Figure 14. Examine Details for Job BKEALCP3 screen space

You can now continue exploring BKEALCP3 by running an INSPECT of the job (enter s beside F INSPECT). The INSPECT produces the results shown in Figure 15 on page 85.

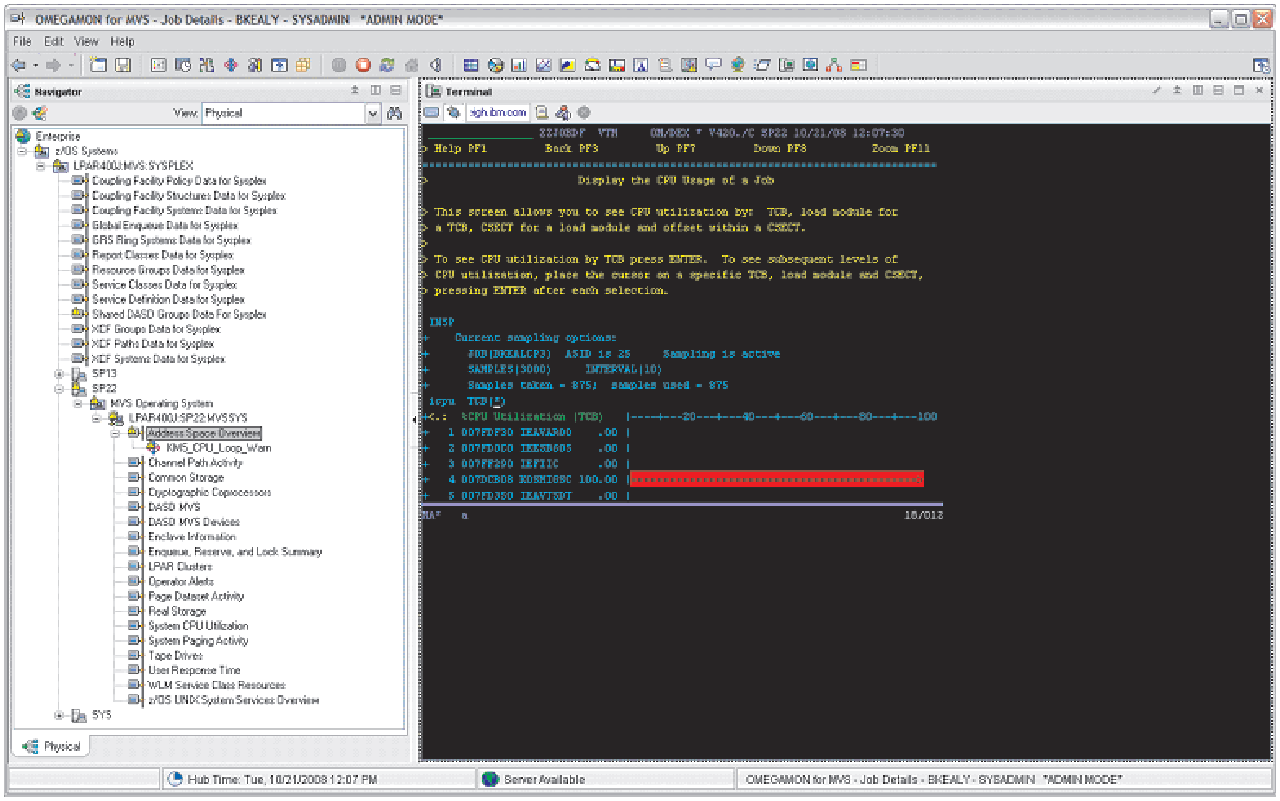
OMEGAMON for MVS - Job Details - BKEALY - SYSADMIN *ADMIN MODE*
File Edit View Help
Navigator
View: Physical
Enterprise
z/OS Systems
LPAR400J:MVS:SYSPLEX
Coupling Facility Policy Data for Sysplex
Coupling Facility Structures Data for Sysplex
Coupling Facility Systems Data for Sysplex
Global Enqueue Data for Sysplex
GRS Ring Systems Data for Sysplex
Report Classes Data for Sysplex
Resource Groups Data for Sysplex
Service Classes Data for Sysplex
Service Definition Data for Sysplex
Shared DASD Groups Data For Sysplex
XCF Groups Data for Sysplex
XCF Paths Data for Sysplex
XCF Systems Data for Sysplex
SP13
SP22
MVS Operating System
LPAR400J:SP22:MVSSYS
Address Space Overview
KM5_CPU_Loop_Warn
Channel Path Activity
Common Storage
Cryptographic Coprocessors
DASD MVS
DASD MVS Devices
Enclave Information
Enqueue, Reserve, and Lock Summary
LPAR Clusters
Operator Alerts
Page Dataset Activity
Real Storage
System CPU Utilization
System Paging Activity
Tape Drives
User Response Time
WLM Service Class Resources
z/OS UNIX System Services Overview
SYS
Physical
Terminal
sigh.ibm.com
Hub Time: Tue, 10/21/2008 12:07 PM
Server Available
OMEGAMON for MVS - Job Details - BKEALY - SYSADMIN *ADMIN MODE*

```
                  ZZ/OBDF  VTM     OM/DEX * V420./C SP22 10/21/08 12:07:30
> Help PF1          Back PF3          Up PF7          Down PF8          Zoom PF11
================================================================================
>                        Display the CPU Usage of a Job

> This screen allows you to see CPU utilization by:  TCB, load module for
> a TCB, CSECT for a load module and offset within a CSECT.
>
> To see CPU utilization by TCB press ENTER.  To see subsequent levels of
> CPU utilization, place the cursor on a specific TCB, load module and CSECT,
> pressing ENTER after each selection.

 INSP
+    Current sampling options:
+      JOB(BKEALCP3)  ASID is 25      Sampling is active
+      SAMPLES(3000)      INTERVAL(10)
+      Samples taken = 875;  samples used = 875
 icpu  TCB(*)
+<.:  %CPU Utilization (TCB)    |----+---20---+---40---+---60---+---80---+---100
+   1 007FDF30 IEAVAR00    .00 |
+   2 007FD0C0 IEESB605    .00 |
+   3 007FF290 IEFIIC      .00 |
+   4 007DCB08 KOSMIGSC 100.00 |==================================================>
+   5 007FD350 IEAVTSDT    .00 |
MA*   a                                                                    18/012
```

Figure 15. Inspect results for BKEALCP3

So you know that all the CPU is being generated in the CSECT KOSMIGSC. You narrow this to specific instruction ranges by moving the cursor over the hot task and changing the OFFset and Granularity values (see Figure 16 on page 86). So far you see that all the CPU is within the narrow range of instructions from offset x'4A0' to x'4F0'.

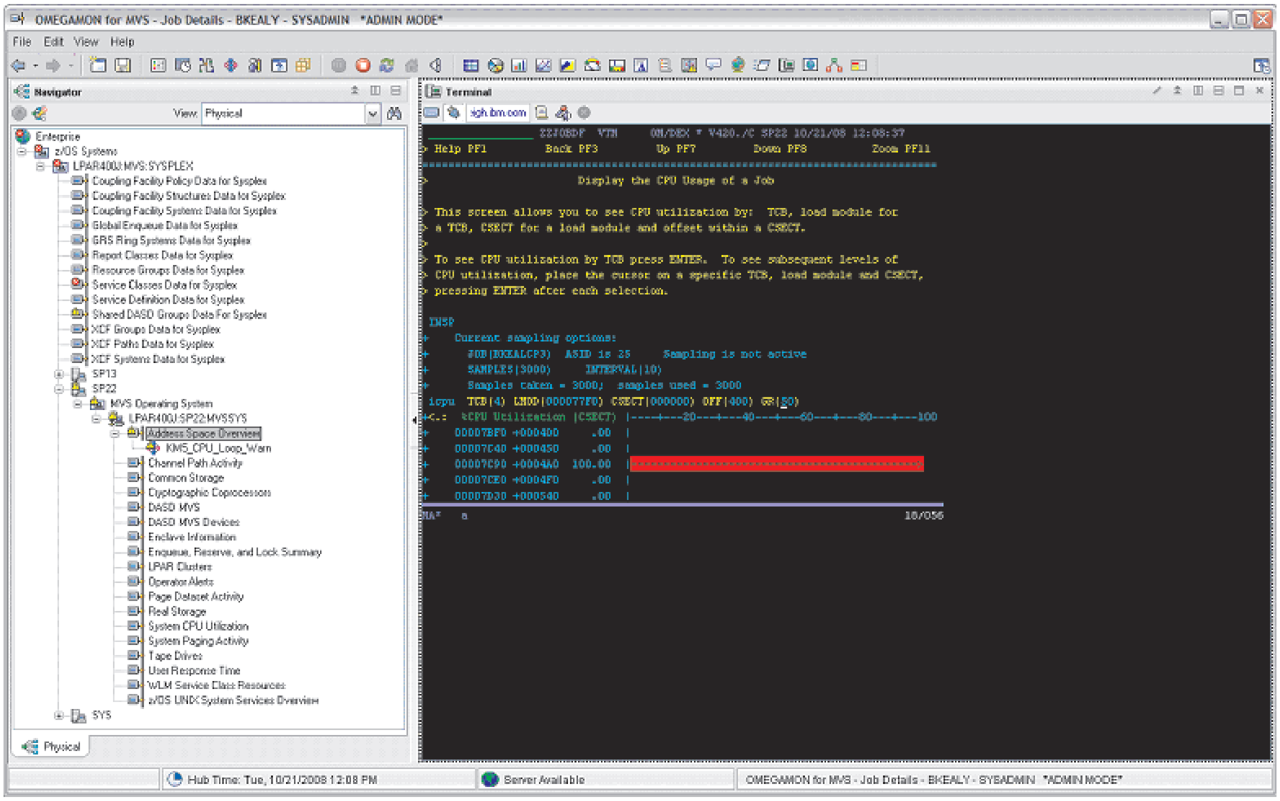
OMEGAMON for MVS - Job Details - BKEALY - SYSADMIN *ADMIN MODE*
File
Edit
View
Help
Navigator
View: Physical
Enterprise
z/OS Systems
LPAR400J:MVS:SYSPLEX
Coupling Facility Policy Data for Sysplex
Coupling Facility Structures Data for Sysplex
Coupling Facility Systems Data for Sysplex
Global Enqueue Data for Sysplex
GRS Ring Systems Data for Sysplex
Report Classes Data for Sysplex
Resource Groups Data for Sysplex
Service Classes Data for Sysplex
Service Definition Data for Sysplex
Shared DASD Groups Data For Sysplex
XCF Groups Data for Sysplex
XCF Paths Data for Sysplex
XCF Systems Data for Sysplex
SP13
SP22
MVS Operating System
LPAR400J:SP22:MVSSYS
Address Space Overview
KM5_CPU_Loop_Warn
Channel Path Activity
Common Storage
Cryptographic Coprocessors
DASD MVS
DASD MVS Devices
Enclave Information
Enqueue, Reserve, and Lock Summary
LPAR Clusters
Operator Alerts
Page Dataset Activity
Real Storage
System CPU Utilization
System Paging Activity
Tape Drives
User Response Time
WLM Service Class Resources
z/OS UNIX System Services Overview
SYS
Physical
Terminal
sgh.ibm.com
Hub Time: Tue, 10/21/2008 12:08 PM
Server Available
OMEGAMON for MVS - Job Details - BKEALY - SYSADMIN *ADMIN MODE*

```
                ZZJOBDP  VTM     ON/DEX * V420./C SP22 10/21/08 12:08:37
> Help PF1       Back PF3       Up PF7        Down PF8         Zoom PF11
=============================================================================
>                     Display the CPU Usage of a Job

> This screen allows you to see CPU utilization by:  TCB, load module for
> a TCB, CSECT for a load module and offset within a CSECT.
>
> To see CPU utilization by TCB press ENTER.  To see subsequent levels of
> CPU utilization, place the cursor on a specific TCB, load module and CSECT,
> pressing ENTER after each selection.

 INSP
+    Current sampling options:
+      JOB(BKEALCP3)  ASID is 25     Sampling is not active
+      SAMPLES(3000)      INTERVAL(10)
+      Samples taken = 3000;  samples used = 3000
 icpu  TCB(4) LMOD(000077F0) CSECT(000000) OFF(400) GR(50)
+<.:  %CPU Utilization (CSECT) |----+---20---+---40---+---60---+---80---+---100
+    00007BF0 +000400      .00  |
+    00007C40 +000450      .00  |
+    00007C90 +0004A0   100.00  |------------------------------------------------>
+    00007CE0 +0004F0      .00  |
+    00007D30 +000540      .00  |
MA*    a                                                                 18/056
```

Figure 16. Drill down of INSPECT data

You narrow the granularity even further (Figure 17 on page 87). Now you find that the looping instructions lie between offsets x’4D5’ and x’4DF’. With this detail you can go back to the application source and find the offending instructions. You can also take action from this interface to cancel the looping job.

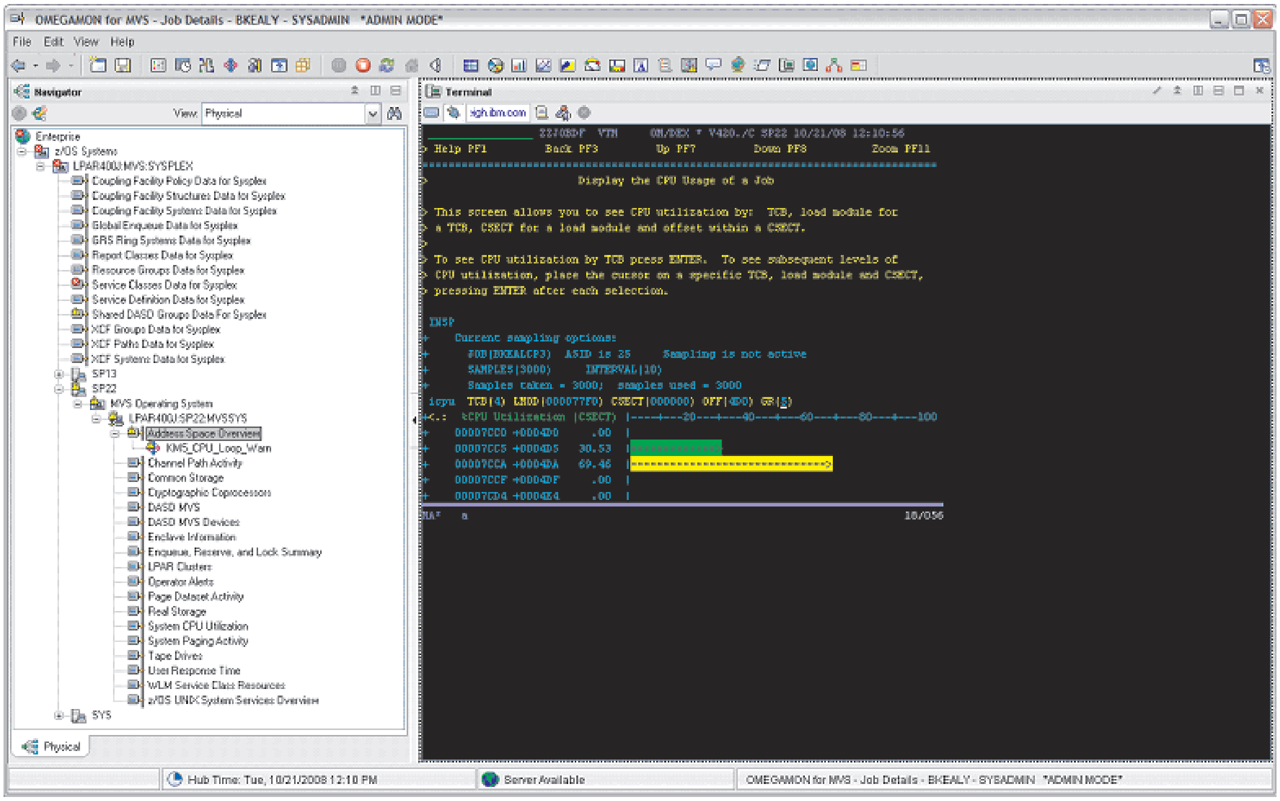
OMEGAMON for MVS - Job Details - BKEALY - SYSADMIN *ADMIN MODE*
File Edit View Help
Navigator
View: Physical
Enterprise
z/OS Systems
LPAR400J:MVS:SYSPLEX
Coupling Facility Policy Data for Sysplex
Coupling Facility Structures Data for Sysplex
Coupling Facility Systems Data for Sysplex
Global Enqueue Data for Sysplex
GRS Ring Systems Data for Sysplex
Report Classes Data for Sysplex
Resource Groups Data for Sysplex
Service Classes Data for Sysplex
Service Definition Data for Sysplex
Shared DASD Groups Data For Sysplex
XCF Groups Data for Sysplex
XCF Paths Data for Sysplex
XCF Systems Data for Sysplex
SP13
SP22
MVS Operating System
LPAR400J:SP22:MVSSYS
Address Space Overview
KM5_CPU_Loop_Warn
Channel Path Activity
Common Storage
Cryptographic Coprocessors
DASD MVS
DASD MVS Devices
Enclave Information
Enqueue, Reserve, and Lock Summary
LPAR Clusters
Operator Alerts
Page Dataset Activity
Real Storage
System CPU Utilization
System Paging Activity
Tape Drives
User Response Time
WLM Service Class Resources
z/OS UNIX System Services Overview
SYS
Physical
Terminal
sigh.ibm.com
Hub Time: Tue, 10/21/2008 12:10 PM
Server Available
OMEGAMON for MVS - Job Details - BKEALY - SYSADMIN *ADMIN MODE*

```
                         ZZJOBDF  VTM     OM/DEX * V420./C SP22 10/21/08 12:10:56
> Help PF1          Back PF3          Up PF7          Down PF8          Zoom PF11
===============================================================================
>                        Display the CPU Usage of a Job

> This screen allows you to see CPU utilization by:  TCB, load module for
> a TCB, CSECT for a load module and offset within a CSECT.
>
> To see CPU utilization by TCB press ENTER.  To see subsequent levels of
> CPU utilization, place the cursor on a specific TCB, load module and CSECT,
> pressing ENTER after each selection.

 INSP
+    Current sampling options:
+      JOB(BKEALCP3)  ASID is 25      Sampling is not active
+      SAMPLES(3000)      INTERVAL(10)
+      Samples taken = 3000;  samples used = 3000
 icpu  TCB(4) LMOD(000077F0) CSECT(000000) OFF(4D0) GR(5)
+<.:  %CPU Utilization (CSECT) |----+---20---+---40---+---60---+---80---+---100
+    00007CC0 +0004D0      .00  |
+    00007CC5 +0004D5    30.53  |=============>
+    00007CCA +0004DA    69.46  |==================================>
+    00007CCF +0004DF      .00  |
+    00007CD4 +0004E4      .00  |
MA*    a                                                                18/056
```

Figure 17. Further drill down on INSPECT data

To cancel the job, you navigate to the main menu for OMEGAMON for MVS by using the PF3 key until you arrive at the OMEGAMON main menu (Figure 18 on page 88).

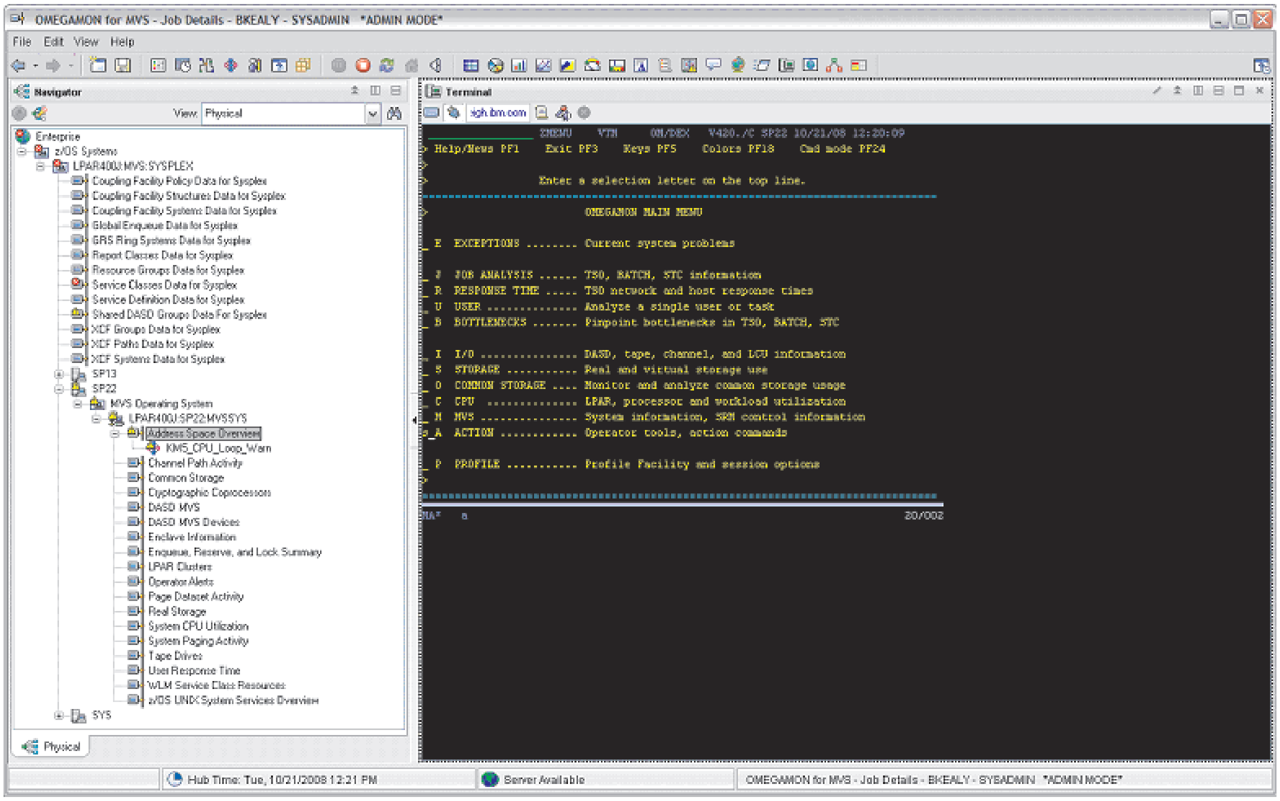
OMEGAMON for MVS - Job Details - BKEALY - SYSADMIN *ADMIN MODE*
File
Edit
View
Help
Navigator
View: Physical
Enterprise
z/OS Systems
LPAR400J:MVS:SYSPLEX
Coupling Facility Policy Data for Sysplex
Coupling Facility Structures Data for Sysplex
Coupling Facility Systems Data for Sysplex
Global Enqueue Data for Sysplex
GRS Ring Systems Data for Sysplex
Report Classes Data for Sysplex
Resource Groups Data for Sysplex
Service Classes Data for Sysplex
Service Definition Data for Sysplex
Shared DASD Groups Data For Sysplex
XCF Groups Data for Sysplex
XCF Paths Data for Sysplex
XCF Systems Data for Sysplex
SP13
SP22
MVS Operating System
LPAR400J:SP22:MVSSYS
Address Space Overview
KM5_CPU_Loop_Warn
Channel Path Activity
Common Storage
Cyptographic Coprocessors
DASD MVS
DASD MVS Devices
Enclave Information
Enqueue, Reserve, and Lock Summary
LPAR Clusters
Operator Alerts
Page Dataset Activity
Real Storage
System CPU Utilization
System Paging Activity
Tape Drives
User Response Time
WLM Service Class Resources
z/OS UNIX System Services Overview
SYS
Physical
Terminal
sgh.ibm.com
Hub Time: Tue, 10/21/2008 12:21 PM
Server Available
OMEGAMON for MVS - Job Details - BKEALY - SYSADMIN *ADMIN MODE*

```
                  ZMENU    VTM      OM/DEX   V420./C SP22 10/21/08 12:20:09
> Help/News PF1    Exit PF3    Keys PF5    Colors PF18    Cmd mode PF24
>
>                 Enter a selection letter on the top line.
===============================================================================
>                       OMEGAMON MAIN MENU

_ E  EXCEPTIONS ........ Current system problems

_ J  JOB ANALYSIS ...... TSO, BATCH, STC information
_ R  RESPONSE TIME ..... TSO network and host response times
_ U  USER .............. Analyze a single user or task
_ B  BOTTLENECKS ....... Pinpoint bottlenecks in TSO, BATCH, STC

_ I  I/O ............... DASD, tape, channel, and LCU information
_ S  STORAGE ........... Real and virtual storage use
_ O  COMMON STORAGE .... Monitor and analyze common storage usage
_ C  CPU  .............. LPAR, processor and workload utilization
_ M  MVS ............... System information, SRM control information
a_A  ACTION ............ Operator tools, action commands

_ P  PROFILE ........... Profile Facility and session options
>
===============================================================================
MA*    a                                                                20/002
```

Figure 18. OMEGAMON main menu

From here you select ACTIONS to get to the Action Commands screen (Figure 19 on page 89).

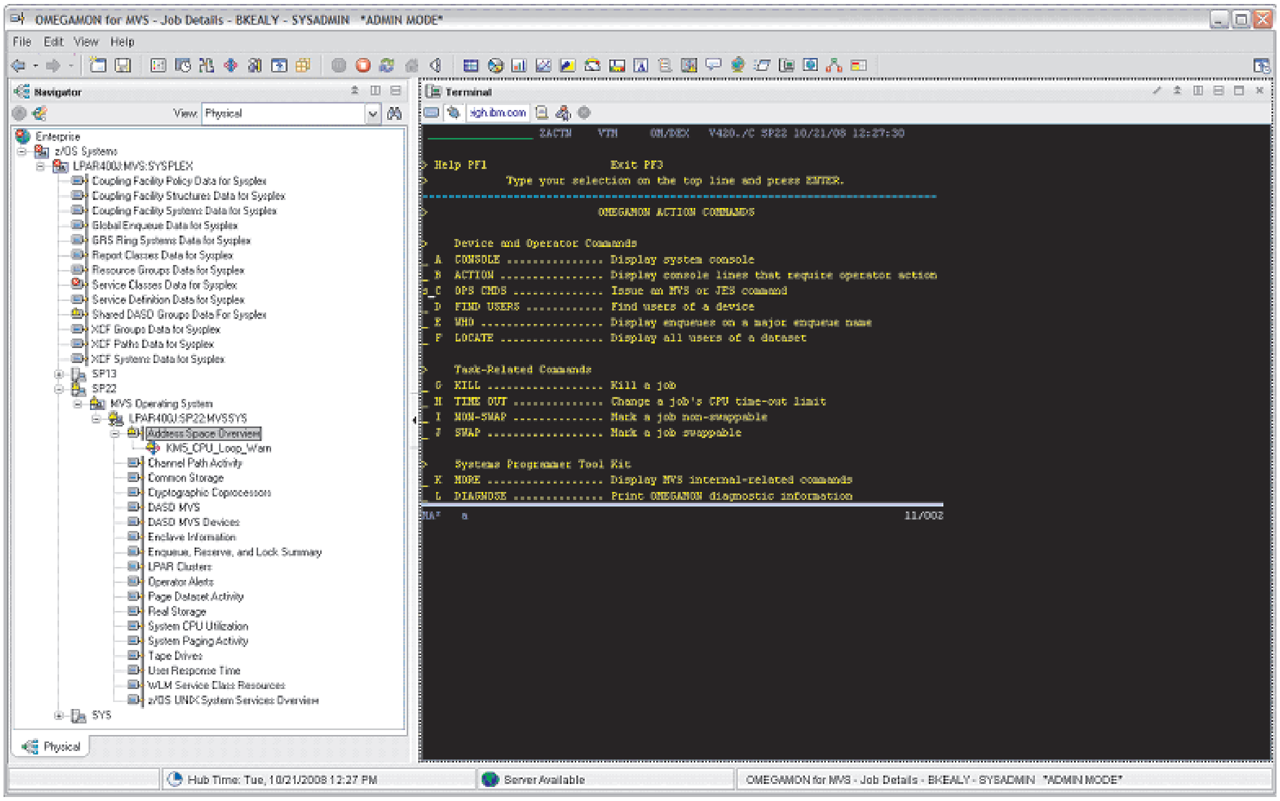
OMEGAMON for MVS - Job Details - BKEALY - SYSADMIN *ADMIN MODE*
File Edit View Help
Navigator
View: Physical
Enterprise
z/OS Systems
LPAR400J:MVS:SYSPLEX
Coupling Facility Policy Data for Sysplex
Coupling Facility Structures Data for Sysplex
Coupling Facility Systems Data for Sysplex
Global Enqueue Data for Sysplex
GRS Ring Systems Data for Sysplex
Report Classes Data for Sysplex
Resource Groups Data for Sysplex
Service Classes Data for Sysplex
Service Definition Data for Sysplex
Shared DASD Groups Data For Sysplex
XCF Groups Data for Sysplex
XCF Paths Data for Sysplex
XCF Systems Data for Sysplex
SP13
SP22
MVS Operating System
LPAR400J:SP22:MVSSYS
Address Space Overview
KM5_CPU_Loop_Warn
Channel Path Activity
Common Storage
Cryptographic Coprocessors
DASD MVS
DASD MVS Devices
Enclave Information
Enqueue, Reserve, and Lock Summary
LPAR Clusters
Operator Alerts
Page Dataset Activity
Real Storage
System CPU Utilization
System Paging Activity
Tape Drives
User Response Time
WLM Service Class Resources
z/OS UNIX System Services Overview
SYS
Physical
Terminal
rtgh.ibm.com
Hub Time: Tue, 10/21/2008 12:27 PM
Server Available
OMEGAMON for MVS - Job Details - BKEALY - SYSADMIN *ADMIN MODE*

```
                 ZACTN     VTM      ON/DEX    V420./C SP22 10/21/08 12:27:30

> Help PF1                    Exit PF3
>             Type your selection on the top line and press ENTER.
===============================================================================
>                             OMEGAMON ACTION COMMANDS

>    Device and Operator Commands
_ A  CONSOLE ................ Display system console
_ B  ACTION ................. Display console lines that require operator action
s_C  OPS CMDS ............... Issue an MVS or JES command
_ D  FIND USERS ............. Find users of a device
_ E  WHO .................... Display enqueues on a major enqueue name
_ F  LOCATE ................. Display all users of a dataset

>    Task-Related Commands
_ G  KILL ................... Kill a job
_ H  TIME OUT ............... Change a job's CPU time-out limit
_ I  NON-SWAP ............... Mark a job non-swappable
_ J  SWAP ................... Mark a job swappable

>    Systems Programmer Tool Kit
_ K  MORE ................... Display MVS internal-related commands
_ L  DIAGNOSE ............... Print OMEGAMON diagnostic information
MA*    a                                                                 11/002
```

Figure 19. OMEGAMON Action Commands screen

And now you can select the OPS CMDS screen space (Figure 20 on page 90).

Note: Most commands, including the z/OS CANCEL command, require the user to be authorized. In OMEGAMON for MVS you can gain authorization by typing in the **/pwd** command in the top line on the screen and then entering the site specific password.).

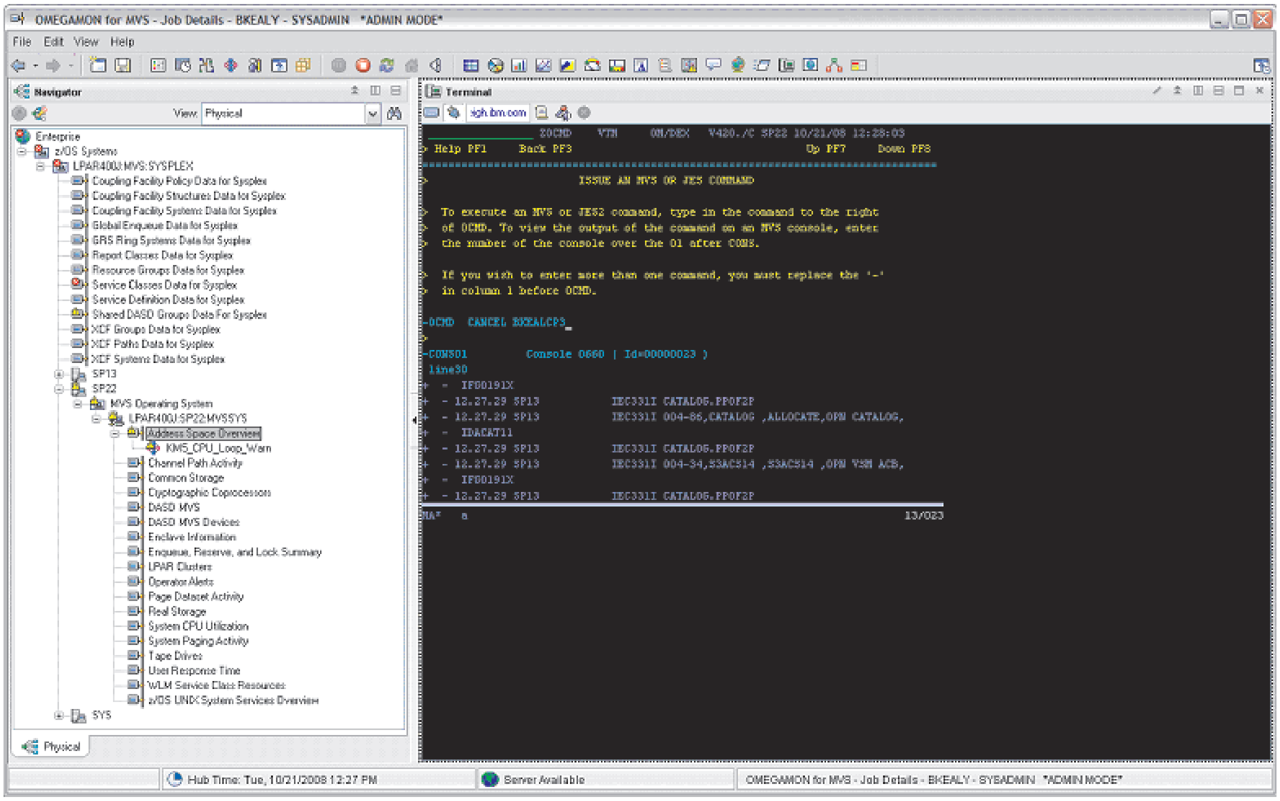

OMEGAMON for MVS - Job Details - BKEALY - SYSADMIN *ADMIN MODE*
File Edit View Help
Navigator
View: Physical
Enterprise
z/OS Systems
LPAR400J:MVS:SYSPLEX
Coupling Facility Policy Data for Sysplex
Coupling Facility Structures Data for Sysplex
Coupling Facility Systems Data for Sysplex
Global Enqueue Data for Sysplex
GRS Ring Systems Data for Sysplex
Report Classes Data for Sysplex
Resource Groups Data for Sysplex
Service Classes Data for Sysplex
Service Definition Data for Sysplex
Shared DASD Groups Data For Sysplex
XCF Groups Data for Sysplex
XCF Paths Data for Sysplex
XCF Systems Data for Sysplex
SP13
SP22
MVS Operating System
LPAR400J:SP22:MVSSYS
Address Space Overview
KM5_CPU_Loop_Warn
Channel Path Activity
Common Storage
Cryptographic Coprocessors
DASD MVS
DASD MVS Devices
Enclave Information
Enqueue, Reserve, and Lock Summary
LPAR Clusters
Operator Alerts
Page Dataset Activity
Real Storage
System CPU Utilization
System Paging Activity
Tape Drives
User Response Time
WLM Service Class Resources
z/OS UNIX System Services Overview
SYS
Physical
Terminal
sgh.ibm.com
Hub Time: Tue, 10/21/2008 12:27 PM
Server Available
OMEGAMON for MVS - Job Details - BKEALY - SYSADMIN *ADMIN MODE*

```
                  ZOCMD     VTM      OM/DEX    V420./C SP22 10/21/08 12:28:03
> Help PF1     Back PF3                                     Up PF7     Down PF8
===============================================================================
>                        ISSUE AN MVS OR JES COMMAND

>  To execute an MVS or JES2 command, type in the command to the right
>  of OCMD. To view the output of the command on an MVS console, enter
>  the number of the console over the 01 after CONS.

>  If you wish to enter more than one command, you must replace the '-'
>  in column 1 before OCMD.

-OCMD  CANCEL BKEALCP3_
>
-CONS01         Console 0660 | Id=00000023 )
 line30
+  -  IFG0191X
+  - 12.27.29 SP13           IEC331I CATALOG.PPOF2P
+  - 12.27.29 SP13           IEC331I 004-86,CATALOG ,ALLOCATE,OPN CATALOG,
+  -  IDACAT11
+  - 12.27.29 SP13           IEC331I CATALOG.PPOF2P
+  - 12.27.29 SP13           IEC331I 004-34,S3AC514 ,S3AC514 ,OPN VSM ACB,
+  -  IFG0191X
+  - 12.27.29 SP13           IEC331I CATALOG.PPOF2P
MA*   a                                                                   13/023
```

Figure 20. OPS CMDS screen space

On the OPS CMDS screen, you enter the z/OS operator command to cancel job BKEALCP3. You can now use the ⇦ Back button on the Tivoli Enterprise Portal interface (upper left above the navigation view) to return to the Address Space Bottleneck Summary view.

The BKEALCP3 job is no longer active (Figure 21 on page 91).

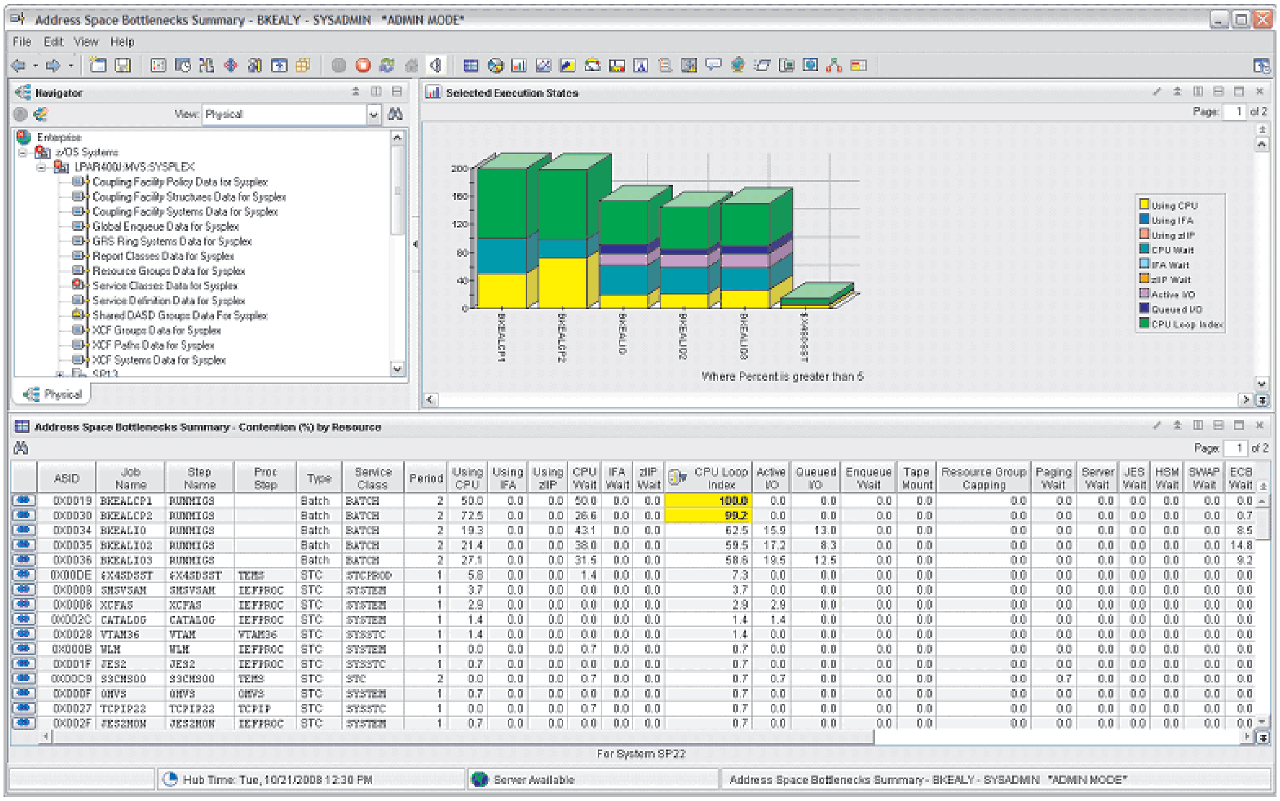
Address Space Bottlenecks Summary - BKEALY - SYSADMIN *ADMIN MODE*
File Edit View Help
Navigator
View: Physical
Enterprise
z/OS Systems
LPAR400J:MVS:SYSPLEX
Coupling Facility Policy Data for Sysplex
Coupling Facility Structures Data for Sysplex
Coupling Facility Systems Data for Sysplex
Global Enqueue Data for Sysplex
GRS Ring Systems Data for Sysplex
Report Classes Data for Sysplex
Resource Groups Data for Sysplex
Service Classes Data for Sysplex
Service Definition Data for Sysplex
Shared DASD Groups Data For Sysplex
XCF Groups Data for Sysplex
XCF Paths Data for Sysplex
XCF Systems Data for Sysplex
SP13
Physical
Selected Execution States
Page: 1 of 2
200
160
120
80
40
0
BKEALCP1
BKEALCP2
BKEALIO
BKEALIO2
BKEALIO3
$X4SDSST
Where Percent is greater than 5
Using CPU
Using IFA
Using zIIP
CPU Wait
IFA Wait
zIIP Wait
Active I/O
Queued I/O
CPU Loop Index

Address Space Bottlenecks Summary - Contention (%) by Resource

Page: 1 of 2

ASID	Job Name	Step Name	Proc Step	Type	Service Class	Period	Using CPU	Using IFA	Using zIIP	CPU Wait	IFA Wait	zIIP Wait	CPU Loop Index	Active I/O	Queued I/O	Enqueue Wait	Tape Mount	Resource Group Capping	Paging Wait	Server Wait	JES Wait	HSM Wait	SWAP Wait	ECB Wait
0X0019	BKEALCP1	RUNMIGS		Batch	BATCH	2	50.0	0.0	0.0	50.0	0.0	0.0	100.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0X0030	BKEALCP2	RUNMIGS		Batch	BATCH	2	72.5	0.0	0.0	26.6	0.0	0.0	99.2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.7
0X0034	BKEALIO	RUNMIGS		Batch	BATCH	2	19.3	0.0	0.0	43.1	0.0	0.0	62.5	15.9	13.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	8.5
0X0035	BKEALIO2	RUNMIGS		Batch	BATCH	2	21.4	0.0	0.0	38.0	0.0	0.0	59.5	17.2	8.3	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	14.8
0X0036	BKEALIO3	RUNMIGS		Batch	BATCH	2	27.1	0.0	0.0	31.5	0.0	0.0	58.6	19.5	12.5	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	9.2
0X00DE	$X4SDSST	$X4SDSST	TEMS	STC	STCPROD	1	5.8	0.0	0.0	1.4	0.0	0.0	7.3	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0X0009	SMSVSAM	SMSVSAM	IEFPROC	STC	SYSTEM	1	3.7	0.0	0.0	0.0	0.0	0.0	3.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0X0006	XCFAS	XCFAS	IEFPROC	STC	SYSTEM	1	2.9	0.0	0.0	0.0	0.0	0.0	2.9	2.9	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0X002C	CATALOG	CATALOG	IEFPROC	STC	SYSTEM	1	1.4	0.0	0.0	0.0	0.0	0.0	1.4	1.4	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0X0028	VTAM36	VTAM	VTAM36	STC	SYSSTC	1	1.4	0.0	0.0	0.0	0.0	0.0	1.4	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0X000B	WLM	WLM	IEFPROC	STC	SYSTEM	1	0.0	0.0	0.0	0.7	0.0	0.0	0.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0X001F	JES2	JES2	IEFPROC	STC	SYSSTC	1	0.7	0.0	0.0	0.0	0.0	0.0	0.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0X00C9	S3CMS00	S3CMS00	TEMS	STC	STC	2	0.0	0.0	0.0	0.7	0.0	0.0	0.7	0.7	0.0	0.0	0.0	0.0	0.7	0.0	0.0	0.0	0.0	0.0
0X000F	OMVS	OMVS	OMVS	STC	SYSTEM	1	0.7	0.0	0.0	0.0	0.0	0.0	0.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0X0027	TCPIP22	TCPIP22	TCPIP	STC	SYSSTC	1	0.0	0.0	0.0	0.7	0.0	0.0	0.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
0X002F	JES2MON	JES2MON	IEFPROC	STC	SYSTEM	1	0.7	0.0	0.0	0.0	0.0	0.0	0.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

For System SP22

Hub Time: Tue, 10/21/2008 12:30 PM

Server Available

Address Space Bottlenecks Summary - BKEALY - SYSADMIN *ADMIN MODE*

Figure 21. Address Space Bottlenecks Summary workspace without BKEALCP3

Part 3. Reference

This sections contains descriptions of the product-provided attributes, workspaces, and situations that are displayed in the Tivoli Enterprise Portal interface. The information in this section is also available as online help in the Tivoli Enterprise Portal when application support for OMEGAMON XE on z/OS is installed on the Tivoli Enterprise Portal Server.

Chapter 10, “Attributes,” on page 95 contains descriptions of all the OMEGAMON XE on z/OS attributes. The attribute groups are divided into system-level and sysplex-level, and the attribute descriptions are organized alphabetically by attribute group within those two categories.

Chapter 11, “Workspaces,” on page 219 contains descriptions of the product-provided workspaces. This section also describes the organization of the sysplex- and system-level workspaces, and the attribute group or groups displayed in each workspace.

Chapter 12, “Situations,” on page 279 contains descriptions of the product-provided situations.

Chapter 10. Attributes

Attributes are characteristics or properties of the logical and physical objects monitored by OMEGAMON XE on z/OS (for example, the amount of allocated virtual storage). Related attributes are organized into groups (which are referred to as *tables* in the context of historical data collection).

These attributes are used to define the *queries* that collect the information displayed in tables and charts in the OMEGAMON XE on z/OS workspaces and to create *situations* that trigger alerts and automated actions in response to specified conditions.

Attribute groups are divided into two types, system and sysplex, corresponding to the two levels of OMEGAMON XE on z/OS nodes in the Tivoli Enterprise Portal Navigator.

Related topics: "Sysplex attribute groups" on page 97, "System attribute groups" on page 135, "Prerequisites."

Prerequisites

Some attributes display data only if specific conditions are met:

Data is available for	Only if
4 Hour MSUs attribute in the System CPU Utilization attributes group	A defined capacity is used as a basis for pricing and the z/OS system is *not* running as a guest on z/VM.
Channel Path attributes	The Resource Measurement Facility (RMF) has been started.
Common Storage attributes	The Common Storage Area Analyzer (CSA Analyzer) is started. **Note:** The CSA Analyzer is shipped and installed with OMEGAMON XE on z/OS. It is configured as part of the configuration of the OMEGAMON II for MVS component and is started as a separate started task.
Coupling facility and cross coupling facility (XCF) data collected by the Resources Management Facility (RMF) Distributed Data Server (DDS)	• You are running z/OS V1.8 or higher with the following RMF components activated: – RMF Control Task (RMF)–one instance on each system – RMF Monitor III Gatherer (RMFGAT)–one instance on each system – RMF Distributed Data Server (GPMSERVE)–one instance per sysplex • You have enabled RMF data collection as described in *OMEGAMON XE on z/OS: Planning and Configuration Guide*.
Cryptographic attributes	At least one IBM cryptographic coprocessor is installed and configured and the KM5EXIT3 exit is installed in the Integrated Cryptographic Service Facility (ICSF). **Note:** The KM5EXIT3 exit is shipped and installed with OMEGAMON XE on z/OS. See *IBM Tivoli OMEGAMON XE on z/OS: Planning and Configuration Guide* for more information.
DASD MVS workspace and DASD MVS Devices attributes	The Resource Measurement Facility (RMF) has been started.
Dynamically-bound PAV aliases	The system you are monitoring is running z/OS V1.8 or above.
GRS Ring Systems attributes	The global resource serialization (GRS) complex is in ring mode. (If the complex is in star mode, only the name, status, and ring acceleration of each system are available.)
HiperDispatch Management and HiperDispatch Logical Processors attributes	HiperDispatch Management mode is On.

Integrated Facility for Applications (IFA) on CP resource times at the address space and service class period level	Either • z/Series Application Assist Processors are configured on the systems, or • Java applications are started using a switch (-Xifa:force)
LPAR cluster attributes	The z/OS system is not running as a guest on z/VM.
Model Permanent Capacity ID and Rating and Model Temporary ID and Rating	System hardware is z10 or above.
Promoted Percent	The z/OS V1.9 Workload Manager blocked workload capability is enabled.
Sysplex DASD attributes (Sysplex DASD Device, Sysplex DASD Group, Sysplex DASD)	A DASD filter situation is enabled.
Suspend lock and spin lock data	• You are running z/OS V1.10 or higher with the following RMF components activated: – RMF Control Task (RMF)–one instance on each system – RMF Monitor III Gatherer (RMFGAT)–one instance on each system – RMF Distributed Data Server (GPMSERVE)–one instance per sysplex • You have enabled RMF data collection as described in *OMEGAMON XE on z/OS: Planning and Configuration Guide*. • Lock data collection is enabled on RMF.
User Response Time attributes	The End to End (ETE) Response Time collector is started. **Note:** The ETE response time collector is shipped and installed with the OMEGAMON XE on z/OS product. It is configured as part of the configuration of the OMEGAMON II for MVS component and is started as a separate started task.
zFS attributes	zFS is specified as the file system on the monitored system (that is, FILESYSTYPE TYPE(ZFS) is specified in SYS1.PARMLIB(BPXPRM*xx*).) **Note:** OMEGAMON XE on z/OS uses an address space name of ZFS, unless the parameter KM3KZFSASNM=*xxxxxxxx* (where *xxxxxxxx* is the started task (STC) name of the zFS address space) has been added to the &*rhilev*.&*rte*.RKANPARU(KDSENV).
z/OS UNIX System Services attributes	The address space where the OMEGAMON XE on z/OS product is running has SUPER USER authority. This level of authority is equivalent to root (UID=0).

Attributes and workspaces

Most workspaces display data for a single attribute group. For a tabular representation of the relationships between the predefined workspaces and the attribute groups, see “Attribute groups used by the system-level predefined workspaces” on page 226 or “Attribute groups used by sysplex-level predefined workspaces” on page 221.

Attributes and queries

Chart and table views use queries to request attribute values from the OMEGAMON XE on z/OS agent for display. You can modify predefined views by editing the queries on which they are based, or design custom views by creating your own queries that collect data for just those attributes you specify.

For information on creating custom queries, see Custom Queries Overview and related topics in the Tivoli Enterprise Portal online help.

Attributes and situations

You can use the OMEGAMON XE on z/OS attributes to create your own situations to monitor or troubleshoot the performance of your systems, analyze their status and alert you to problems.

For information on creating situations, see Situations and Events Overview and related topics in the Tivoli Enterprise Portal online help. For information on the predefined situations available with this product, see Chapter 12, “Situations,” on page 279.

Sysplex attribute groups

Sysplex attribute groups provide detailed information about sysplex-level components such as Workload Manager, coupling facilities, cross-system coupling facilities (XCF), global enqueues and shared DASD.

CF Clients attributes

The CF Clients attributes group monitors user connections to coupling facility structures.

Address Space Address space name. Valid format is a simple text string of 1 through 8 characters; for example, JES2.

AllowAlter Specifies whether or not this connection allows structure alter to be initiated for the structure. Possible values are Yes and No.

AllowRebuild Specifies whether or not this connection allows user-managed structure rebuild to be initiated for the structure. Possible values are Yes and No.

Allow System Managed Duplexing Specifies whether or not this connection allows system-managed processes to be intiated for the structure. Possible values are Yes and No.

Allow User Managed Duplexing Specifies whether or not this connection allows user-managed structure rebuild with duplexing to be initiated for the structure. Possible values are Yes, No, and Not available.

ASID Hexadecimal address space ID.

CF Name Name of the coupling facility. Valid format is a simple text string from 1 through 8 characters; for example, COUPLER1.

Connection Name Name of the connection between the address space and the CF structure. Valid format is a simple text string of 1 through 16 characters; for example, JES2_SYSA.

Connection Problem Flag Flag indicating whether or not this connection is in a problem state. A Connection Status of FAILEDPERS or FAILING is considered a problem. Valid values are 1 (problem) or 0 (no problem).

Connection Status Status of the current coupling facility connection. Valid values are:

Active	This connection is currently active.
FailedPers	This connection has failed but is being maintained because the connector specified connection disposition of “keep”. It is expected that the failed connector will attempt to reconnect.
Failing	The connection is attempting to fail but some connected users of the structure have not yet responded to the fail/disconnect event.
Disconnecting	The user is attempting to disconnect from the structure but some connected users of the structure have not yet responded to the fail/disconnect event.

ConnectedKeep	Connected with CONDISP=KEEP (that is, a persistent connection).
ConnectedRebuild	Connected to both structures during structure rebuild.

Note: Typically, you do not see any of the following statuses for a structure. They exist for very brief periods of time. If one of these statuses persists for more than five minutes, it would indicate a severe problem with either the structure or the monitoring facility.

RebuildActive	Structure rebuild process is in progress
RebuildStopped	A structure rebuild process has been stopped usually by operator command.
RebuildOldStr	This is the original structure in a rebuild process.
RebuildNewStr	This is the new structure in a rebuild process.
RebuildQuiesce	A structure rebuild has been initiated. Connections need to stop usage of the structure and confirm. This phase will be complete when all connections have issued IXLEERSP for the Rebuild Quiesce event.
RebuildComplete	A structure rebuild is in progress. Connections can connect to new structure. This phase will be complete when all connections have issued IXLREBLD REQUEST=COMPLETE.
RebuildCleanup	A structure rebuild is in progress. Connections have completed their part of the process and final cleanup is in progress. This phase will be complete when all connections have issued IXLEERSP for the Rebuild Cleanup event.
RebuildStartOperator	A rebuild process was started by the operator.
RebuildStartConnector	A rebuild process was started by one of the connecting applications.
RebuildLostOldFacility	A rebuild process was stopped because connectivity was lost to the coupling facility containing the original structure.
RebuildFailed	A rebuild process failed.
RebuildStoppedOperator	A rebuild process was stopped by operator command.
RebuildStoppedConnector	A rebuild process was stopped by one of the connected users.
RebuildLostNewFacility	A rebuild process was stopped because connectivity was lost to the coupling facility containing the new structure.
RebuildOldStrFailed	A rebuild process stopped because the original structure failed.

Managed System A sysplex in your enterprise that is being monitored by OMEGAMON XE on z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format *plexname*:MVS:SYSPLEX.

Structure Name Name of the coupling facility structure. Valid format is a simple text string of 1 through 16 characters; for example, IXCSTR1.

Structure Type Type of structure. Valid values are Cache, List, or Lock.

Suspend Specifies if this connection can tolerate suspension of work units during system-managed processing for a structure. Possible values are Yes, No, and Not available.

System Name Name of the coupling facility system; in this case, the name of the z/OS system where the address space is located. Valid format is a simple text string of 1 through 8 characters; for example, SYSA.

CF Path attributes

The CF Path attributes monitors channel path activity between z/OS systems and coupling facilities.

CF Name Name of the coupling facility system to which this z/OS system is connected. Valid format is a simple text string of from 1 through 8 characters; for example, COUPLER1.

Contention Percent Percentage of requests started on this path that were delayed due to path busy. Valid value is a numeric value in the range of 0.0 through 100.0, to one decimal point precision. An example is 10.2.

I/Os Per Second I/O rate for the channel path. Valid value is a numeric value in the range of 0 through 2147483647, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.2.

Managed System A sysplex in your enterprise that is being monitored by Tivoli OMEGAMON XE on z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

Status Current status of the path. Valid values are:

- NotOperational–The path, although defined, cannot be activated.
- Inactive–The path is currently offline but could be brought into service.
- Active–The path is currently available for service.

Sub Channel Path ID Hexadecimal channel path ID.

System Name Name of the z/OS image connected to the facility. Valid format is a simple text string of 1 through 8 characters; for example, SYSA.

CF Policy attributes

The CF Policy attribute group monitors resource consumption of address spaces with service classes.

Date Time Activate Date and time that the coupling facility policy was activated. Valid format is mm/dd/yy hh:mm:ss.

Policy Name Coupling facility policy name. Valid format is a simple text string of 1 through 8 characters; for example, CFRMPOL1.

Managed System A sysplex in your enterprise that is being monitored by Tivoli OMEGAMON XE for z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

Reformat Required Indicates whether a reformat of the Coupling Facility Resource Manager is required. Valid values are Yes or No.

CF Structures attributes

The CF Structures attribute group to create situations that monitor the request activity from a z/OS image to a coupling facility structure.

Asynchronous Requests Per Minute Rate of asynchronous operations for this structure during the last interval. An asynchronous request is a request where the application does not wait for the request, but waits to be notified of the completion of the request either by an ECB (event control block) post or an entry into an exit routine. An asynchronous request is issued when the request is expected to take significant time to complete. Valid value is numeric and is a real number, with a maximum of 4 digits, to one decimal point precision; for example, 839.2.

`AUTOALT`

AutoAlter Specifies the status of system-initiated alters for this structure. Possible values are:

Yes	Specifies that system-initiated alters (automatic alter) are allowed for this structure.
No	Specifies that system-initiated alters (automatic alter) are not allowed for this structure.

CF Name Name of the coupling facility where the structure is located. Valid format is a simple text string of 1 through 8 characters; for example, COUPLER1.

CF Preferences Indicates the names in order of preference where the structure would like to be allocated. Valid format is a simple text string of 1 through 35 characters.

Data Area Element Number of data area elements in the cache structure. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 880.

Data Element Size Size in bytes of data elements in the structure. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 2048.

Directory Element Ratio Directory-to-element target ratio for cache structures. Valid format is a numeric ratio from 1 through 10 characters; for example, 224:880.

Directory Entry Count Number of directory entries in the cache structure. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 224.

Dump Contention Count of operations that were queued for dump serialization. Valid value is a numeric value in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 0.

Dump Table Size Size in 4K blocks of dump table space used for this structure. Valid value is a numeric value in the range of 0 through 2147483647, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 0.

DUPLEX

Duplex Indicates the status of duplexing rebuild for this structure. Possible values are:

Disabled	Neither user-managed nor system-managed duplexing rebuild can be started for this structure.
Allowed	Application can initiate its own user-managed or system-managed duplexing rebuild. z/OS will make no attempt to maintain duplexing status.
Enabled	The system will initiate and attempt to maintain a user-managed or system-managed duplexing rebuild for this structure.

Element Count Number of data elements in a cache or queues in a list. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 880.

Entry Element Ratio Entry-to-element target ratio for list structures. Valid format is a numeric ratio from 1 through 10 characters; for example 193:167.

False Lock Table Entry Contention Count Number of lock requests that encountered contention due to storage constraints. If this number is high, more list table entries should be allocated. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 4.

First Castout Class Name of the first castout class found in the cache structure. Valid value is an integer in the range of 0 through 32767, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. Castout classes are integers (0, 1, 2, etc.). This attribute lists the first one assigned to this structure, if any. An example is 0.

In Use List Set Entry Count Number of list set entries currently in use in the structure. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 19.

Last Castout Class Name of the last castout class found in the cache structure. Castout classes are integers (0, 1, 2, etc.). This attribute lists the last one assigned to this structure, if any. An example is 0.

List Set Element Count Number of elements in the list structure currently in use. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 24.

Lock Table Entry Contention Count Number of lock requests that encountered contention on a lock table entry. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 0.

Lock Table Entry Count Number of lock table entries in the list structure. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 0.

Managed System A sysplex in your enterprise that is being monitored by Tivoli OMEGAMON XE for z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

Maximum Data List Entry Size Maximum number of data elements per list entry. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 10.

Maximum List Set Elements Maximum number of data elements in the list structure. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 167.

Maximum List Set Entry Count Maximum number of list set entries in the structure. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 193.

Maximum Structure Size Maximum size of the structure. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. The value is expressed in 4K pages. An example is 256.

Maximum Users Maximum number of users that may connect to the structure. Valid value is an integer in the range of 0 through 32767, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 64.

Minimum Structure Size Minimum size of the structure. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. The value is expressed in 4K pages for attribute purposes. The number is scaled to K or Megapages. An example is 128.

Percent CF Storage Size Percentage of the total coupling facility allocated to the structure. Valid value is a numeric value in the range 0.0 through 100.0, to one decimal point precision.

Percent Converted Percentage of total requests for this structure in the last interval that were converted from synchronous to asynchronous requests. Valid value is a numeric value in the range 0.0 through 100.0, to one decimal point precision.

Percent Queued Requests Percentage of total requests for this structure in the last interval that were queued. Valid value is a numeric value in the range 0.0 through 100.0, to one decimal point precision.

Percent Total Requests Asynch Percentage of the total requests that were asynchronous operations for this structure during this interval period. Valid value is a percentage that is expressed to one decimal place.

Percent Total Requests Synch The percentage of the total requests that were synchronous operations for this structure during this interval period. Valid value is a percentage that is expressed to one decimal place.

Problem Users Number of users that are connected to the structure in an exception state. The address space is not in an exception state, but the connection is. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 2.

Rebuild Percent Percentage of failed weighted connections at which rebuild should be automatically started by the system. Valid values are numeric values in the range 0.0 through 100.0, to one decimal point precision,. An example is 35.5.

Storage Classes Highest storage class used by the structure. Storage classes are in the range 0 to this value. Valid value is an integer in the range 0 through 255.

Storage Size Storage size currently allocated to the structure. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. The value is expressed in 4K pages for attribute purposes. 128 pages = 512.0 K bytes. An example is 128.

Structure Exclusions Indicates the structure names in order of preference that the structure would like to avoid having allocated in the coupling facility system. Valid format is a simple text string of 1 through 67 characters.

Structure Name Name of the structure. Valid format is a simple text string of 1 through 16 characters; for example, IXCSTR1.

Structure Status Current status of the structure. Valid values are:

- **ActiveDuplexPrimary**
- **ActiveDuplexAlternate**
- **ActiveInUse**
- **ActivePersistent**
- **Defined**
- **PolicyChange**
- **Transitional**
- **Failed**

- **HoldingConnectivity**
- **HoldingDumpPend**
- **Maintmode**
- **RebuildActive**
- **RebuildStopped**
- **RebuildOldStr**
- **RebuildNewStr**
- **RebuildQuiesce**
- **RebuildComplete**
- **RebuildCleanup**
- **RebuildStartOperator**
- **RebuildStartConnector**
- **RebuildLostOldFacility**
- **RebuildFailed**
- **RebuildStoppedOperator**
- **RebuildStoppedConnector**
- **RebuildLostNewFacility**
- **RebuildOldStrFailed**

Structure Type Type of structure (cache, list or lock).

Synch Asynch Convs Per Min Rate of synchronous operations that were converted to asynchronous requests for this structure during a one minute period. Valid value is a count that is expressed to one decimal place.

Synchronous Requests Per Minute Rate of synchronous operations for this structure during the last interval. A synchronous request is one that is expected to complete so quickly that the requesting application would rather wait for completion than do other work. Valid value is a number in the range of 0 through 2147483647, to one decimal point precision, including the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 9.2.

Total Queued Requests Count of all operations that were queued in the last interval. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 1389.

Total Request Cnt This Interval Total number of synchronous and asynchronous operations for this structure during this interval period. Valid value is an integer in the range of 0 through 2147483647, to one decimal point precision, including the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 9.2.

Total Users Total number of users connected to the structure. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 8.

Utilized Storage Size Percentage of the structure storage actually in use. Value is a numeric value in the range of 0.0 through 100.0, to one decimal point precision.

CF Structure to MVS attributes

The CF Structure To MVS System attribute group monitors the request activity from a z/OS image to a coupling facility structure.

Average Asynchronous Request Time Average time (in microseconds) that an asynchronous request takes to complete. An asynchronous request is a request where the requesting application does not remain idle while waiting for the request, but waits to be notified of the completion of the request either by an ECB post or an entry into an exit routine. An asynchronous request is issued when the request is expected to take significant time to complete. Valid value is a numeric value in the range of 0 through 2147483647, to one decimal point precision. For example, a number such as 1234.0 represents 1,234.0 microseconds.

Average Queued Request Time Average time that a request was queued to this z/OS image. Valid value is a numeric value in the range of 0 through 2147483647, to one decimal point precision. For example, a number such as 1234.0 represents 1,234.0 microseconds.

Average Synchronous Request Time Average time (in microseconds) that a synchronous request takes to complete. A synchronous request is one that is expected to complete so quickly that the requesting application would rather wait for completion than do other work. Valid value is a numeric value in the range of 0 through 2147483647, to one decimal point precision. For example, a number such as 1234.0 represents 1,234.0 microseconds.

CF Name Name of the coupling facility. Valid format is a simple text string from 1 through 8 characters; for example, COUPLER1.

Managed System A sysplex in your enterprise that is being monitored by OMEGAMON XE for z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

Request Rates Number of requests per second from this system to this structure. Valid value is a numeric value in the range of 0 through 2147483647, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 12.8.

Requests Converted Total number of requests from this system to this structure over the last observation period that were issued as synchronous and were converted to asynchronous. Valid value is a numeric value in the range of 0 through 2147483647, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.

Structure Name Name of the structure that has connected users in the associated z/OS image. Valid format is a simple text string from 1 through 16 characters; for example, IXSTR1.

System Name Name of the z/OS image connected to the Sysplex. Valid format is a simple text string from 1 through 16 characters; for example, SYSA.

Total Requests Number of requests from this system to this structure over the last observation interval. Valid value is a numeric value in the range of 0 through 2147483647, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 182.0.

User Count Number of users on this image connected to the structure. Valid value is numeric in the range of 0 through 2147483647, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 4.0.

CF Systems attributes

The CF Systems attribute group monitors coupling facility CPU and storage demands.

Allocated Storage Amount of storage allocated to the Coupling Facility currently in use. Allocated storage is the total number of 4K pages assigned to the coupling facility. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 59328.

Asynchronous Request Per Minute Rate of asynchronous operations for this system during this interval period. Valid value is a numeric value to one decimal point precision.

Average Failed Operation Time Average service time for unsuccessful operations. Valid value is an integer in the range of 0 through 2147483647, to one decimal point precision, and can include the use of the *MAX, *MIN, or *SUM functions. An example is 211.

Average Sync Operation Delay Average waiting time for a subchannel to be available for a synchronous operation. Valid value is a number in the range of 0 through 2147483647, to one decimal point precision, and can include the use of the **MAX, *MIN, or *SUM functions. An example is 38.7.

CF Level Coupling facility architected function level. Valid value is an integer in the range of 0 through 2147483647. An example is 3.

CF Name Name of the coupling facility. Valid format is a simple text string from 1 through 8 characters; for example, COUPLER1.

Connected MVS System Count The number of z/OS images connected to the facility. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 6.

CPU Percent Percentage of CPU utilization by the coupling facility. Valid value is a number in the range of 0.0 through 100.0, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 37.5.

Dump Table in Use Amount (a count of 4K pages) of reserved dump storage in the coupling facility currently holding dumps. This value is not a percentage. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 0.

Dump Table Size Amount of storage in the coupling facility reserved for dump tables. The size is expressed as a 4K page count. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 512.

Free Control Space Amount of control storage in the named coupling facility that is not assigned to structures or dump space, given as a 4K page count. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 4416.

Highest 1 Asynch Request Structure Name of the structure with the highest Asynchronous Requests per Minute count followed by its associated count value. Valid value is a 26 character string.

Highest 2 Asynch Request Structure Name of the structure with the second highest Asynchronous Requests per Minute count followed by its associated count value. Valid value is a 26 character string.

Highest 3 Asynch Request Structure Name of the structure with the third highest Asynchronous Requests per Minute count followed by its associated count value. Valid value is a 26 character string.

Highest 1 Synch Request Structure Name of the structure with the highest Synchronous Requests per Minute count followed by its associated count value. Valid value is a 26 character string.

Highest 2 Synch Request Structure Name of the structure with the second highest Synchronous Requests per Minute count followed by its associated count value. Valid value is a 26 character string.

Highest 3 Synch Request Structure Name of the structure with the third highest Synchronous Requests per Minute followed by its associated count value.

I/Os Per Second Sustained I/O rate (I/O requests per second) by the coupling facility. Valid value is a number in the range of 0 through 2147483647, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 27.6.

Managed System A sysplex in your enterprise that is being monitored by OMEGAMON XE for z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

Maximum Sub Channels Count of subchannels that, if they were available, could be used for operations. Valid value is an integer in the range 0 to 2 billion. Values are rarely above 10. An example is 4.

Percent Converted Percentage of the total requests that were converted from synchronous operations to asynchronous operations for this system during this interval period. Valid value is a percentage with one decimal point precision.

Percent of Total Requests Asynch Percentage of the total requests that were asynchronous operations for this system during this interval period. Valid value is a percentage to one decimal point precision.

Percent of Total Requests Synch Percentage of the total requests that were synchronous operations for this system during this interval period. Valid value is a percentage to one decimal point precision.

Percent Utilized Storage Percentage of storage currently being used. Valid value is a numeric value in the range of 0 through 100, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 63.5.

Status Current status of the facility. Valid values are:

Okay	This coupling facility is functioning properly.
Failed	This coupling facility has failed.
Reconcile	The coupling facility to CFRM policy reconcile process is in progress. When this bit is on, connections to structures in this coupling facility using the IXLCONN macro are not permitted.
Maintmode	This coupling facility is in maintenance mode.
PolChg	Policy change pending which will delete this coupling facility from the CFRM active policy when all allocated structures are gone from this coupling facility.

Structure Count In Policy Number of structures in the coupling facility that are contained in the current Coupling Facility Resource Management policy. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 31.

Structure Count Out Policy Number of structures in the coupling facility that cannot be added to the current coupling facility resource management policy. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 0.

Sub Channels Allocated Count of subchannels that this facility has been allocated. Valid value is an integer in the range 0 to 2 billion. Values are rarely above 10. An example is 4.

Sub Channels In Use Count of subchannels available for use. Valid value is an integer in the range 0 to 2 billion. Values are rarely above 10. An example is 4.

Synchronous Requests Per Minute Rate of synchronous operations for this system during this interval period. Valid value is a number with one decimal point precision.

Synch to Asynch Conversions Per Min Rate of synchronous operations that were converted to asynchronous requests for this system during a one minute period. Valid value is an integer with one decimal point precision.

Total Control Space Total amount of storage in the coupling facility that can be used as control storage. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 59238.

Total Request Count This Interval Total number of synchronous and asynchronous operations for this system during this interval period.

Volatility Status 1-byte flag representing the volatility state of storage in the LPAR housing the coupling facility. Valid values are:

Volatile	Storage in the LPAR housing the coupling facility is volatile.
Not Volatile	Storage in the LPAR housing the coupling facility is not volatile.

DASD Device Collection Filtering attributes

The attributes in this group are used to filter collection for DASD devices. When you create a situation for DASD device collection filtering, the program limits the amount of DASD device data sent to the CMS to the devices that meet the criteria of the situation. For an illustration of how to use DASD device collection filtering, see the chapter on “Monitoring Shared DASD” in *IBM Tivoli OMEGAMON XE on z/OS: User's Guide*.

Address Device address for the system in hexadecimal.This attribute is not recommended for use in selecting devices. DASD volumes may be reorganized over device addresses from time to time.

Average Response Time Average I/O response time, in milliseconds, of the device as seen by the system. The valid format is numeric values, including the use of the *AVG, *MAX, *MIN, or *SUM functions in the range of 0 to 2147483647, one decimal point precision. Use this attribute to select devices that might have problems.

Cache Status Cache status of the device for the system. Valid values include:

- Active
- Inactive
- PendingActive
- PendingInactive
- TimedOut

Connect Time Average connect time, in milliseconds, for a typical I/O. The valid format is numeric values, including the use of the *AVG, *MAX, *MIN, or *SUM functions in the range of 0 to 2147483647 with one decimal point precision.

Control Unit Type Control unit type. The valid values include:

- Unknown
- 2105
- 3880-23
- 3990-1
- 3990-2
- 3990-3

- 3990-6

Device Contention Index Device contention index value. Device contention is an indication of I/O delay for this group caused by inter system contention for the device. The larger the number the more severe is the contention. The valid format is numeric values, including the use of the *AVG, *MAX, *MIN, or *SUM functions. In the range of 0 to 2147483647 with two decimal point precision; for example, 1.35.

Device Type The device type. Valid format is 3380 or 3390.

Disconnect Time Average disconnect time, in milliseconds, for a typical I/O. Disconnect time is the time that the I/O searches for the data that has been requested. It includes moving the device to the requested cylinder and track (SEEK + SET SECTOR) waiting for the record to rotate under the head (LATENCY) possible rotational delay as the device waits to reconnect to the channel (RPS delay). The valid format is numeric values, including the use of the *AVG, *MAX, *MIN, or *SUM functions. In the range of 0 to 2147483647 with one decimal point precision; for example, 11.3.

Group Name Name of the DASD group. The valid format is simple text strings of 1 through 30 characters; for example, SGDB2X.

I/O Rate Average I/O rate for the device. The valid format is numeric values, including the use of the *AVG, *MAX, *MIN, or *SUM functions. In the range of 0 to 2147483647 with one decimal point precision; for example, 12.8.

Indexed VTOC Status of indexed VTOC as seen from the system. The valid values include:

- DataUnavailable
- Enabled
- Disabled
- NotIndexed

On heavily used volumes with many data sets, an index to a volume table of contents may improve I/O performance.

IOSQ Average IOS queue time, in milliseconds, for a typical I/O. The valid format is numeric values, including the use of the *AVG, *MAX, *MIN, or *SUM functions. In the range of 0 to 2147483647 with one decimal point precision; for example, 15.5.

Mount Status Mount status of the device for the system. The valid values include:

- Public
- Private
- Storage

Pending Time Average pending time, in milliseconds, for a typical I/O. Although shared DASD often causes high pending time, it can be caused by channel, control unit, or head of string being busy. Use this attribute for filtering. The valid format is numeric values, including the use of the *AVG, *MAX, *MIN, or *SUM functions. In the rang e of 0 to 2147483647 with one decimal point precision; for example, 5.5.

Reconfiguration Status Dynamic reconfiguration status of device for the system. The valid values include:

- Static
- Dynamic
- InstallationStatic
- PinStatusUnknownStatic
- PinStatusUnknownDynamic

- PinStatusUnknownInstallationStatic
- PinnedStatic
- PinnedDynamic
- PinnedInstallationStatic

System Name of the system that shares access to this device. The valid format is simple text strings of 1 through 32 characters; for example, SYSA.

True Percent Busy Device true busy value aggregated over all systems in the Sysplex. The valid format is numeric values, including the use of the *AVG, *MAX, *MIN, or *SUM functions. In the range of 0 to 100 with one decimal point precision; for example, 25.5.

Vary Status Vary status of device for the system. The valid values include:

- Online
- Offline
- PendingOffline
- Boxed
- NotReady

Volume Serial Number Volume serial number. The valid format is a simple text string of 1 through 6 characters; for example, VOL001. When used in a situation, ensures that certain devices are included for collection. It should be specified as a predicate that is joined by the OR operator to all other predicates. You can also include a group of devices by ending the volume serial number with an asterisk (*). For example, SYS*, would select all devices beginning with the characters SYS*.

Enterprise Global Enqueue attributes

The Enterprise Global Enqueue workspace contains the following columns.

ASID Hexadecimal address space ID.

Enqplex Name The name for the enqplex. Valid format is a simple text string from 1 through 8 characters; for example, SYSA.

Major Name Major name (qname) of the resource. Valid format is a simple text string of 1 through 8 characters; for example, SYSDSN.

Maximum Wait Time Longest time that the task has been waiting for a resource. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 45.

Minor Name Minor name (rname) of the resource. Valid format is a simple text string of 1 through 255 characters; for example, SYS1.PROCLIB.

Owning Address Space Name of the job containing the task that owns the resource. Valid format is a simple text string from 1 through 8 characters; for example, PRODJOB1.

Owning Task Count Number of tasks owning the resource. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 3.

Swapped Swap status of the address space in which the task is executing. Valid values are Swapped or Not Swapped.

Sysplex Name The name of the sysplex. Valid format is a simple text string with 1 through 8 characters; for example, SYSPLEX1.

System Name SYSNAME or SMF ID of the z/OS image where the task is executing. Valid format is a simple text string from 1 through 8 characters; for example, SYSA.

True Owner Distinguishes between ordinary jobs, which acquire resources for their own purposes, and resource integrity managers, which generally acquires resources for purely protective purposes. This column is always YES for owners and Not available for waiters. Use True Owner in a filter to eliminate from reports those owners of a resource which have acquired ownership for protective purposes only.

Type Type of ENQ the task has issued for the resource. Valid values are:

Exclusive	The requester requires that no other processes have access to this resource.
Share	The requester is capable of sharing access to the resource with other processes.

Wait Time Time in seconds that the task has been waiting for the resource. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 45.

Waiting Address Space Name of the address space currently waiting on the resource. Valid format is a simple text string from 1 through 8 characters; for example, PRODJOB2.

Waiting Task Count Number of tasks waiting for the resource. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 1.

Global Enqueues attributes

The Global Enqueues attribute group monitors enqueue conflicts within the sysplex. The Global Enqueues attribute group contains these attributes.

ASID Hexadecimal address space ID.

Major Name Major name (qname) of the resource. Valid format is a simple text string of 1 through 8 characters; for example, SYSDSN.

Managed System A sysplex in your enterprise that is being monitored by OMEGAMON XE for z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

Maximum Wait Time Longest time that the task has been waiting for a resource. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 45.

Minor Name Minor name (rname) of the resource. Valid format is a simple text string of 1 through 255 characters; for example, SYS1.PROCLIB.

Owning Address Space Name of the job containing the task that owns the resource. Valid format is a simple text string from 1 through 8 characters; for example, PRODJOB1.

Owning Task Count Number of tasks owning the resource. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 3.

Swapped Swap status of the address space in which the task is executing. Valid values are:

- Swapped
- Not swapped

System Name SYSNAME or SMF ID of the z/OS image where the task is executing. Valid format is a simple text string from 1 through 8 characters; for example, SYSA.

Type Type of ENQ the task has issued for the resource. Valid values are:

Exclusive	The requester requires that no other processes have access to this resource.
Share	The requester is capable of sharing access to the resource with other processes.

Wait Time Time in seconds that the task has been waiting for the resource. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 45.

Waiting Address Space Name of the address space currently waiting on the resource. Valid format is a simple text string from 1 through 8 characters; for example, PRODJOB2.

Waiting Task Count Number of tasks waiting for the resource. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 1.

GRS attributes

The GRS (global resource serialization) attribute group monitors ring status and ring message delays.

Actual Response Time Actual measured ring response time in milliseconds. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 32.

Current RSA Message Size Current size, in bytes, of the outgoing RSA message. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 144.

Expected Response Time Anticipated (or average) time, in milliseconds, that a requester must wait for access to a global resource. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 78.

GRS Mode The serialization propagation topology. Valid values are: Ring, Star, or None.

Managed System A sysplex in your enterprise that is being monitored by OMEGAMON XE for z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

Maximum RSA Message Size Maximum size of the RSA message in bytes. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 32768.

Resident Milliseconds Minimum amount of time, in milliseconds, the RSA message is delayed in each z/OS image in the GRS complex. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 3.

Ring Acceleration Number of consecutive z/OS images in the GRS complex that must process the RSA message before GRS grants access to a global resource. Valid value is an integer in the range of 0 through 65535, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 2.

Status Status of each system in the global resource serialization complex. Valid values are:

- Inactive
- Active
- Joining
- Restarting
- Quiesced

System Name Name of the z/OS image connected to the facility. Valid format is a simple text string from 1 through 8 characters; for example, SYSA.

Tolerance Interval Maximum amount of time, in milliseconds, GRS waits before it detects an overdue RSA message. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 75000.

Report Classes attributes

The Report Class attribute group monitors workload activity accumulated by report class rather than service class.

Address Space Count Number of address spaces for this report class. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 211.

I/O Rate Number of I/Os per second experienced by the report class during the current interval. Valid value is a numeric value in the range of 0 through 2147483647, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 1.4.

Managed System A sysplex in your enterprise that is being monitored by OMEGAMON XE for z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

Page Rate Number of page faults per second experienced by all address spaces in the report class during the report period. Valid value is a numeric value in the range of 0 through 2147483647, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 2.5.

Report Class Name of the report class. Valid format is a simple text string from 1 through 8 characters; for example, REPORTA.

Velocity Percentage of processor utilization for all address spaces in this report class. Valid value is a numeric value in the range of 0 through 100, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.5.

Resource Groups attributes

The Resource Groups attribute group monitors service class resource limits.

Description Description of the resource group. Valid format is a simple text string from 1 through 32 characters. An example is “Resource group to allow 4 MIPs”.

Group Capacity Utilization Number of unweighted service units that are currently being consumed per second by the resource group. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.

IFA on CP Utilization Service units consumed by work units within a resource group performing work eligible for a zSeries Application Assist (zAAP) processor on a standard processor.

IFA Utilization Service units consumed by work units within a resource group performing work on an Integrated Facility for Applications processor.

Managed System A sysplex in your enterprise that is being monitored by OMEGAMON XE for z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

Maximum Capacity Maximum unweighted number of service units that can be consumed per second by all service classes within the resource group. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 200.

Minimum Capacity Minimum unweighted number of service units that can be consumed per second by all service classes within the resource group. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 50.

Name Name of the resource group. Valid format is a simple text string from 1 through 8 characters; for example, FOURMIP.

Percent Maximum Group Capacity Used Percentage of CPU service units defined as the maximum consumable by the resource group and currently being consumed by all service classes associated with the resource group. Valid value is a numeric value in the range of 0.0 through 100.0, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.5.

Service Class Name of the service class that is in the resource group. Valid format is a simple text string from 1 through 8 characters; for example, STC.

zIIP On CP Utilization Service units consumed by work units within a resource group performing work eligible for a System z9 Integrated Information Processor on a standard processor.

zIIP Utilization Service units consumed by work units within a resource group performing work on a System z9 Integrated Information Processor.

Service Class Address Spaces attributes

The Service Class Address Spaces attribute group monitors resource consumption of address spaces with service classes.

Address Space Name of the address space that is running work in the service class. Valid format is a simple text string from 1 through 8 characters; for example, PLDSOS1.

Address Space Type Indicates the type of address space. The valid values are:

APPC	Advanced Program-to-Program Communication address space
Server	Services one or more transaction service classes
Batch	Batch address space
TSO	TSO address space
STC	Started task address space

ASID Hexadecimal address space ID.

Class Flag Categorizes the service class based on the general types of transactions it manages. The valid values are:

Transaction	Service class manages subsystem transactions such as CICS transaction or IMS transactions.
Address space	Service class handles workloads such as batch jobs and started tasks.
TSO	Service class handles TSO users.

CPU Percent Percentage of CPU consumed by the address space over the current interval. Valid value is a numeric value in the range 0.0 through 100.0, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.5.

CSTOR Frames Number of working set frames in central storage for the address space. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 213.

ESTOR Frames Number of working set frames in expanded storage for the address space. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 0.

IFA Percent Percentage of time the address space executed on an Integrated Facility for Applications (IFA) processor.

IFA On CP Percent Percentage of time the address space executed Integrated Facility for Applications (IFA) eligible work on a standard processor.

I/O Rate Number of I/Os per second experienced by the address space over the current interval. Valid value is a numeric value in the range of 0 through 2147483647, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.2.

Managed System A sysplex in your enterprise that is being monitored by OMEGAMON XE for z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

Page Rate Number of page faults per second experienced by the address space over the reporting period. Valid value is a numeric value in the range of 0 through 2147483647, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 10.2.

Period Service class period in which the address space is currently running. Valid format is numeric in the range 0 through 65535; for example, 1.

Report Class Name of the report class to which this address space is assigned. Valid format is a simple text string from 1 through 8 characters; for example, REPORTA.

Service Class Name of the service class to which the associated address space is assigned. Valid format is a simple text string from 1 through 8 characters; for example, STC.

System Name SMF ID of the system that is associated with the address space. Valid format is a simple text string from 1 through 8 characters; for example, SYSA.

Velocity Speed achieved for the work running in the service class from the address space. Valid value is a numeric value in the range of 0.0 through 100.0, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.5.

zIIP On CP Percent Percentage of time the address space performed work eligible for System z9 Integrated Information Processor resources on standard processors. The maximum value is *n*00%, where n represents the number of processors assigned to the sysplex.

zIIP Percent Percentage of time the address space performed work on System z9 Integrated Information Processor resources. The maximum value is *n*00%, where n represents the number of processors assigned to the sysplex.

Service Class Enqueue Workflow Analysis attributes

The Service Class Enqueue Workflow Analysis attribute group monitors service class delays due to specific enqueue resources.

Level Aggregation at a z/OS image level. Data can be aggregated for an individual z/OS system or for an entire Sysplex. Valid values are:

SYSTEM	To aggregate at a z/OS image level
SYSPLEX	To aggregate over an entire Sysplex

Major Queue Name Resource major name. Valid format is a simple text string from 1 through 8 characters; for example, SYSDSN.

Managed System A sysplex in your enterprise that is being monitored by OMEGAMON XE for z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

Period Period number. Valid value is an integer in the range 0 through 8, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. If the situation is to be at the service class level rather than the period level, use Period = SUM. An example is 1.

Queued Percent Percentage of time queued for a major resource for a service class. This indicates the percentage of time transactions for this service class spent waiting to acquire this enqueue. Valid value is a numeric value in the range 0.0 through 100.0, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.5.

Service Class Service class name. Valid format is a simple text string from 1 through 8 characters; for example, STC.

System Name Name of the z/OS system being reported on. Valid format is a simple text string from 1 through 8 characters; for example, SYSA.

Service Class I/O Workflow Analysis attributes

The Service Class I/O Workflow Analysis attribute group monitors service class delays due to specific device I/O activity.

Device Number Number of the device that the service class is currently using. Valid format is an integer in the range 0 through 65535. For situations, the value must be entered in its decimal equivalent. For example, device X’150’ should be entered as its decimal equivalent 336.

Device Type Type of device that the service class is currently using. Valid format is a simple text string from 1 through 4 characters; for example, DASD.

I/O Active Percent Percentage of active I/O rate for the service class. This is the percentage of the transaction elapsed time spent actively doing I/O on the subject device. Valid value is a numeric value in the range 0.0 through 100.0, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.5.

I/O Wait Percent Percentage of I/O wait for the service class. This is the percentage of the transaction elapsed time spent waiting for I/O on the subject device. Valid value is a numeric value in the range 0.0 through 100.0, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.5.

Level Aggregation at a z/OS image level. Data can be aggregated for an individual z/OS system or for an entire Sysplex. Valid values are:

SYSTEM	To aggregate at a z/OS image level
SYSPLEX	To aggregate over an entire Sysplex

Managed System A sysplex in your enterprise that is being monitored by OMEGAMON XE for z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

Period Service class period number. Valid value is an integer in the range 0 through 8, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. If the situation is to be at the service class level rather than the period level, use Period = SUM. An example is 1.

Service Class Service class name. Valid format is a simple text string from 1 through 8 characters; for example, STC.

System Name Name of the z/OS system being reported on. Valid format is a simple text string from 1 through 8 characters; for example, SYSA.

Volser Volume serial number of the device that the service class is using. Valid format is a simple text string from 1 through 6 characters; for example, LST004.

Service Class Subsystem Workflow Analysis attributes

The Service Class Subsys Workflow Analysis attribute group monitors the percentage of time that transaction-oriented service classes spend in specific execution states.

Active Percent Percentage of samples in active state. Valid value is a numeric value in the range 0.0 through 100.0, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.5.

Continuation In Network Percent Percentage of samples in a state representing transactions for which there are logical continuations somewhere in the network. Valid value is a numeric value in the range 0.0 through 100.0, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.5.

Continuation In Sysplex Percent Percentage of samples in a state representing transactions for which there are logical continuations from another z/OS image in the Sysplex. Valid value is a numeric value in the range 0.0 through 100.0, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.5.

Continuation On Local MVS Percent Percentage of samples in a state representing transactions for which there are logical continuations on a z/OS image. Valid value is a numeric value in the range 0.0 through 100.0, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.5.

Conversation Wait Percent Percentage of samples waiting in a conversation state. Valid value is a numeric value in the range 0.0 through 100.0, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.5.

Distribution Wait Percent Percentage of samples waiting in a distributed request state. Valid value is a numeric value in the range 0.0 through 100.0, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.5.

I/O Wait Percent Percentage of samples waiting in I/O state. This is the percentage of the transaction elapsed time spent waiting for I/O on the subject device. Valid value is a numeric value in the range 0.0 through 100.0, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.5.

Idle Percent Percentage of samples in idle state. This is the percentage of tasks in this class that are waiting for terminal output. Valid value is a numeric value in the range 0.0 through 100.0, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.5.

Level Aggregation at a z/OS image level. Data can be aggregated for an individual z/OS system or for an entire Sysplex. Valid values are:

SYSTEM	To aggregate at a z/OS image level
SYSPLEX	To aggregate over an entire Sysplex

Local Session Wait Percent Percentage of samples waiting for a local session to be established. Valid value is a numeric value in the range 0.0 through 100.0, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.5.

Lock Wait Percent Percentage of samples waiting in lock state. Valid value is a numeric value in the range 0.0 through 100.0, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.5.

Managed System A sysplex in your enterprise that is being monitored by OMEGAMON XE for z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

Network Session Wait Percent Percentage of samples waiting for a session to be established somewhere in the network. Valid value is a numeric value in the range 0.0 through 100.0, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.5.

Other Product Wait Percent Percentage of samples waiting for another product state. Valid value is a numeric value in the range 0.0 through 100.0, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.5

Phase Identifies the transaction states sampled in the phase of transactions defined by RCAEEFLG. Valid values are:

- Begin-to-End
- Execution

Ready Percent Percentage of samples in ready state. Valid value is a numeric value in the range 0.0 through 100.0, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.5.

Service Class Name Name of the service class being reported on. Valid format is a simple text string from 1 through 8 characters; for example, CICSTRN.

Subsystem Type of subsystem where the samples are processed. Valid format is a simple text string from 1 through 4 characters; for example, CICS.

Sysplex Session Wait Percent Percentage of samples waiting for a session to be established somewhere in the Sysplex. Valid value is a numeric value in the range 0.0 through 100.0, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.5.

System Name Name of the z/OS system being reported on. Valid format is a simple text string from 1 through 8 characters; for example, SYSA.

Timer Wait Percent Percentage of samples waiting in timer state. Valid value is a numeric value in the range 0.0 through 100.0, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.5.

Unidentified Resource Wait Percent Percentage of samples waiting for an unidentified resource. Valid value is a numeric value in the range 0.0 through 100.0, to one decimal point precision, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.5.

Service Definition attributes

The Service Definition attribute group displays Workload Manager (WLM) service definition and service policy activation dates.

Definition Install Date Date that the service definition was installed.

Definition Install Time Time that the service definition was installed.

Managed System A sysplex in your enterprise that is being monitored by OMEGAMON XE for z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

Policy Activate Date Date and time the service policy was activated.

Policy Activate System System where the service policy was activated. Valid format is a simple text string from 1 through 8 characters; for example, SYSA.

Policy Activate Time Time that the service policy was activated. Valid format is a time in the format hhmmssttt, where:

- hh - hour
- mm - minutes
- ss - seconds
- ttt - thousandths of seconds

Policy Activate User User that activated the service policy. Valid format is a simple text string from 1 through 8 characters; for example, JSMITH.

Service Definition Name Name of the service definition. Valid format is a simple text string from 1 through 8 characters; for example, PLEXADEF.

Service Policy Name of the service policy. Valid format is a simple text string from 1 through 8 characters; for example, POLPRIME.

Sysplex DASD attributes

The Sysplex DASD attribute group monitors sysplex-wide activity of groups of DASD devices.

Average Device Contention Index Average of the device level contention indexes for devices of the group. Device contention is an indicator of I/O delay for this group, caused by intersystem contention for the device. The larger the number, the more severe is the contention. Valid value is a numeric value in the range 0 through 2147483647, to two decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 1.25.

Average True Percent Busy Average of true busy calculated for devices in the group. True percent busy is a measure of the busy time accumulated for the group by all systems in the Sysplex. Valid value is a numeric value in the range 0.0 through 100.0, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 12.5.

Bypass Filtering This attribute is for information only. You can use this attribute in a situation.

Group Name Name of the DASD group. Valid format is a simple text string from 1 through 30 characters; for example, SGDB2X.

Highest Device Contention Index Highest of the device-level contention indexes for devices of the group. Valid value is a numeric value in the range 0 through 2147483647, to two decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 4.25.

Highest Device Contention Index Volser Volume serial number of the device contributing the largest device contention index. Valid format is a simple text string from 1 through 6 characters; for example, VOL001.

Highest True Percent Busy Largest true busy for a device in the group. Valid value is a numeric value in the range 0.0 through 100.0, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.5.

Highest True Percent Busy Volser Volume serial number of the device contributing the largest true busy. Valid format is a simple text string from 1 through 6 characters; for example, VOL001.

Managed System A sysplex in your enterprise that is being monitored by OMEGAMON XE for z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

Volume Serial Number The volume serial number. The valid format is a simple text string with 1 through 6 characters, for example, VOL001

Sysplex DASD Device attributes

The Sysplex DASD Device attribute group monitors z/OS image-by-image activity for individual DASD devices.

Address Device address for the system in hexadecimal format.

Average Response Time Average I/O response time, in milliseconds, of the device as seen by the system. Valid value is a numeric value in the range 0 through 2147483647, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 55.2.

Bypass Filtering This attribute is for information only. You can use this attribute in a situation.

Cache Status Cache status of the device for the system. Valid values include:

- Active
- Inactive
- PendingActive
- PendingInactive
- TimedOut

Connect Time Average connect time in milliseconds, for a typical I/O. Valid value is a numeric value in the range 0 through 2147483647, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 21.2.

Disconnect Time Average disconnect time, in milliseconds, for a typical I/O. Valid value is a numeric value in the range 0 through 2147483647, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 12.8.

I/O Rate Average I/O rate for the device. Valid value is a numeric value in the range 0 through 2147483647, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 12.8.

Indexed VTOC Status of the indexed volume table of contents (VTOC) as seen from the system. Valid values are:

- DataUnavailable
- Enabled
- Disabled
- NotIndexed

IOSQ Average IOS queue time, in milliseconds, for a typical I/O. Valid value is a numeric value in the range 0 through 2147483647, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 15.5.

Managed System A sysplex in your enterprise that is being monitored by OMEGAMON XE for z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

Mount Status Mount status of the device for the system. Valid values are:

- Public
- Private
- Storage

Pending Time Average pending time, in milliseconds, for a typical I/O. Valid value is a numeric value in the range 0 through 2147483647, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 5.5.

Reconfiguration Status Dynamic reconfiguration status of the device for the system. Valid values are:

- Static
- Dynamic
- InstallationStatic
- PinStatusUnknownStatic
- PinStatusUnknownDynamic
- PinStatusUnknownInstallationStatic
- PinnedStatic

- PinnedDynamic
- PinnedInstallationStatic

System Name of the system that shares access to this device. Valid format is a simple text string from 1 through 32 characters; for example, SYSA.

True Percent Busy Device true busy value aggregated over all systems in the Sysplex. Valid value is a numeric value in the range 0 through 100, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.5.

Vary Status Vary status of the device for the system. Valid values are:
- Online
- Offline
- PendingOffline
- Boxed
- NotReady

Volume Serial Number Volume serial number of the device. Valid format is a simple text string from 1 through 6 characters; for example, VOL001.

Sysplex DASD Group attributes

The Sysplex DASD Group attribute group monitors sysplex-wide activity of individual shared DASD devices.

Bypass Filtering This attribute is for information only. You can use this attribute in a situation.

Control Unit Type Type of control unit. Valid values are:
- Unknown
- 3880-23
- 3990-1
- 3990-2
- 3990-3
- 3990-6

Cumulative I/O Rate Sum of the I/O rate in seconds as seen from each sharing system. Valid value is a numeric value in the range 0 through 2147483647, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 45.1.

Device Contention Index Indication of I/O delay for this group caused by intersystem contention for the device. The larger the number, the more severe is the contention. Valid value is a numeric value in the range 0 through 2147483647, to two decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 1.35.

Device Type Device type. Valid values are:
- **3380**
- **3390**

Group Name Name of the DASD group. Valid format is a simple text string from 1 through 30 characters; for example, SGDB2X.

Managed System A sysplex in your enterprise that is being monitored by OMEGAMON XE for z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

True Percent Busy Device true busy value. Valid value is a numeric value in the range 0 through 100, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.5.

Volume Serial Number Volume serial number of the device. Valid format is a simple text string from 1 through 6 characters; for example, VOL001.

Worst Disconnect Time Worst disconnect time, in milliseconds, as seen from one system. Valid value is a numeric value in the range 0 through 100, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 2.5.

Worst Disconnect Time System System where the worst disconnect time was observed. Valid value is a simple text string from 1 through 32 characters; for example, SYSA.

Worst Response Time Worst response time, in milliseconds, as seen from one system. Valid value is a numeric value in the range 0-2147483647, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 242.5.

Worst Response Time System System where the worst response time was observed. Valid value is a simple text string from 1 through 32 characters; for example, SYSA.

Sysplex WLM Service Class Period attributes

The Sysplex WLM Service Class Period Group attribute group monitors service class period performance relative to goals.

Active I/O Indicates the percent of time transactions in this service class period spent doing active I/O over all devices.

Actual Host Measured result for this service class period, either response time (in milliseconds) or velocity (between 0 and 100), depending on the goal type.

Valid value is a numeric value in the range 0 through 2147483647, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. If you are setting thresholds for response time goals, express the value in milliseconds. For example, show 1.5 seconds as 1500 milliseconds. For velocity goals, the value should be between 0 and 100; for example, 85.3.

Actual Network Average network response time, in milliseconds, for the service class. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. If used in a situation, this value should be entered in milliseconds; for example, 1200.

Actual Total Total actual value, in milliseconds, for the service class period. Actual Total presents the measured result including network delay where applicable (for response time-oriented goals). The average measured network delay is added to the measured host response time. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. If used in a situation, this value should be entered in milliseconds; for example, 2200.

Average Response Time Average host response time of transactions in this service class. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. If used as a threshold in a situation, this value should be entered in milliseconds; for example, 1200.

Class Flag Categorizes the service class based on the general types of transactions that are managed by it. Valid values are:

Transaction	This service class manages subsystem transactions such as CICS or IMS transactions.
Address Space	This service class handles workloads such as batch jobs and started tasks.
TSO	This service class handles TSO users.

Common Page-in Wait Wait time for the PLPA or common page-in. This indicates the percentage of time transactions in this service class spent waiting for page-in delays for common storage pages. Valid value is a numeric value in the range 0 through 2147483647, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 10.3.

CPU Wait Total wait time on the active CPU dispatching queue. This indicates the percentage of time transactions in this service class are waiting for CPU access. Valid value is a numeric value in the range 0 through 2147483647, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 21.8.

Cross Memory Wait Total wait time for a cross-memory request. This indicates the percentage of time transactions in this service class wait for cross memory request delays. Valid value is a numeric value in the range 0 through 100, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 3.4.

Crypto AMP Wait This value indicates the percentage of time workloads assigned to this service class spent waiting for Crypto Async Message services. Valid value is a numeric value in the range 0 through 100, to one decimal point precision.

Crypto AP Wait This value indicates the percentage of time workloads assigned to this service class spent waiting for Crypto Assist processor services. Valid value is a numeric value in the range 0 through 100, to one decimal point precision.

Duration Number of CPU service units the period may use before work is passed to the next period. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 200.

ECB Wait Total time waiting for some unknown event. This indicates the percentage of time transactions in this service class wait for an unspecified event. Valid value is a numeric value in the range 0 through 100, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 30.7.

End Date Time The data collection interval end date and time. This attribute has been provided for historical support purposes only and is currently not being used.

Enqueue Wait Indicates the percentage of time transactions in this service class spent waiting for enqueue delays. Valid value is a numeric value in the range 0 through 100, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 5.7.

Goal Describes the service class period's goal as a complete text string. Valid format is a simple text string from 1 through 30 characters; for example, Pct Resp 95% < 400.0 ms.

Goal Importance Importance level of the goal for the service class period. Goal importance for the service class is set as part of the Workload Manager (WLM) policy. When WLM cannot satisfy all goals, it will try to meet goals for more important service class periods first. Valid values are:

- Highest
- High

- Medium
- Low
- Lowest

Goal Percentile Percentage of work in the service class period that completed within the expected response time. Goal percentile is valid for percentile response time goals only. It is the percentage of transactions that should complete within the goal value. This value is set by the system administrator for this service class when the WLM policy is defined. It is not a measured result. Valid value is an integer in the range 0 through 65535, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 90.

Goal Type Type of goal for the service class period. Valid values are:

SysGoal	z/OS system task goal assigned by z/OS itself.
Discret	Discretionary goal get whatever resources remain after all more important workloads are satisfied.
Velocity	Velocity goals define the acceptable amount of delay for work when work is ready to run. Velocity goals are intended for subsystems which use address spaces or enclaves to represent individual work requests.
AveResp	Average Response time is the average time for transactions in this class to complete.
PctResp	Percentile Response time requires two thresholds, the desired transaction response time and the percentage of all transactions completing for this service class. The goal is met if the specified percentage of transactions complete by or below the desired response time. For example 90% of transactions complete in less than or equal to 2 seconds.

Hiperspace™ Page-in Wait Percentage of time transactions in this service class spent waiting for page-in delays for hiperspace pages. Valid value is a numeric value in the range 0 through 100, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 3.4.

HSM Backup/Migrate Total wait time for HSM backup and/or migrate requests. This value indicates the percentage of time transactions in this service class wait for Hierarchical Storage Manager backup and migrate delays. Valid value is a numeric value in the range 0 through 100, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 1.4.

HSM Other Total wait time for other Hierarchical Storage Manager requests and processing. This value indicates the percentage of time transactions in this service class wait for Hierarchical Storage Manager miscellaneous delays such as recovery, JES3 control, interval locate, read control record, etc. Valid value is a numeric value in the range 0 through 100, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 0.4.

HSM Recall Indicates the percentage of time transactions in this service class spent waiting for Hierarchical Storage Manager recall delays. Valid value is a numeric value in the range 0 through 100, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 2.4.

IFA Wait Percentage of time transactions in this service class spent waiting for Integrated Facility for Applications (IFA) resources across the sysplex.

I/O Wait I/O wait rate for the service class period. This value indicates the percentage of time transactions in this service class wait for I/O over all devices. Valid value is a numeric value in the range 0 through 100, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 12.1.

JES Wait Total wait time for JES. This value indicates the percentage of time transactions in this service class wait for JES delays. Valid value is a numeric value in the range 0 through 100, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 2.4.

Level Aggregation of a z/OS image level. Data can be aggregated for an individual z/OS system or for an entire sysplex. Valid values are:

- SYSTEM - to aggregate at a z/OS image level
- SYSPLEX - to aggregate over an entire Sysplex

MVS Lock Wait time for z/OS locks. This value indicates the percentage of time transactions in this service class wait on the z/OS lock. Valid value is a numeric value in the range 0 through 100, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 1.1.

Managed System A sysplex in your enterprise that is being monitored by OMEGAMON XE for z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

New Other Waits Indicates the percentage of time workloads assigned to this service class spent using or waiting for execution states that fall under under the category of Recently Supported Using/Wait states. Currently, this represents the sum of wait/using samples observed for Crypto and WLM Server states. Valid value is a numeric value in the range 0.0 through 100.0.

Other Waits Total wait time for other reasons. This value indicates the percentage of time transactions in this service class spent waiting for miscellaneous other conditions such as disk mounts, task quiescence, mass storage devices, etc. Valid value is a numeric value in the range 0 through 100, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 3.1.

Performance Index Measure of how the class and period are performing in relation to the goal. Values less than or equal to 1.00 indicate that the goal has been achieved. Values greater than 1.0 indicate that the goal has been missed. The larger the index, the further the service class period from the goal. Valid value is a numeric value in the range 0 through 2147483647, to two decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 0.85.

Period Service class period number. If specifying a situation where the values should represent the service class as a whole, then specify Period = SUM. This value is used when building a situation for workflow thresholds at the service class level. Valid value is an integer in the range 1 through 8, and includes the user of the *AVG, *MAX, *MIN, or *SUM functions. An example is 1.

Private Page-in Wait Wait time for a private area page-in. This value indicates the percentage of time transactions in this service class wait for page-in delays for private area pages. Valid value is a numeric value in the range 0 through 100, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 0.5.

Resource Group Capping Total wait time for CPU due to resource group maximum. This value indicates the percentage of time transactions in this service class wait for resource group capping delays. Valid value is a numeric value in the range 0 through 100, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 0.3.

Service Class Name of the service class being reported on. Valid format is a simple text string from 1 through 8 characters; for example, TSO.

Service Class Name Name of the service class being reported on. Valid format is a simple text string from 1 through 8 characters; for example, TSO.

Start Date Time Contains a 16-byte string that represents the data collection interval start date and time. This attribute has been provided for historical support purposes only and is currently not being used.

Stimer Total percentage of time waiting for the stimer. This value indicates a percentage of time transactions in this service class are waiting for a time limit. Valid value is a numeric value in the range 0 through 100, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 10.5.

Stimer ECB Wait Total wait time for some unknown event and the stimer. This value indicates the percentage of time transactions in this service class are waiting for an unspecified event or a time limit. Valid value is a numeric value in the range 0 through 100, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 65.1.

SWP Not Ready Time spent swapped out and not ready for swap-in. This value is the percentage of time transactions in this service class are waiting in a not-ready swapped state. Valid value is a numeric value in the range 0 through 100, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 0.1.

SWP Ready Total time waiting and ready for swap-in. this value is the percentage of time transactions in this service class are waiting in a read swapable state. Valid value is a numeric value in the range 0 through 100, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 0.5.

System Name Name of the z/OS system to which the data in this row applies. Valid format is a simple text string from 1 through 8 characters; for example, SYSA.

Tape Mount Total percentage of time waiting for tape mounts. This value indicates the percentage of time transactions in this service class are waiting for tapes to be mounted. Valid value is a numeric value in the range 0 through 100, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 3.6.

Transaction Rate Number of transactions per second currently being processed in the service class period. Valid value is a numeric value in the range 0 through 2147483647, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 105.5.

Using CPU Total percentage of time transactions in this service class spent using the CPU. Valid value is a numeric value in the range 0 through 100, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 2.3.

Using Crypto AMP This value indicates the percentage of time workloads assigned to this service class spent using Crypto Async Message processor services. Valid value is a numeric value in the range 0 through 100, to one decimal point precision.

Using Crypto AP This value indicates the percentage of time workloads assigned to this service class spent using Crypto Assist processor services. Valid value is a numeric value in the range 0 through 100, to one decimal point precision.

Using IFA Percentage of time transactions in this service class were actively using Integrated Facility for Applications (IFA) resources across the sysplex.

Using IFA on CP Percentage of time transactions in the service class were actively performing work eligible for Integrated Facility for Applications (IFA) resources on standard processors across the sysplex.

Using zIIP Percentage of time transactions in the service class were actively using System z9 Integrated Information Processor (zIIP) resources. The maximum value is *n*00%, where *n* represents the number of processors assigned to the sysplex.

Using zIIP on CP Percentage of time transactions in the service class were actively performing work eligible for System z9 Integrated Information Processors (zIIPs) on standard processors across the sysplex. The maximum value is *n*00%, where n represents the number of processors assigned to the sysplex.

Velocity Velocity of the service class period. Valid value is a numeric value in the range 0 through 100, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 65.2.

VIO Wait Wait time for virtual I/O processing. This value indicates the percentage of time transactions in this service class wait for virtual I/O delays. Valid value is a numeric value in the range 0 through 100, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 0.5.

WLM Server Paging This value indicates the percentage of time workloads assigned to this service class spent waiting for a WLM Server address space due to paging delay. This attribute represents a roll-up of the Server paging delay counts (listed below):

- Private area paging delay samples
- VIO Space paging delay samples
- Hiperspace paging delay samples

Valid value is a numeric value in the range 0 through 100, to one decimal point precision.

WLM Server MPL This value indicates the percentage of time workloads assigned to this service class spent waiting for a WLM Server address space due to Multi-programming Level delay. Valid value is a numeric value in the range 0 through 100, to one decimal point precision.

WLM Server Swapin This value indicates the percentage of time workloads assigned to this service class spent waiting for a WLM Server address space due to swap-in delay. Valid value is a numeric value in the range 0 through 100, to one decimal point precision.

Workload Name of the workload that contains the service class. Valid format is a simple text string from 1 through 8 characters; for example, BATCH.

Worst Performance Index Performance index value scaled to hundredths, calculated for this service class on an individual z/OS image. Valid value is a numeric value in the range 0 through 2147483647, to two decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 2.75.

WTOR Wait Wait time for a reply to operator message. This value is the percentage of time transactions in this service class wait for virtual I/O delays. Valid value is a numeric value in the range 0 through 100, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 2.2.

zIIP Wait Percentage of time transactions in this service class spent waiting for System z9 Integrated Information Processor (zIIP) resources across the sysplex. The maximum value is *n*00%, where *n* represents the number of processors assigned to the sysplex.

XCF Group attributes

The XCF Group attributes provides connected user counts and problem counts for XCF groups.

Group Name Name of the group to which the member belongs. Valid format is a simple text string from 1 through 8 characters; for example, SYSJES.

Managed System A sysplex in your enterprise that is being monitored by OMEGAMON XE for z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

Member Count Number of members in the group. Valid value is a numeric value in the range 0 through 2147483647, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 4.

Problem Count Number of members in the group reporting a problem status. Valid value is a numeric value in the range 0 through 2147483647, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 1.

XCF Members attributes

The XCF Members attribute group monitors the status and activity of the address spaces of XCF group members.

Group Name Name of the group to which the member belongs. Valid format is a simple text string from 1 through 8 characters; for example, SYSJES.

Job Name String that describes the address space job name of the task using this XCF member name. Valid format is a simple text string from 1 through 8 characters, starting with an alphabetic character and using only special characters as described in the z/OS JCL manual; for example, JESXCF.

Managed System A sysplex in your enterprise that is being monitored by OMEGAMON XE for z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

Member Name Name of the member itself. Valid format is a simple text string from 1 through 16 characters; for example, SYSA.

Signals Received Number of signals received by the member. Valid value is a numeric value in the range 0 through 2147483647, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 4.

Signals Sent Number of signals sent by the member. Valid value is a numeric value in the range 0 through 2147483647, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 2.

Status Status of the member. Valid values are:

Missing	Information about the status of this member is missing.
Created	A member in a created state is known to XCF, but cannot use XCF services. It can become an active user through the IXCJOIN macro.
ActiveState	An active member is known to XCF and can use XCF services.
Quiesced	A member in the quiesced state is disassociated from XCF services, but XCF still maintains a record of the member's existence. The IXCQUIES macro places a member in the quiesced state.
Failed	A member in the failed state is one whose associated task, job step task, address space, or system terminated before the member was explicitly deactivated by invoking an XCF service. When a member is in the failed state, other members can infer that the member did not have an opportunity to clean up its own resources, and another member should take recovery action.
MonitorRemoved	Monitoring has been removed for this member.
SysTermination	XCF system containing member is terminating.

Status Checking Interval An optional parameter that specifies the status checking interval in seconds. It determines the length of time that can elapse with no change to the status field before the user status routine is scheduled. Valid value is numeric and is expressed in seconds. It must be a multiple of 100 scaled to two decimal point precision. For example, a value of 3000 seconds is shown as 3000.00.

System Name Name of the z/OS image where the member is executing. Valid format is a simple text string from 1 through 8 characters; for example, SYSA.

XCF Paths attributes

The XCF Paths attribute group monitors the status and activity of XCF signaling paths between z/OS images.

Destination Device Represents the destination of the path. Destination devices may be coupling facility structures as well as CTC devices. In this case, the designation can be either ListStructure (which implies that the System To value is unknown), or CFList. It is also possible that the destination device is unknown at the time the data is collected. Valid format is a simple text string of 1 through 4 characters. To designate a structure in a situation for Destination Device:

- Use *CFST*to represent ListStructure, or
- Use *List*to represent CFList

Examples: CTC Device - 07F6. Structure - List

Managed System A sysplex in your enterprise that is being monitored by OMEGAMON XE for z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

Message Limit Maximum number of messages allowed to be sent by the signaling path. This is the maximum number of messages concurrently in progress. Valid value is an integer in the range 0 through 2147483647 and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 750.

Origin Device Device from which the path originates. Origin devices may be Coupling Facility structures as well as CTC devices. In this case, the designation can be either ListStructure (which implies that the System To value is unknown), or CFList. Valid format is a simple text string of 1 through 4 characters. To designate a structure in a situation for Origin Device:

- Use *CFST*to represent ListStructure, or
- Use*List*to represent CFList

Examples: CTC Device - 07F6. Structure - List

Restart Count Number of times XCF restarted the signaling path. Valid value is an integer in the range 0 through 2147483647 and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 1.

Retry Limit Retry limit for the signaling path. Valid value is an integer in the range 0 through 2147483647 and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 60.

Retry Percent Percentage of the retry limit reached by the signaling path. Valid value is an integer in the range 0.0 through 100.0, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 25.5.

Signals Pending Transfer Number of messages waiting to be sent through the signaling path. Valid value is an integer in the range 0 through 2147483647 and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 2.

Signals Received Number of messages received through the signaling path. Valid value is an integer in the range 0 through 2147483647 and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 1582.

Signals Sent Number of messages sent through the signaling path. Valid value is an integer in the range 0 through 2147483647 and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 2176.

Status Status of the signaling path. Valid values are:

Starting	Validating and initializing path hardware.
Restarting	Making path ready (again) for use.
Working	Path is capable of being used.
Stopped	Stopping use, in the process of being removed from service.
WaitingForComp	Waiting for completion of initial protocol used to establish communication link.
NotOperational	Not operational. Path defined to XCF but not usable until hardware and/or definition problems are resolved.
Failed	Stop failed, intervention required.
Rebuilding	In the process of being rebuilt.
Quiescing	Quiescing the use of.
Quiesced	Use was quiesced.

Storage In Use Amount of storage that is in use by the path. Valid value is an integer in the range 0 through 2147483647 and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. The value is expressed as a count of 1K buffers. For example, 280 KB is expressed as Storage_In_Use=280.

System From Name of the z/OS image from which the signaling path originates. Valid format is a simple text string from 1 through 8 characters; for example, SYSA.

System Name Path as seen from this system. Each path has a sending and receiving system. This attribute indicates which of these systems is providing the data observation.

System To Name of the z/OS image that is the destination of the signaling path. Valid format is a simple text string from 1 through 8 characters; for example, SYSB.

Times Buffer Unavailable Number of times message receipt was delayed because no buffer was available. Valid value is an integer in the range 0 through 2147483647 and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 1.

Times Path Busy Number of sends issued when the path was busy. Valid value is an integer in the range 0 through 2147483647 and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 1.

Transport Class Name of a transport class associated with the signaling path. Valid format is a simple text string from 1 through 8 characters; for example, DEFAULT.

XCF System attributes

The XCF System attribute group monitors the status of z/OS images in a sysplex as they relate to the XCF Communication feature.

Managed System A sysplex in your enterprise that is being monitored by OMEGAMON XE for z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

Monitoring Interval Length of time, in seconds, it takes XCF to detect a failure in the Sysplex. Valid value is a numeric value in the range 0 through 2147483647, with two decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. When used in situations, specify seconds with precision to hundredths of seconds; for example, 90.00.

Operator Interval Length of time, in seconds, it takes XCF to notify the operator of a failure in the Sysplex. Valid value is a numeric value in the range 0 through 2147483647, with two decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. When used in situations, specify seconds with precision to hundredths of seconds; for example, 90.00.

Status Status of the z/OS image. Valid values are:

Active	This z/OS system is active in the sysplex.
BeingRemoved	This z/OS system is partitioning out of the sysplex.
Missing	Status-update missing.
Local	Single system, no coupling dataset, sysplex.
Cleanup	System has completed Sysplex Partitioning but is still in the process of Cleanup.
Unknown	The status is not known.

System Level z/OS release level running on this system image. Valid format is a simple text string from 1 through 16 characters; for example, 02.05.00.

System Name Name of the z/OS image in the sysplex. Valid format is a simple text string from 1 through 8 characters; for example, SYSA.

XCF System Statistics attributes

The XCF System Statistics attributes monitors XCF signals (messages) sent between z/OS systems in the sysplex and transport class buffer utilization. The situations you create using these attributes let you identify where changes to transport class definitions can be made to optimize XCF performance and resource utilization.

Buffer Length Buffer length for the indicated transport class. Valid value is an integer.

Default Sequence A sequence number for sorting rows of this group in a meaningful order. This attribute is provided for information only. It cannot be used to create situations or queries, or to set thresholds and filters.

Managed System A sysplex in your enterprise that is being monitored by OMEGAMON XE for z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

Percent Degraded Percent of messages sent where the message length was larger than the transport class buffer length and caused degraded performance. Valid value is a percentage with one decimal point precision.

Percent Fit Percent of messages sent where the message length fit the transport class buffer length. Valid value is a percentage with one decimal point precision.

Percent Larger Percent of messages sent where the message length was larger than the transport class buffer length. Valid value is a percentage with one decimal point precision.

Percent Smaller Percent of messages sent where the message length was smaller than the transport class buffer length. Valid value is a percentage with one decimal point precision.

Signals Received Number of messages received. Valid value is an integer.

Signals Sent Number of messages sent. Valid value is an integer.

System From Name of the z/OS System where the messages originated. Valid value is an 8 character string.

System Name of the z/OS System on which data was collected. Valid value is an 8 character string.

System To Name of the z/OS System that is the destination of the messages. Valid value is an 8 character string.

Transport Class Transport class used for messages sent. Transport classes are used to segregate message traffic according to message size or performance needs, and may be used to dedicate signalling paths to particular workloads. Valid value is an 8 character string.

Times Buffer Unavailable Number of times a ’no buffer’ condition occurred for the indicated transport class. Valid value is an integer.

Times Path Unavailable Number of times a ’no path’ condition occurred for the indicated transport class. Valid value is an integer.

Sysplex Coupling Facility Columns

The Sysplexes Coupling Facility view of the Sysplex Overview workspace contains the following columns.

CF Name Name of the coupling facility. Valid format is a simple text string from 1 through 8 characters; for example, COUPLER1.

Managed System A sysplex in your enterprise that is being monitored by OMEGAMON XE for z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

Problem Users Number of users that are connected to the structure in an exception state. The address space is not in an exception state, but the connection is. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 2.

Proxy Host The managed system name of the Tivoli Enterprise Monitoring Server which is currently hosting the sysplex proxy. The valid value is a simple text string with a maximum of 32 characters.

Sysplex Name The name of the sysplex. Valid format is a simple text string with 1 through 8 characters; for example, SYSPLXA.

Structure Name Name of the structure. Valid format is a simple text string of 1 through 16 characters; for example, IXCSTR1.

Total Users Total number of users connected to the structure. Valid value is an integer in the range of 0 through 2147483647 and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 8.

Sysplexes Shared DASD Columns

The Sysplexes Shared DASD view contains the following columns.

Average Device Contention Index Average of the device level contention indexes for devices of the group. Device contention is an indicator of I/O delay for this group, caused by intersystem contention for the device. The larger the number, the more severe is the contention. Valid value is a numeric value in the range 0 through 2147483647, to two decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 1.25.

Group Name Name of the DASD group. Valid format is a simple text string from 1 through 30 characters; for example, SGDB2X.

Highest Device Contention Index Highest of the device-level contention indexes for devices of the group. Valid value is a numeric value in the range 0 through 2147483647, to two decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 4.25.

Highest Device Contention Index Volser Volume serial number of the device contributing the largest device contention index. Valid format is a simple text string from 1 through 6 characters; for example, VOL001.

Managed System A sysplex in your enterprise that is being monitored by OMEGAMON XE for z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

Proxy Host The managed system name of the CMS which is currently hosting the Sysplex Proxy CMS. The valid value is a simple text string with a maximum of 32 characters.

Sysplex Name The name of the sysplex. Valid format is a simple text string with 1 through 8 characters; for example, SYSPLXA.

Sysplexes Global Enqueue Columns

The Sysplexes Global Enqueue workspace contains the following columns.

Major Name Major name (qname) of the resource. Valid format is a simple text string of 1 through 8 characters; for example, SYSDSN.

Managed System A sysplex in your enterprise that is being monitored by OMEGAMON XE for z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

Maximum Wait Time Longest time that the task has been waiting for a resource. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 45.

Minor Name Minor name (rname) of the resource. Valid format is a simple text string of 1 through 255 characters; for example, SYS1.PROCLIB.

Proxy Host The managed system name of the CMS which is currently hosting the Sysplex Proxy CMS. The valid value is a simple text string with a maximum of 32 characters.

Sysplex Name The name of the sysplex. Valid format is a simple text string with 1 through 8 characters; for example, SYSPLXA.

Waiting Task Count Number of tasks waiting for the resource. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 1.

Sysplexes GRS Columns

The Sysplexes GRS (Global Resource System) view contains the following columns.

Actual Response Time Actual measured ring response time in milliseconds. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 32.

Expected Response Time Anticipated (or average) time, in milliseconds, that a requester must wait for access to a global resource. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 78.

Managed System A sysplex in your enterprise that is being monitored by OMEGAMON XE for z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

Proxy Host The managed system name of the CMS which is currently hosting the sysplex Proxy CMS. The valid value is a simple text string with a maximum of 32 characters.

Resident Milliseconds Minimum amount of time, in milliseconds, the RSA message is delayed in each z/OS image in the GRS complex. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 3.

Sysplex Name The name of the sysplex. Valid format is a simple text string with 1 through 8 characters; for example, SYSPLXA.

Sysplexes Workloads Columns

The Sysplexes Workloads workspace contains the following columns.

Actual Host Measured result for this service class period. Depending on the goal type, it is either a response time or a velocity. Valid value is a numeric value in the range 0 through 2147483647, to one decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions.

Goal Describes the service class period's goal as a complete text string. Valid format is a simple text string from 1 through 30 characters; for example, Pct Resp 95% < 400.0 ms.

Goal Importance Importance level of the goal for the service class period. Goal importance for the service class is set as part of the WLM policy. When WLM cannot satisfy all goals, it will try to meet goals for more important service class periods first. Valid values are:

- Highest
- High
- Medium
- Low
- Lowest

Managed System A sysplex in your enterprise that is being monitored by OMEGAMON XE for z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

Performance Index Measure of how the class and period are performing in relation to the goal. Values less than or equal to 1.00 indicate that the goal has been achieved. Values greater than 1.0 indicate that the goal has been missed. The larger the index, the further the service class period from the goal. Valid value is a numeric value in the range 0 through 2147483647, to two decimal point precision, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 0.85.

Proxy Host The managed system name of the CMS which is currently hosting the Sysplex Proxy CMS. The valid value is a simple text string with a maximum of 32 characters.

Service Class Name of the service class to which the associated address space is assigned. Valid format is a simple text string from 1 through 8 characters; for example, STC.

Sysplex Name The name of the sysplex. Valid format is a simple text string with 1 through 8 characters; for example, SYSPROD1.

Sysplexes XCF Columns

The Sysplexes XCF workspace contains the following columns. This workspace does not have an associated attribute group.

Group Name Name of the group to which the member belongs. Valid format is a simple text string from 1 through 8 characters; for example, SYSJES.

Managed System A sysplex in your enterprise that is being monitored by OMEGAMON XE for z/OS agents. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:MVS:SYSPLEX.

Member Count Number of members in the group. Valid value is a numeric value in the range 0 through 2147483647, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 4.

Problem Count Number of members in the group reporting a problem status. Valid value is a numeric value in the range 0 through 2147483647, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 1.

Proxy Host The managed system name of the CMS which is currently hosting the sysplex proxy CMS. The valid value is a simple text string with a maximum of 32 characters.

Sysplex Name The name of the sysplex. Valid format is a simple text string with 1 through 8 characters; for example, SYSPLXA.

System attribute groups

System attribute groups provide system-level data on address space CPU, storage and bottleneck usage, channel activity, common storage usage, single image DASD device usage, WLM service class resource usage, and LPAR cluster activity. Three groups of attributes monitor the status and configuration of IBM cryptographic coprocessors installed in zSeries servers.

Address Space Bottleneck attributes

This attribute group provides address space-level wait state (bottleneck) information.

Active I/O The percentage of time the address space has waited for I/O requests that are being processed by the I/O subsystem. Valid range is 0.0 - 100.0.

ASID Hexadecimal address space ID. The address space ID is a unique number assigned to an address space when it is created. It is used by z/OS to differentiate between address spaces.

Common Page-in The percentage of time the address space has waited due to common area (CSA or LPA) paging. Valid range is 0.0 - 100.0.

ADXLPIDX

CPU Loop Index The sum of all CPU, zIIP, zIIP on CP, zAAP, and zAAP on CP using and waiting counts, divided by total sample count. For CPU looping jobs, this value is usually above 98%. Expressed as a percentage to one decimal place.

CPU Wait The percentage of time the address space has waited for access to the CPU. Valid range is 0.0 - 100.0.

Cross Memory Page-in The percentage of time the address space has waited due to paging activity to and from a cross-memory address space. Valid range is 0.0 - 100.0.

Crypto Assist Proc Wait The percentage of time the address space has waited for access to the Cryptographic Assist (AP) Message Processor. Valid range is 0.0 - 100.0.

Crypto Async Msg Proc Wait The percentage of time the address space has waited for access to the Cryptographic Asynchronous Message (CAM) Processor. Valid range is 0.0 - 100.0.

ECB Wait The percentage of time the address space has waited due to an ECB (Event Control Block) wait. An ECB wait is a voluntary wait that an address space goes into when it is waiting for an event or a series of events to occur. Valid range is 0.0 - 100.0.

Enqueue Wait The percentage of time the address space has waited due to enqueues. Valid range is 0.0 - 100.0.

HSM Backup The percentage of time the address space has waited for an HSM dataset backup. Valid range is 0.0 - 100.0.

HSM CDS Read The percentage of time the address space has waited for an HSM control dataset (CDS) read. Valid range is 0.0 - 100.0.

HSM Delete The percentage of time the address space has waited for an HSM dataset delete. Valid range is 0.0 - 100.0.

HSM JES3 C/I Locate The percentage of time the address space has waited for an HSM JES3 C/I Locate. Valid range is 0.0 - 100.0.

HSM Migrate The percentage of time the address space has waited for an HSM dataset migration. Valid range is 0.0 - 100.0.

HSM Recall The percentage of time the address space has waited for an HSM dataset recall from auxiliary storage. Valid range is 0.0 - 100.0.

HSM Recover The percentage of time the address space has waited for an HSM dataset recovery. Valid range is 0.0 - 100.0.

HSM TSO CLIST The percentage of time the address space has waited for an HSM TSO CLIST. Valid range is 0.0 - 100.0.

HSM Wait The percentage of time the address space has waited due to HSM. This is a roll-up of all HSM wait states. Valid range is 0.0 - 100.0.

Hiperspace Page-in The percentage of time the address space has waited due to hiperspace paging. Valid range is 0.0 - 100.0.

IFA Wait The percentage of all observed execution states where the address space was found waiting to use an IFA. Valid range is 0.0 - 100.0.

JES Wait The percentage of time the address space has waited due to JES2. This is a rollup of all JES2 wait states. Valid range is 0.0 - 100.0.

JES2 Job Cancel The percentage of time the address space has waited for a JES2 job cancel. Valid range is 0.0 - 100.0.

JES2 Job Delete The percentage of time the address space has waited for a JES2 job delete. Valid range is 0.0 - 100.0.

JES2 Job Requeue The percentage of time the address space has waited for a JES2 job requeue. Valid range is 0.0 - 100.0.

JES2 Job Status The percentage of time the address space has waited for a JES2 job status request. Valid range is 0.0 - 100.0.

JES2 SYSOUT The percentage of time the address space has waited for JES2 SYSOUT processing. Valid range is 0.0 - 100.0.

Job Name Address space name (job or started task name, or TSO user ID). Valid value is an 8-byte character string.

Managed System A z/OS operating system in your enterprise that is being monitored by a Tivoli OMEGAMON XE on z/OS agent. Valid value is a character string with a maximum length of 32 bytes in the format <plexname>:<smfid>:MVSSYS.

MVS Lock Wait The percentage of time the address space has waited due to a z/OS lock. Locking is a way of reserving access to a resource. Typically, locks are used to limit access to system and address space control blocks, particularly when those control blocks are being manipulated or updated. Valid range is 0.0 - 100.0.

Paging Wait The percentage of time the address space has waited due to paging activity. This is a roll-up of all paging wait states. Valid range is 0.0 - 100.0.

Period Period number. Valid value is a 2-byte integer.

Private Page-in The percentage of time the address space has waited due to private (local) area paging. Valid range is 0.0 - 100.0.

Proc Step Procedure step name. Valid value is an 8-byte character string.

Queued I/O The percentage of time the address space has waited for I/O requests that are queued to the I/O subsystem. Valid range is 0.0 - 100.0.

Report Class Report class name. Valid value is an 8-byte character string.

Resource Group Capping The percentage of time the address space has waited because the system has imposed a cap on the resource group to which this workload belongs. Valid range is 0.0 - 100.0.

Server MPL Delay The percentage of time the address space has waited for a server due to the System Resource Manager (SRM) adjusting the server multiprogramming level (MPL) to correct a resource shortage. Valid range is 0.0 - 100.0.

Server Paging The percentage of time the address space has waited due to server paging. Valid range is 0.0 - 100.0. This is a roll-up of the following server paging wait states:

- Private area paging delay
- VIO space paging delay
- Hiperspace paging delay

Server Swap-In The percentage of time the address space has waited due to server swap-in delay. Valid range is 0.0 - 100.0.

Server Wait The percentage of time the address space has waited for a server. This is a roll-up of all server wait states. Valid range is 0.0 - 100.0.

Service Class Service class name. Valid value is an 8-byte character string.

Shared Pages The percentage of time the address space has waited for shared storage paging. Valid range is 0.0 - 100.0.

Step Name Step name. Valid value is an 8-byte character string.

Stimer ECB Wait The percentage of time the address space has waited due to an ECB wait with a timer event. An ECB wait is a voluntary wait that an address space goes into when it is waiting for an event or a series of events to occur. Valid range is 0.0 - 100.0.

Stimer Wait The percentage of time the address space has waited due to a STIMER event. A STIMER wait is a voluntary wait condition that an address space issues when it expects to be waiting for a long period of time. Valid range is 0.0 - 100.0.

SWAP APPC The percentage of time the address space was swapped out waiting for service from an APPC/z/OS service. Valid range is 0.0 - 100.0.

SWAP Aux Stor The percentage of time the address space has waited for a swap to correct a shortage of auxiliary storage. Valid range is 0.0 - 100.0.

SWAP Central Stor The percentage of time the address space was swapped out to correct a shortage of storage availability. Valid range is 0.0 - 100.0.

SWAP Detected The percentage of time the address space was swapped out due to being in a detected wait state. A wait state is defined as a wait (without the LONG parameter) that exceeds a duration limit set by the System Resource Manager (SRM). Valid range is 0.0 - 100.0.

SWAP Enq Exchange The percentage of time the address space was swapped out to expedite processing of an enqueue by its owner. Valid range is 0.0 - 100.0.

SWAP Exchange The percentage of time the address space was swapped out due to the System Resource Manager (SRM) adjusting the multi-programming level (MPL) to correct a resource shortage. Valid range is 0.0 - 100.0.

SWAP In-Real The percentage of time the address space was swapped out due to the reallocation of real storage frames. Valid range is 0.0 - 100.0.

SWAP Long The percentage of time the address space was swapped out because it was in a long wait state. Valid range is 0.0 - 100.0.

SWAP OMVS Input The percentage of time the address space was swapped out due to an OpenMVS (UNIX System Services) input wait. Valid range is 0.0 - 100.0.

SWAP OMVS Output The percentage of time the address space was swapped out due to an OpenMVS (UNIX System Services) output wait. Valid range is 0.0 - 100.0.

SWAP Out Too Long The percentage of time the address space was swapped out to enable the swap-in of an address space that has been swapped out too long. Valid range is 0.0 - 100.0.

SWAP MPL Delay The percentage of time an address space was swapped out and ready to be swapped in, but the multiprogramming level (MPL) could not be increased to allow the address space to be swapped in. Valid range is 0.0 - 100.0.

SWAP Page-Ins The percentage of time the address space was swapped out waiting for its working set to be swapped in. Valid range is 0.0 - 100.0.

SWAP Real Stor The percentage of time the address space was swapped out waiting to correct a real storage shortage. Valid range is 0.0 - 100.0.

SWAP Request The percentage of time the address space was swapped out due to the requested release of the real storage frames it occupied. Valid range is 0.0 - 100.0.

SWAP Sys Paging The percentage of time the address space was swapped out to improve the system paging rate. Valid range is 0.0 - 100.0.

SWAP Terminal In The percentage of time the address space was swapped out due to a terminal input wait. Valid range is 0.0 - 100.0.

SWAP Terminal Out The percentage of time the address space was swapped out due to a terminal output wait. Valid range is 0.0 - 100.0.

SWAP Transition The percentage of time the address space was swapped out due to a change from swappable to non-swappable status. Valid range is 0.0 - 100.0.

SWAP Unilateral The percentage of time the address space was swapped out due to the number of users within the domain exceeding the target level set by the System Resource Manager (SRM). Valid range is 0.0 - 100.0.

SWAP Wait The percentage of time the address space has waited due to swapping. This is a roll-up of all swap wait states. Valid range is 0.0 - 100.0.

Tape Mount The percentage of time the address space has waited for a tape to be mounted/loaded. Valid range is 0.0 - 100.0.

Type Type of address space. Valid values are Server, Batch, TSO, STC, and APPC.

Using CPU The percentage of all observed execution states where the address space was found using CPU. Valid range is 0.0 - 100.0.

Using Crypto Assist Proc The percentage of time the address space was found using the Cryptographic Assist (AP) Processor. Valid range is 0.0 - 100.0.

Using Crypto Async Msg Proc The percentage of time the address space was found using the Cryptographic Asynchronous Message (CAM) Processor. Valid range is 0.0 - 100.0.

Using IFA The percentage of all observed execution states where the address space was found using an IFA. Valid range is 0.0 - 100.0.

Using IFA on CP The percentage of all observed execution states where the address space was found executing IFA work on a standard CP.

Using zIIP The percentage of all observed execution states where the address space was found using a System z9 Integrated Information Processor.

ADXAASCP

Using zIIP on CP The percentage of all observed execution states where the address space was found executing work eligible for a System z9 Integrated Information Processor work on a standard processor.

VIO Wait The percentage of time that the address space has waited due to virtual I/O paging. Valid range is 0.0 - 100.0.

WTOR Wait The percentage of time that the address space has waited for an operator reply to a console message. Valid range is 0.0 - 100.0.

zIIP Wait The percentage of all observed execution states where the address space was found waiting to use a System z9 Integrated Information Processor.

Address Space Bottlenecks Portrait attributes

This group provides bottleneck analysis details (multi-row) for one address space ID. These attributes are for information only, and should not be used to create queries or situations.

ASID The address space ID in hexadecimal.

Attribute The name of the attribute that is displayed on the tabular report column heading. Only attributes that have a percentage value greater than zero will be displayed in this table. Valid value is character string with a maximum length of 32 bytes.

Enqueue_Count Number of unique enqueue waits. Valid value is a 2-byte integer.

I/O_Dev_Count Number of unique I/O waits. Valid value is a 2-byte integer.

Job Name The name of the address space (either job name, started task name, or TSO user ID). Valid value is character string with a maximum length of 8 bytes.

Link_Values Values passed in the link to DASD Device or Enqueue probes. The contents will be either DASD volume serial numbers (maximum of 20) or enqueue major names (maximum of 15).

Managed System A z/OS operating system in your enterprise that is being monitored by a Tivoli OMEGAMON XE on z/OS agent. Valid value is a character string with a maximum length of 32 bytes in the format <plexname>:<smfid>:MVSSYS.

Percent The percentage of time the address space was in the state defined by the associated attribute name. Valid value in the range 0 - 100.

Resource Device type, volser, and devices address, for I/O wait attributes; enqueue major name for the enqueue wait attribute. Blank for all other attributes.

Address Space ComStor Owned attributes

This attribute group provides information about common service area (CSA), extended common service area (ECSA), system queue area (SQA), and extended system queue area common storage (ESQA) information at the address space level.

Note: The CSA Analyzer must be running for data to be available for these attributes. For information on configuring the CSA Analyzer, see IBM *Tivoli OMEGAMON XE for z/OS Configuration Guide*.

% of Total CSA Percentage of total common service area (CSA) storage in use by this job. Valid value is a number in the range 0 - 100.

% of Total ECSA Percentage of total extended common service area (ECSA) storage in use by this address space. Valid value is a number in the range 0 - 100.

% of Total ESQA Percentage of total extended system queue area (ESQA) storage in use by this address space. Valid value is a number in the range 0 - 100.

% of Total SQA Percentage of total system queue area (SQA) storage in use by this address space. Valid value is a number in the range 0 - 100.

ASID Address space ID in hexadecimal.

CSA in Use Amount of common service area (CSA) storage in use by this address space, in bytes. Valid value is a numeric value in the range 1 - 2147483647.

CSA Orphaned Indicates whether or not the storage is currently owned. Valid values are:

Not available	Address space did not allocate CSA storage
Yes	CSA storage was orphaned by this address space and is currently unowned
No	Address space is active and this CSA storage is owned

ECSA In Use Amount of extended common service area (ECSA) storage in use by this address space.

ECSA Orphaned Indicates whether or not the storage is currently owned. Valid values are:

Not available	Address space did not allocate ECSA storage
Yes	ECSA storage was orphaned by this address space and is currently unowned
No	Address space is active and this ECSA storage is owned

ESQA In Use Amount of extended system queue area (ESQA) storage in use by this address space.

ESQA Orphaned Indicates whether or not the storage is currently owned. Valid values are:

Not available	Address space did not allocate ESQA storage
Yes	ESQA storage was orphaned by this address space and is currently unowned
No	Address space is active and this ESQA storage is owned

Job Name Name of the address space (either job name, started task name, or TSO user ID). Valid value is a string of up to eight characters.

Managed System A z/OS operating system in your enterprise that is being monitored by a Tivoli OMEGAMON XE on z/OS agent. Valid value is a character string with a maximum length of 32 bytes, in the format <plexname>:<smfid>:MVSSYS.

SQA In Use Amount of system queue area (SQA) storage in use by this address space.

SQA Orphaned Indicates whether or not the storage is currently owned. Valid values are:

Not available	Address space did not allocate SQA storage
Yes	SQA storage was orphaned by this address space and is currently unowned
No	Address space is active and this SQA storage is owned

Address Space ComStor Owned Detail attributes

This attribute group provides common storage (CSA, ECSA, SQA, ESQA) detail information for areas owned by the selected address space.

Note: The CSA Analyzer must be running for data to be available for these attributes. For information on configuring the CSA Analyzer, see *IBM Tivoli OMEGAMON XE on z/OS: Planning and Configuration Guide*.

% of Total Percentage of the total common storage area in use by this address space. Valid values in the range 0-100.

Age Time since the area was allocated.

Age Units The units in which the age is reported, either hours, days, or minutes. In the format DDD:HH, HH:MM, or MM:SS.

Area Common storage area. Valid value is one of the following:

- CSA (common service area)
- ECSA (extended common service area)
- SQA (system queue area)
- ESQA (extended system queue area)

ASID Address space ID in hexadecimal. The address space ID is a unique number assigned to an address space when it is created. It is used by z/OS to differentiate between address spaces.

End Address Ending address of the common storage area. Valid value is a 4-byte hexadecimal number.

Fixed Indicates whether the area is from a fixed, non-pageable subpool (Yes, No).

Job Name Address space name (JOB/STC name, TSO User ID). Valid value is a character string with a maximum length of 8 bytes. A job name of *SYSTEM* indicates that the issuer of the request is unknown.

Managed System A z/OS operating system in your enterprise that is being monitored by a Tivoli OMEGAMON XE on z/OS agent. Value is a character string with a maximum length of 32 bytes in the format <plexname>:<smfid>:MVSSYS.

Requestor Return Address The return address of the GETMAIN request, which is the next instruction following the GETMAIN. If the return address cannot be determined (the storage was obtained before the VSM service was active), the address is given as Unknown. Valid value is a 4-byte hexadecimal number.

Size Size of the common storage area. Valid value is a 4-byte integer.

Start Address Starting address of the common storage area. Valid value is a 4-byte hexadecimal number.

Storage Key Unique key assigned to the frame to protect it from unauthorized use.

Subpool Subpool of the common storage space area. For CSA, valid values are subpools 227, 228, 231, and 241. For SQA, valid values are subpools 226, 239, and 245. Although related by common area, these subpools have different characteristics. CSA subpools 227 and 228, for example, are fixed in storage, while subpools 231 and 241 are pageable.

Address Space ComStor Trends attributes

The attributes in this group provide common storage (CSA, ECSA, SQA, ESQA) usage information for the selected address space over time.

Note: The CSA Analyzer must be running for data to be available for these attributes. For information on configuring the CSA Analyzer, see *IBM Tivoli OMEGAMON XE on z/OS: Planning and Configuration Guide* .

% of Total The percentage of the total common storage area in use, to one decimal precision. Valid value is in the range 0.0 to 100.0.

Area The common area, either CSA (common service area), ECSA (extended common service area), SQA (system queue area), or ESQA (extended system queue area).

ASID Address space ID in hexadecimal.

Date The month and day the sample was taken, in the format mm/dd.

In Use Amount of the common storage area in use, in bytes. Valid value is a number in the range 0 through 2147483647.

Job Name The name of the address space (job name, STC name, or TSO user ID). Valid value is a character string with a maximum length of 8 bytes.

Managed System A z/OS system in your enterprise that is being monitored by a Tivoli OMEGAMON XE for z/OS agent. Valid value is a character string with a maximum length of 32 bytes in the format <plexname>:<smfid>:MVSSYS.

Time The hour and minute the sample was taken, in the format hh:mm.

Address Space ComStor Unowned attributes

Reports virtual storage address of unowned common storage areas (CSA, SQA, ECSA, or ESQA), as well as the name and ID of the address space that allocated the storage.

% of Total Percentage of the total common storage area in use by this address space.

Age Time since the area was allocated (see age units).

Age Units The units in which the age is reported, either hours, days, or minutes. In the format DDD:HH, HH:MM, or MM:SS.

Area Common storage area. Valid value is one of the following:

- CSA (common service area)
- ECSA (extended common service area)
- SQA (system queue area)
- ESQA (extended system queue area)

ASID Address Space ID in hexadecimal. The address space ID is a unique number assigned to an address space when it is created. It is used by z/OS to differentiate between address spaces.

End Address Ending address of the common storage area. Valid value is an 4-byte hexadecimal number.

Fixed Indicates whether the area is from a fixed, non-pageable subpool (Yes, No).

In Use Amount of the common storage area in use by this address space.

Job Name Address space name (JOB/STC name, TSO User ID). Valid value is a character string of up to 8 bytes.

Managed System z/OS operating system in your enterprise that Tivoli Enterprise Portal is monitoring using an OMEGAMON XE on z/OS agent. Value is a character string with a maximum length of 32 bytes in the format <plexname>:<smfid>:MVSSYS.

Requestor Return Address The return address of the GETMAIN request, which is the next instruction following the GETMAIN. Valid value is a 4-byte hexadecimal number.

Size Size of the common storage area. Valid value is a 4-byte integer.

Start Address Starting address of the common storage area. Valid value is a 4-byte integer.

Storage Key Unique key assigned to the frame to protect it from unauthorized use.

Subpool Subpool of the common storage space area. For CSA, valid values are subpools 227, 228, 231, and 241. For SQA, valid values are subpools 226, 239, and 245. Although related by common area, these subpools have different characteristics. CSA subpools 227 and 228, for example, are fixed in storage, while subpools 231 and 241 are pageable.

Address Space IA Impact attributes

Use the Address Space IA Impact attributes to determine how various workloads within your system may be interfering with each other. Impact analysis assists you in arriving at quick and accurate solutions to problems caused by workload contention.

ASID Address space ID in hexadecimal. The address space ID is a unique number assigned to an address space when it is created. It is used by z/OS as an index to find the main control block of the address space or to differentiate between address spaces.

CPU Percent The percentage of the total delay to the favored workload that this Impactor contributed during the collection period that resulted from CPU contention. Valid value is a 2-byte integer in the range 0.0 - 100.0.

Enqueue Percent The percentage of the total delay to the favored workload that this Impactor contributed during the collection period that resulted from waiting for an enqueue. Valid value is a 2-byte integer in the range 0.0 - 100.0.

Favored ASID The hexadecimal address space ID of the favored address space.

Favored Job Name The job name of the favored address space. Valid value is a string, with a maximum of 8 characters.

Favored Type The type of the favored address space. The value in this column should be AS, identifying this as an address space. Valid value is a string, with a maximum of 2 characters.

IFA Percent The amount that this impactor contributed towards the delay of the favored address space due to IFA contention during the collection period. Valid range is 0.0 - 100.0.

IFC Percent The amount that this Impactor contributed towards the delay of the favored address space due to IFA on CP contention during the collection period. Valid value is a 4-byte integer in the range is 0.0 - 100.0.

Impact Percent The percentage of the total delay, for any reason, to the favored workload that this Impactor contributed during the collection period. Valid value is a 4-byte integer in the range from 0.0 to 100.0.

Impactor IFA Count The number of samples that this impactor has delayed the favored address space during the collection period for which the delay was due to IFA contention. Valid value is a 4-byte integer in the range 0-2147483647.

Impactor IFC Count The number of samples that this impactor has delayed the favored address space during the collection period for which the delay was due to IFA on CP contention. Valid value is a 4-byte integer in the range 0-2147483647.

Impactor zIIP Count The number of samples this Impactor has delayed the favored address space during the collection period for which the delay was due to contention for System z9 Integrated Information Processor resources.

I/O Percent The percentage of the total delay to the favored address space that this Impactor contributed during the collection period that resulted from I/O contention. Valid value is a 4-byte integer in the range from 0.0 to 100.0.

Job Name The job name of the impactor that is delaying the favored workload. Valid value is an 8-byte character string.

Job Number The JES job ID of the Impactor. Valid value is an 8-byte character string.

Managed System A z/OS operating system in your enterprise that a Tivoli OMEGAMON XE for z/OS agent is monitoring. Valid value is a character string with a maximum length of 32 bytes in the format <plexname>:<smfid>:MVSSYS.

Performance Index Performance index defined for the impactor's service class. Valid value is a 2-byte integer.

Service Class Name of the service class in which the Impactor is running. Valid value is an 8-byte character string.

Total IFA Count The total number of samples that have been taken for the favored address space during the collection period for which the delay was due to IFA contention. Valid value is a 4-byte integer

Total IFC Count Total number of samples that have been taken for the favored address space during the collection period for which the delay was due to IFA on CP contention. Valid value is a 4-byte integer in the range 0-2147483647.

Total zIIP Count Total number of samples taken for the favored address space during the collection period for which the delay was due to contention for System z9 Integrated Information Processor resources.

zIIP Percent The amount that this impactor contributed towards the delay of the favored address space due to contention for System z9 Integrated Information Processor resources during the collection period.

IAM_SUC

zIIP on CP Percent The amount that this impactor contributed towards the delay of the favored address space due to System z9 Integrated Information Processor (zIIP) on CP contention during the collection period. The value can exceed 100%. The maximum value is *n*00% where *n* is the number of processors (CPUs) assigned to the LPAR or image.

Address Space CPU Utilization attributes

The Address Space CPU Utilization Attribute Group provides basic address space identifying information and CPU usage information for all address spaces. It also provides enclave data.

Active Enclave Count The total number of dependent, independent, or unknown enclaves owned by the address space that are currently active. Valid range is 0-2147483647.

ASID The address space ID in hexadecimal. The address space ID is a unique number assigned to an address space when it is created. It is used by z/OS as an index to find the main control bock of the address space or to differentiate between address spaces.

CPU Percent The percentage of CPU utilized by the address space over the current interval. The value can exceed 100%. The maximum value is *n* 00%, where *n* represents the number of processors (CPUs) assigned to the LPAR or image.

Dependent Active Enclave Count The total number of dependent enclaves owned by the address space that are currently active. Valid range is 0-2147483647.

Dependent Enclave CPU Percent Total CPU time for all dependent enclaves owned by the address space, expressed as a percentage. The value can exceed 100%. The maximum value is *n* 00%, where *n* represents the number of processors (CPUs) assigned to the LPAR or image.

Dependent Enclave IFA % The percentage of IFA utilized by dependent enclaves owned by the address space over the interval. The value can exceed 100%. The maximum value is *n*00%, where *n* is the number of processors (IFAs) assigned to the LPAR or image.

Dependent Enclave IFA % On CP The percentage of regular processor time used by dependent enclaves owned by the address space performing work eligible for a zSeries Application Assist Processor. The value can exceed 100%. The maximum value is *n*00%, where *n* is the number of processors (IFAs) assigned to the LPAR or image.

Dependent Enclave zIIP % The percentage of System z9 Integrated Information Processor (zIIP) time consumed by dependent enclaves owned by the address space. The maximum value is *n*00%, where *n* represents the number of processors assigned to the LPAR or image.

Dependent Enclave zIIP % On CP The percentage of standard processor time consumed by dependent enclaves owned by the address space performing work eligible for a System z9 Integrated Information Processor (zIIP). The maximum value is *n*00%, where *n* represents the number of processors assigned to the LPAR or image.

Dependent Inactive Enclave Count The total number of dependent enclaves owned by the address space that are currently inactive. Valid range is 0-2147483647.

I/O Rate The average number of reads and writes (I/Os) per second during the last sample interval, to one decimal place. (By default, the sample interval is 2.3 seconds.)

IFA Percent The percentage of IFA utilized by tasks in the address space over the interval. The value can exceed 100%. The maximum value is *n*00%, where *n* is the number of processors (IFAs) assigned to the LPAR or image.

IFA on CP The percentage of CPU utilized by tasks performing IFA work in the address space over the interval. The value can exceed 100%. The maximum value is *n* 00%, where *n* represents the number of processors (CPUs) assigned to the LPAR or image.

Inactive Enclave Count The total number of dependent, independent, and unknown enclaves owned by the address space that are currently inactive. Valid range is 0-2147483647.

Independent Enclave CPU% CPU time for all independent enclaves owned by the address space, expressed as a percentage. The value can exceed 100%. The maximum value is *n* 00%, where *n* represents the number of processors (CPUs) assigned to the LPAR or image.

Independent Enclave IFA % The percentage of IFA utilized by independent enclaves owned by the address space over the interval. The value can exceed 100%. The maximum value is *n* 00%, where *n* represents the number of processors (CPUs) assigned to the LPAR or image.

Independent Enclave IFA % On CP The percentage of CPU utilized by independent enclaves owned by the address space performing IFA work. The value can exceed 100%. The maximum value is *n*00%, where *n* is the number of processors (IFAs) assigned to the LPAR or image.

Independent Enclave zIIP % The percentage of System z9 Integrated Information Processor (zIIP) time consumed by independent enclaves owned by the address space. The value can exceed 100%. The maximum value is n00%, where n represents the number of processors assigned to the LPAR or image.

Independent Enclave zIIP % On CP The percentage of standard processor time consumed by independent enclaves owned by the address space performing work eligible for a System z9 Integrated Information Processor (zIIP). The maximum value is n00%, where n represents the number of processors assigned to the LPAR or image.

Independent Active Enclave Count The total number of independent enclaves owned by the address space that are currently active. Valid range is 0-2147483647.

Independent Inactive Enclave Count The total number of independent enclaves owned by the address space that are currently inactive. Valid range is 0-2147483647.

JESJOBID The JES job ID of the address space. This column can be blank if the address space is not using JES services. (Note that tasks started using SUB=MSTR generally do not use JES services and therefore do not have a JESJOBID. Address spaces started prior to JES initialization will generally not use JES services. However, a JESJOBID commonly does exist for the Master Scheduler, since it uses a JES service for SYSLOG.) Valid value is a string, with a maximum of eight characters.

Job Additional SRB Service Percent Cumulative additional SRB service percentage, to one decimal place, for the address space since collection started. This is (Job_Additional_SRB_Service_Time x 1000) / Job_Elapsed_Time. Since multiple CPs may be active, this value can exceed 100%.

Job Additional SRB Service Time Cumulative additional SRB service time (ASST) in seconds with two decimal places for the address space since it started. CPU time is accumulated here for this address space's preemptable SRBs and for client related SRBs for which this address space is the client.

Job CPU Percent Cumulative CPU percentage, to one decimal place, for the address space since collection started. This will be (Job_CPU_Time x 1000) / Job_Elapsed_Time. Since multiple CPs may be active, this value can exceed 100.0%.

Job CPU Time Cumulative CPU time in seconds with two decimal places for the address space since collection started. This is a sum of Job TCB Time, Job SRB Time, and Job Additional SRB Service Time. Together, these fields plus Owned Enclave time make up the CPU time that this address space is accountable for.

Job Elapsed Time Elapsed time since collection for this job started, in seconds with two decimal places.

Job Name The name of the job, started task, TSO user, APPC address space, and so on, consuming CPU cycles. Valid value is a string, with a maximum of eight characters.

Job Preemptable Home SRB Service Percent Cumulative preemptable home SRB service percentage to one decimal place for the address space since collection started. This is (Job_Preemptable_Home_SRB_Service _Time x 1000) / Job_Elapsed_Time. Since multiple CPs may be active, this value can exceed 100.0%.

Job Preemptable Home SRB Service Time Cumulative preemptable home SRB service time (PHTM) in seconds with two decimal places for the address space since it started. The CPU time for all types of preemptable SRBs (PSRB, CRSRB, ESRB) executing within this address space as their home space is accumulated here. CPU time is accumulated here for SRBs dispatched in this address space on behalf of a client address space.

Job SRB Percent Cumulative SRB percent to one decimal place for the address space since collection started. This is (Job_SRB_Time x 1000) / Job_Elapsed_Time. Since multiple CPs may be active this value can exceed 100.0%.

Job SRB Time Cumulative SRB time in seconds with two decimal places for the address space since it started.

Job Start Time The date and time of day that Tivoli OMEGAMON XE on z/OS agent started tracking this address space. If the address space was active before the agent was started then this is not the address space's start time.

Job TCB Percent Cumulative TCB percentage, to one decimal place, for the address space since collection started. This is (Job_TCB_Time x 1000) / Job_Elapsed_Time. Since multiple CPs may be active, this value can exceed 100.0%.

Job TCB Time Cumulative TCB time in seconds with two decimal places for the address space since it started.

Managed System A z/OS operating system in your enterprise that is being monitored by a Tivoli OMEGAMON XE on z/OS agent. Valid value is a character string with a maximum length of 32 bytes.

Proc Step The step name of the procedure. Valid value is a string, with a maximum of eight characters.

SRB Percent The percentage of CPU utilized by SRBs in the address space over the interval. The value can exceed 100%. The maximum value is *n*00%, where *n* represents the number of processors (CPUs) assigned to the LPAR or image.

Start Up Monitored Indicates whether or not the start of the job was monitored. “No” indicates that total CPU times may not reflect the entire job. Valid values are: Yes, No.

Step Name The name of the step. Valid value is a string, with a maximum of eight characters.

SvcClass The service class name. This field will contain Not available if the system is in compatibility mode. Valid value is a string, with a maximum of eight characters.

SvcClass Period The service class period. This field will contain Not available if the system is in compatibility mode. Valid range is 0-2147483647.

TCB Percent The percentage of CPU utilized by tasks in the address space over the interval. The value can exceed 100%.The maximum value is *n* 00%, where *n* represents the number of processors (CPUs) assigned to the LPAR or image.

Total Enclave Count The total number of dependent, independent, and unknown enclaves owned by the address space that are currently active or inactive. ("Unknown" enclaves refers to an enclave not being identifiable as either a dependent or an independent enclave. This can occur if an operating system level introduces a new type of enclave classification and the version of Tivoli OMEGAMON XE on z/OS monitoring agent does not support this type.) Valid range is 0-2147483647.

Type Address space type flag. The valid types are:

- Batch
- TSO
- STC
- APPC

Unknown Active Enclave Count The total number of active enclaves owned by the address space that could not be classified as dependent or as independent. Valid range is 0-2147483647.

Unknown Inactive Enclave Count The total number of inactive enclaves owned by the address space that could not be classified as dependent or as independent. Valid range is 0-2147483647.

zIIP Percent The percentage of System z9 Integrated Information Processor (zIIP) utilized by tasks in the address space over the interval. The value can exceed 100%. The maximum value is *n*00% where *n* is the number of processors (zIIPs) assigned to the LPAR or image.

SUPCPPCT

zIIP on CP Percent The percentage of CPU utilized by tasks performing System z9 Integrated Information Processor (zIIP) work in the address space over the interval. The value can exceed 100%. The maximum value is *n*00% where *n* is the number of processors (CPUs) assigned to the LPAR or image.

Address Space Real Storage attributes

The Address Space Real Storage Attribute Group provides information about the real storage allocated to an address space in terms of various types of frame counts and slot counts, as well as the management status of a given address space.

Address Space Name The name of the job, TSO user, started task, or APPC address space reported on in the row. Valid value is a simple text string of from 1 to 8 characters.

ASID The identifier of the address space for which data is being displayed.

Central Frame Count The total number of central storage frames, including any frames that are backing user dataspaces. Valid value is a 4-byte integer.

Central Shared Pages Shared pages in central storage that are valid in the current address space.

Dispatching Priority The dispatching priority of an address space. Dispatching priorities are set either in the IEAIPS member of PARMLIB or without an IPS, through JCL, or a corresponding TSO logon. Dispatching priorities control the access of an address space to the CPU relative to other jobs and users. Dispatching priorities can range from 1 to 255, with larger numbers indicating higher priority.

Expanded Frame Count The total number of expanded storage frames, including any dataspace frames that have been paged out to expanded storage.

Fixed Frame Count The total number of fixed frames, including fixed frames found on the pageable frame queue and all frames found on the fixed frame queue. Valid value is a 4-byte integer.

High UIC How long the oldest frame of pageable storage has gone without being referenced, in seconds. The unreferenced interval count (UIC) is inversely related to contention for real storage--the lower the UIC, the more quickly frames are being referenced. When SRM determines that there is a shortage of available frames, it uses the high UIC of all address spaces in deciding which pages to steal. Frames with the highest UIC are the most likely to be stolen.

Hiperspace Frame Count The total number of hiperspace frames. Hiperspace allows a program to store and retrieve data directly from expanded storage, thus avoiding the overhead of DASD I/O. Valid value is a 4-byte integer.

Managed System A z/OS operating system in your enterprise that is being monitored by a Tivoli OMEGAMON XE on z/OS agent. Valid value is a character string with a maximum length of 32 bytes.

Management Status The management status of the address space. Valid values are:

Monitored	The Working Set Manager has noticed that the address space is using an excessive amount of CPU for paging and is actively keeping track of its central storage use and paging rates.
Managed	The Working Set Manager has determined that the paging rate of the monitored address space continues to be excessive and some type of action must be taken.
Isolated	The address space is storage isolated. An address space in this state is exempt from Working Set Manager action.
NonSwap	The address space is non-swappable and is therefore beyond Working Set Manager control.
OK	The address space is of no interest to the Working Set Manager.
Unknown	The management status of the address space is unknown.

Non-virtual I/O Slot Count The number of 4K blocks allocated in auxiliary storage for pages other than virtual I/O pages. Valid value is a 4-byte integer.

Other Aux Storage Slots Other virtual storage slots in use by the address space.

Page ins per second The page fault rate, which does not include page reclaims or pages brought in for either VIO datasets or swap-ins. A consistently high number may be an indication that work is being delayed by contention for real and expanded storage.

Page outs per second The number of page-outs per second. It includes those pages that were stolen by the real storage manager (RSM) in an attempt to increase the number of available frames. The number does not include page-outs for VIO or swap-ins requests. Although a high page-out value may not necessarily mean that workload performance is being affected by paging, it may be a sign of contention for real storage.

SPAGVIEWS

Shared Page views Member of a group of virtual addresses that are backed by the same frame of processor storage.

Status The current status of the job or user, including the dispatchability and swap status of the job. Dispatchability is displayed in the first part of the Status field, while swap status is displayed in the last part of the field. Dispatchability may be either CPU (waiting for CPU) or WAT (Waiting). Swap status may be LSW (logically swapped), NSW (Nonswappable), RES (Swapped in; resident), SWP (swapped out), or DLY (waiting for other reasons, which may include an enqueue, SRM delay, or voluntary wait. An asterisk may appear in the middle of the field, indicating that the task is currently in transaction.

Swap Status The swap status of the address space. Valid values are:

IN	The address space is swapped in and occupies central storage.
OUT	The address space is swapped out and does not occupy central storage.
IN NSW	The address space is in central storage and is non-swappable.
OUT LSW	The address space is logically swapped and occupies central storage.
UNKNOWN	Swap status is unknown.

User Key Dataspace The maximum number of 4K virtual storage blocks of user key dataspace that can be used by this address space. Valid value is a 4-byte integer.

Virtual I/O Slot Count The number of 4K blocks allocated in auxiliary storage for virtual I/O pages. Valid value is a 4-byte integer.

Working Set Size The number of frames of real storage an address space has in use. Often defined as the minimum number of pages that an application needs to do work. With expanded storage installed, working set size also includes the number of expanded storage frames that the user owns. In systems with expanded storage, the working set is divided into three parts. The primary working set consists of LSQA pages, fixed pages, and at least one page per segment. The secondary working set is all other referenced pages. The remaining, unreferenced pages are contained in the nonworking set.

Address Space Summary attributes

This attribute group provides summary address space information such as count of active address spaces and total enclave counts.

Active Enclave Count Total number of dependent and independent enclaves owned by all address spaces that are currently active. Valid value is a 4-byte integer.

Address Space Count Count of all active address spaces in this LPAR. Valid value is a 4-byte integer.

APPC Count Count of all advanced program-to-program communication address spaces in this LPAR. Valid value is a 4-byte integer.

Batch Job Count Count of all batch job address spaces in this LPAR. Valid value is a 4-byte integer.

Inactive Enclave Count Total Number of dependent and independent enclaves owned by all address spaces that are currently inactive. Valid value is a 4-byte integer.

Managed System A z/OS operating system in your enterprise that is being monitored by a Tivoli OMEGAMON XE on z/OS agent. Valid value is a character string with a maximum length of 32 bytes.

Started Task Count Count of all started task address spaces in this LPAR. Valid value is a 4-byte integer.

Total Enclave Count Total number of dependent and independent enclaves owned by all address spaces that are currently active or inactive. Valid value is a 4-byte integer.

TSO User Count Count of all TSO User address spaces in this LPAR. Valid value is a 4-byte integer.

Address Space Virtual Storage attributes

The Address Space Virtual Storage Attribute Group provides information about the virtual storage allocated to an address space. The attributes report on low, extended, and large storage. Low storage is that below the 16 Megabyte (MB) line. Extended storage is in the range from 16 MB to 2 Gigabytes (GB). Above 2 GB, storage is referred to as large storage, or “above the bar”. For each category, both fixed and virtual storage are reported on.

ASID The address space identifier as it is known to z/OS in hexadecimal.

Address Space Name The name of the job, TSO user, started task, or APPC address space reported on in the row. Valid value is an alphanumeric string, with a maximum of 8 characters.

Extended Fixed The amount of fixed storage (in megabytes) between 16 MB and 2 GB. Valid value is a 4-byte integer.

Extended Virtual The allocated virtual storage (in megabytes) between 16 MB and 2 GB. Valid value is a 4-byte integer.

Large Fixed The amount of fixed storage above the 2 GB (or above the bar) address range. Storage above the 2 GB boundary is often referred to as storage above the bar. Valid value is a 4-byte integer.

Large Inuse Percent The percent of storage in use above the bar, relative to the maximum allowed. Valid value is a 2-byte integer in the range 0.0 - 100.0.

Large Max The maximum amount of storage above the bar (2 GB boundary) that could be allocated by the address space. Valid value is a 4-byte integer.

Here are the suffixes and names for values in 64-bit:
K (kilobyte, 2^10 = 1,024)
M (megabyte, 2^20 = 1,048,576)
G (gigabyte, 2^30 = 1,073,741,824)
T (terabyte, 2^40 = 1,099,511,627,776)
P (petabyte, 2^50 = 1,125,899,906,842,624)
E (exabyte, 2^60 = 1,152,921,504,606,846,976)
64 bit = 16E = 2^64

Large Virtual The amount of allocated virtual storage above 2 GB (above the bar). Valid value is a 4-byte integer.

Low Fixed The storage below the 16 MB line, which is fixed in Megabytes. Valid value is a 4-byte integer.

Low Virtual The amount of allocated virtual storage, in Megabytes, below the 16 MB line, for this address space. Valid value is a 4-byte integer.

Managed System A z/OS operating system in your enterprise that is being monitored by a Tivoli OMEGAMON XE on z/OS agent. Valid value is a character string with a maximum length of 32 bytes in the format <plexname>:<smfid>:MVSSYS.

Management Status The management status of the address space. Valid values are:

Monitored	The Working Set Manager has noticed that the address space is using an excessive amount of CPU for paging and is actively keeping track of its central storage use and paging rates. Monitored address spaces are associated with a central storage target value of 0. An address space can be swapped out and still be monitored by the Working Set Manager.
Managed	The Working Set Manager has determined that the paging rate of the monitored address space continues to be excessive and some type of action must be taken.
Isolated	The address space is storage isolated. An address space in this state is exempt from Working Set Manager action.
NonSwap	The address space is non-swappable. An address space in this state is beyond Working Set Manager control.
OK	The address space is non-swappable. An address space in this state is beyond Working Set Manager control.
Unknown	This is a “should not happen” value. If the agent can’t determine a value, it defaults to unknown.

Step Name The address space stepname currently executing. Valid value is an alphanumeric string, with a maximum of 8 characters.

Swap Status The swap status of the address space. Valid values are:

IN	The address space is swapped in and occupies central storage.
OUT	The address space is swapped out and does not occupy central storage.
IN NSW	The address space is in central storage and is non-swappable.
OUT LSW	The address space is logically swapped and occupies central storage.
UNKNOWN	Swap status is unknown.

Total Fixed The total amount of fixed storage, in megabytes, for the address space. Valid value is a 4-byte integer.

Total Virtual The total amount of allocated virtual storage, in megabytes, for the address space. Valid value is a 4-byte integer.

Channel Paths attributes

The Channel Paths Attribute group provides information about the utilization of a channel path.

Cluster Name The name of the LPAR Cluster to which the host LPAR is assigned. The name corresponds to the I/O Cluster to which the Channel Path is assigned. The format for the name is xxxxxxxx.ssss.tttt, where:

xxxxxxxx	is the name of the sysplex
ssss	is the last 4 characters of the CPC serial number
tttt	is the first 4 characters of the CPC Model number

This format is consistent with the Cluster Name provided in the LPAR Clusters workspace. A value of Not available is provided if the channel path is not managed by Dynamic Channel Path Management (DCM). Valid value is an 18-byte character string.

Complex Percent The percentage of time that the channel was busy within the complex. Valid value is a 4-byte integer in the range 0.0 to 100.0.

Configured A description of how the channel path was reconfigured during the RMF interval. Valid values:

Unknown	
Add	Channel path was added during the RMF interval.
Mod	Channel path was modified during the RMF interval.
Del	Channel path was deleted during the RMF interval.
Not available	Status cannot be determined because an RMF substitute is being used and no CPDT is present.
_	No data was returned.

CPMF The Channel Path Measurement Facility (CPMF) calculates channel path utilization. Valid values are:

Unknown	
Available	In LPAR mode, CPMF calculates LPAR channel path utilization, which in turn determines complex-wide channel path utilization. If the channel path is shared, complex-wide utilization is calculated by other channel measurement techniques.
Unavailable	CPMF is supported but is unavailable for use. All channel path utilization information is calculated by other channel measurement means. In LPAR mode, for shared channel paths, values for LPAR channel utilization are blank. In basic mode, only complex-wide values are displayed. If you restart CPMF, this status changes to available and channel path utilization is calculated as described under Available.
Not Installed	CPMF is not supported on your system. All channel path utilization information is calculated by other channel measurement techniques. Only complex-wide values are available for both LPAR and basic modes.
Compatibility	
Extended	

DCM Status Indicates whether Dynamic Channel Path Management (DCM) is enabled for this channel path. It is possible that the Intelligent Resource Director (IRD) can change the channel's configuration to help meet Workload Manager goals. Valid values are YES or NO.

LPAR Percent The percentage of time the channel was busy working for the logical partition. This column can be blank or non-blank. The column is blank when: a) the system is running in basic mode; b) the channel path is offline; c) the ESCON® Multiple Image Facility (EMIF) is not installed; or d) CPMF is unavailable or not installed.

Valid values range from 0 to 100% when the system is running in LPAR mode, EMIF is installed, and the channel path is online. If the channel path is shared, this value represents the LPAR's utilization of the channel path from the start of the RMF interval. If the channel path is not shared, the LPAR utilization value is the same as the complex-wide utilization value.

This value can be greater than the utilization of the channel by the entire complex, since the LPAR value is obtained from the hardware, whereas the Complex value is sampled and may not include channel utilization.

Managed System An z/OS operating system in your enterprise that a Tivoli OMEGAMON XE for z/OS agent is monitoring. Valid value is a character string with a maximum length of 32 bytes.

Mode The mode of the channel path. Valid values are:
- Unknown
- Basic (no partitions)
- LPAR (partitioned)

Online An indicator showing whether the channel is online (Y) or offline (N). Valid values are Y or N.

Path ID The channel path ID. Valid value is 1-byte hexadecimal.

RMF Interval Start The time in HH:MM:SS format when the current RMF interval started.

Sample Count The number of samples in this measurement interval. Valid value is a 4-byte integer.

Shared An indicator indicating whether the channel is shared with other logical partitions. Only ESCON channels (ESchn) or directors (ESdir) can be shared channels. Valid values are Y or N.

Type The channel data type. Some of the valid values are shown below. Because the list of possible values is long and changes often, please refer to IBM's RMF Report Analysis Manual for the most current list.

Block	Block multiplexor channel
Byte	Byte multiplexor channel
ESchn	ESCON channel
ESctc	ESCON channel-to-channel adapter
EScnv	Channel connected to an ESCON converter
ESdir	ESCON channel attached to an ESCON director
Unkn	Unknown type

Common Storage attributes

The Common Storage Attribute Group displays information about four important areas of common storage: common service area (CSA), extended CSA, system queue area (SQA), and extended SQA.

Allocation The amount of storage currently allocated for the area. For SQA and ESQA, under normal conditions, this value is the same as totsize. If an overflow has occurred, then the value may exceed totsize. Valid value is a numeric in the range 0 to 2147483647.

Allocation Percent The amount of storage currently allocated as a percentage of total storage in this area. Valid range is 0 - 100.

Area The name of the common storage area presented in this row of the table view. For SQA and ESQA, under normal conditions, this will be 100%. If an overflow has occurred, then the value may exceed 100%. Valid values are:

- CSA — Common Storage Area
- ECSA — Extended Common Storage Area
- SQA — System Queue Area
- ESQA — Extended System Queue Area

ESQA Overflow The amount of ESQA that has overflowed into the ECSA. Valid value is a numeric in the range 0 to 2147483647.

Growth Growth in use during the last interval. Valid value is a numeric in the range 0 to 2147483647.

In Use The amount of storage currently in use, not including the SQA or ESQA overflow. Valid value is a 4-byte integer.

In Use Percent The amount of storage currently in use as a percentage of total storage. Valid value is a numeric in the range 0 to 2147483647.

Managed System An z/OS operating system in your enterprise that a Tivoli OMEGAMON XE for z/OS agent is monitoring. Valid value is a character string with a maximum length of 32 bytes.

SQA Overflow The amount of SQA that has overflowed into the CSA. Valid value is a numeric in the range 0 to 2147483647.

Total Size The total size of each area as specified at IPL time or in the IEASYS member of the PARMLIB data set minus any CSA/ECSA storage used by SQA/ESQA. Valid value is a numeric in the range 0 to 2147483647.

Unowned The amount of allocated storage that is not currently owned by an address space. Valid value is a numeric in the range 0 to 2147483647.

DASD MVS attributes

The DASD MVS Attribute Group assists you in monitoring various z/OS error conditions collectively for a group of devices.

Cache Deactivated The count of volumes that are eligible for cache but do not have cache activated. Valid value is a 4-byte integer.

Dropped Ready Count of DASD showing a "dropped ready" status. The dropped ready condition may indicate a power supply or other hardware problems, or the switch may have been turned off at the device. Valid value is a 4-byte integer.

Indexed VTOC Lost Count of DASD with Indexed VTOC disabled (that is, non-indexed format VTOC). Note that on heavily used volumes with many datasets, an index to a VTOC may significantly improve I/O performance. System routines at times may disable a VTOC index and return the VTOC to its non-indexed format. Use the DSF BUILDIX command to return to the VTOC to indexed format. If the problem persists, you may need to delete and rebuild the index using DSF. Valid value is a 4-byte integer.

Managed System A z/OS operating system in your enterprise a Tivoli OMEGAMON XE for z/OS agent is monitoring. Valid value is a character string with a maximum length of 32 bytes.

No Dynamic Path Reconnect The count of devices with dynamic path reconnect disabled. Valid value is a 4-byte integer.

Not Responding Count of DASD not responding to I/O requests. When a device is not responding to I/O requests, start pending messages may appear on the system console. In multiple systems having shared DASD environments, this may be caused by a RESERVE command issued by another system. If this is the cause of the problem, it may be possible to change the RESERVE to a global enqueue. If you cannot find the cause of any unknown reserves or enqueues, a hardware malfunction is possible, causing the device to show dropped ready. Valid value is a 4-byte integer.

DASD MVS Devices attributes

The DASD MVS Devices Attribute Group permits you to monitor z/OS activity for individual DASD devices.

Address The device address for this DASD volume.

Cache Read Hit Percent This value represents a successful I/O request to read data from the cache. The ratio of successful reads to total read requests is the hit percentage. Valid value is a 4-byte integer.

Cache Status The cache status of the device for the system. Valid values are:

- Unknown
- Active
- Inactive
- PendingActive
- PendingInactive
- TimedOut

Cache Write Hit Percent This value represents a successful I/O request to write temporary data to the cache. The ratio of successful writes to total write requests is the hit percentage. Valid value is a 4-byte integer.

CU Busy Delay Time Control Unit busy delay time in milliseconds. Valid value is a 4-byte integer.

Dev Allocations Average number of allocated datasets that are open on the device.

Dev Busy Delay Time Device busy delay time in milliseconds. Valid value is a 4-byte integer.

Director Port Busy Delay ESCON Director Port busy delay time in milliseconds. Valid value is a 4-byte integer.

Fast Write Hit Percent This value represents write requests that are cached and later placed in nonvolatile storage (NVS) and written to the DASD device. Valid value is a 4-byte integer.

HyperPAV Indicates whether or not HyperPAV is enabled for this device. Valid values are yes and no.

I/O Connect Time The time, in milliseconds, attributable to data transfer (search + transfer). Valid value is a 4-byte integer.

I/O Disconnect Time The time, in milliseconds, that an I/O request spends freed from the channel. This is the time that the I/O searches for the data that has been requested. This includes moving the device to the requested cylinder and track (SEEK + SET SECTOR), waiting for the record to rotate under the head (LATENCY), and possible rotational delay as the device waits to reconnect to the channel (RPS delay). Valid value is a 4-byte integer.

I/O Pending Time The time, in milliseconds, that the I/O is delayed in the path to the device. Pending time may be attributable to the channel, control unit, or head of string being busy, although it is often caused by shared DASD. Valid value is a 4-byte integer.

I/O Rate Number of I/Os per second to the device. Valid value is a 4-byte integer.

IOS Queue Time The average time, in milliseconds, that an I/O waits because the device is busy. Valid value is a 4-byte integer.

LCU Number A device's Logical Control Unit (LCU) number (a number that represents the aggregate paths to the device). Valid value is a 4-byte integer in the range 0 - 255.

Model The control unit model type/number. Valid value is an 8-byte character string.

Managed System A z/OS operating system in your enterprise monitored by an Tivoli OMEGAMON XE on z/OS agent. Valid value is a character string with a maximum length of 32 bytes.

PAV Count The number of Parallel Access Volumes (PAVs) assigned to this DASD volume, or, in the case of base HyperPAV volumes, the average number of HyperPAV aliases assigned over the interval. If none, a value of Not available is displayed. Valid value is a 4-byte integer.

Percent Busy The percentage of time the device was busy doing I/O during the last measurement interval. Valid value is a 4-byte integer in the range 0.0 to 100.0.

Percent Reserve The percentage of time the device was observed to be in a Reserved state. Valid value is a 4-byte integer.

Response The average response time for I/Os to this device, in milliseconds.

Storage Group Name of the storage group to which the device is assigned. Valid value is an 8-byte character string.

Volume The volume serial number of the device. Valid value is a 6-byte character string.

Enclave Detail attributes

The Enclave Detail attributes permit you to view performance and extended classification information for a selected enclave. These attributes cannot be used to create situations or collect history.

Account Information A Workload Manager classification attribute. If this value is unavailable, this column, as well as all other classification attribute columns will display Not available.

Collection A Workload Manager classification attribute. If this value is unavailable, this column, as well as all other classification attribute columns will display Not available.

Connection A Workload Manager classification attribute. If this value is unavailable, this column, as well as all other classification attribute columns will display Not available.

Correlation A Workload Manager classification attribute. If this value is unavailable, this column, as well as all other classification attribute columns will display Not available.

CPU Percent The percentage of CPU recently consumed by the enclave. This value cannot be determined for a very recently created enclave, in which case Not available is displayed. In some unusual cases, restarting a transaction can make the value appear as a negative number. In that case, the value NEG is displayed. This value can exceed 100%. The maximum value is n 00%, where *n* is the number of processors (CPUs) assigned to the LPAR or image.

Export Token The 32-byte export token in hexadecimal.

Function Name A Workload Manager classification attribute. If this value is unavailable, this column, as well as all other classification attribute columns will display Not available.

IFA Percent The percentage of IFA recently consumed by the enclave. The determination of this value is not possible for a very recently created enclave. Unusual circumstances involving the restarting of a transaction can make the value on which this number is based appear negative. In that case, NEG is displayed. The value can exceed 100%. The maximum value is *n* 00%*IFA normalization factor, where n is the number of IFAs configured to the LPAR and IFA normalization factor is a multiplier that allows normalization of IFA speeds to standard CP speeds on “knee-capped” processors.

IFA Percent on CP Average percentage of time the system consumes regular CP resource executing IFA work across all regular CPs configured to this LPAR.

LUName The logical unit name. This column represents a Workload Manager classification attribute. If this value is unavailable, this column, as well as all other classification attribute columns will display Not available.

Managed System An z/OS operating system in your enterprise that is being monitored by a Tivoli OMEGAMON XE on z/OS agent. Valid value is a character string with a maximum length of 32 bytes.

NetID A Workload Manager classification attribute. If this value is unavailable, this column, as well as all other classification attribute columns will display Not available.

Owning ASID ID of the owning address space in hexadecimal. This column may contain 0 if the enclave is foreign, or in unusual recovery situations.

Owning JESJOBID The JES job ID of the address space. In the following circumstances, this information is not available:

blank	The owning address space is not using JES services.
OTHSYS	The report was produced by a system other than the owning system.
GONE	The owning job is no longer executing. This may mean that the owning job’s address space no longer exists, or that the address space is now assigned to a different job. At any one instant in time, both an enclave and its owning address space will or will not exist. However, it is not possible for a Tivoli OMEGAMON XE for z/OS agent to gather all enclave and address space information in one atomic operation. Thus, when an address space owning an enclave has recently ceased to exist, it is possible for the agent to have enclave information and still be unable to locate information associated with the owning address space.

Owning Jobname The job name of the owning address space. This attribute is not available on all levels of z/OS.

Owning System Name of the system where the owning address space is executing. This attribute is not available on all levels of z/OS.

Package A Workload Manager classification attribute. If this value is unavailable, this column, as well as all other classification attribute columns will display Not available.

Perform A Workload Manager classification attribute. If this value is unavailable, this column, as well as all other classification attribute columns will display Not available.

Plan A Workload Manager classification attribute. If this value is unavailable, this column, as well as all other classification attribute columns will display Not available.

Procedure Name A Workload Manager classification attribute. If this value is unavailable, this column, as well as all other classification attribute columns will display Not available.

Process Name A Workload Manager classification attribute. If this value is unavailable, this column, as well as all other classification attribute columns will display Not available.

ResGroup The resource group name.

RptClass The report class name.

Scheduling Environment A Workload Manager classification attribute. If the data is not available, Not available is displayed. Not available is also displayed if this and other classification attributes are not available.

Status The enclave's status. This value will be one of the following:

- Active—if the enclave is active
- Inactive—if the enclave is inactive. An enclave is considered inactive if there is no unit of work currently associated with the enclave.
- Not available—For older or unsupported releases of the system.

Subsystem Collection A Workload Manager classification attribute. Not available is displayed if this and other classification attributes are not available.

Subsystem Name A Workload Manager classification attribute. If this value is unavailable, this column, as well as all other classification attribute columns will display Not available.

Subsystem Parameter A Workload Manager classification attribute. If this value is unavailable, this column, as well as all other classification attribute columns will display Not available.

Subsystem Priority A Workload Manager classification attribute. If this value is unavailable, this column, as well as all other classification attribute columns will display Not available.

Subsystem Type A Workload Manager classification attribute. If this value is unavailable, this column, as well as all other classification attribute columns will display Not available.

Subtype Type of subsystem to which the enclave belongs. This attribute is not available on all levels of z/OS. If the information is not available, the column contains Not available.

SvcClass The service class name.

SvcClass Period The service class period.

Token The 8-byte token name in hexadecimal format. The entire token is displayed including any leading zeroes. If you are relating information in this report to information provided by other products, you should be aware that some products do not display the leading zeroes. The token can be used to RESET the enclave. There is no z/OS command for this function, but it is supported by SDSF.

Total CPU The total CPU consumption in seconds. For a multisystem enclave, this value does not include CPU consumption on systems other than the system that produced the report.

Total IFA Total IFA consumption, in seconds. Valid value is an integer in the range 1 - 2147483647.

Total IFA on CP Total regular CP resource consumption executing IFA work, in seconds. For a multisystem enclave, IFA consumed on systems other than the system which produced the report is not included in this figure. Valid value is an integer in the range 1 - 2147483647.

Total zIIP Total System z9 Integrated Information Processor (zIIP) consumption, in seconds.

Total zIIP on CP Total standard process resource consumed performing work eligible for a System z9 Integrated Information Processor (zIIP), in seconds. For a multisystem enclave, zIIP consumed on systems other than the system which produced the report is not included in this figure.

Transaction Class A Workload Manager classification attribute. If this value is unavailable, this column, as well as all other classification attribute columns will display Not available.

Transaction Program Name A Workload Manager classification attribute. If this value is unavailable, this column, as well as all other classification attribute columns will display Not available.

Type The enclave type. This value will be one of the following:

- Dep — Dependent
- Indep — Independent
- ForeignDep — Foreign Dependent
- ForeignIndep — Foreign Independent
- Unknown — All other types

Userid A Workload Manager classification attribute. If this value is unavailable, this column, as well as all other classification attribute columns will display Not available.

Workload The name of the workload.

zIIP on CP Percent The percentage of regular processor resource recently consumed by the enclave performing work eligible for a System z9 Integrated Information Processor.

zIIP Percent The percentage of System z9 Integrated Information Processor (zIIP) resources recently consumed by the enclave. The determination of this value is not possible for a very recently created enclave. Unusual circumstances involving the restarting of a transaction can make the value on which this number is based appear negative. In that case, NEG is displayed.

This value can exceed 100%. The maximum value is *n*00%*zIIP normalization factor, where *n* is the number of zIIPs configured to the LPAR and the zIIP normalization factor is a multiplier that allows normalization of zIIP speeds to standard processor speeds on ″knee-capped″ processors.

Enclave Table attributes

An enclave is defined as a transaction that can span multiple dispatchable units (SRBs and tasks) in one or more address spaces and is classified and managed as a unit. The Enclave Table attributes permit you to view information about or create situations that report on the status of various resources belonging to the enclave.

CPU Percent The percentage of regular processor time recently consumed by the enclave. This value cannot be determined for a very recently created enclave, in which case Not available is displayed. In some unusual cases, restarting a transaction can make the value appear as a negative number. In that case, the value NEG is displayed. This value can exceed 100%. The maximum value is *n*00%, where *n* is the number of processors (CPUs) assigned to the LPAR or image.

Export Token The 32-byte export token in hexadecimal.

IFA Percent The percentage of IFA recently consumed by the enclave. The determination of this value is not possible for a very recently created enclave. Unusual circumstances involving the restarting of a transaction can make the value on which this number is based appear negative. In that case, NEG is displayed. The value can exceed 100%. The maximum value is *n*00%*IFA normalization factor, where *n* is the number of IFAs configured to the LPAR and IFA normalization factor is a multiplier that allows normalization of IFA speeds to standard CP speeds on “knee-capped” processors.

IFA Percent on CP Average percentage of time the system consumes regular CP resource executing IFA work across all regular CPs configured to this LPAR.

Managed System An z/OS operating system in your enterprise that is being monitored by a Tivoli OMEGAMON XE on z/OS agent. Valid value is a character string with a maximum length of 32 bytes.

Owning ASID ID of the owning address space in hexadecimal. This column may contain 0 if the enclave is foreign, or in unusual recovery situations.

Owning JESJOBID The JES job ID of the address space. In the following circumstances, this information is not available:

blank	The owning address space is not using JES services.
OTHSYS	The report was produced by a system other than the owning system.
GONE	The owning job is no longer executing. This may mean that the owning job’s address space no longer exists, or that the address space is now assigned to a different job. At any one instant in time, both an enclave and its owning address space will or will not exist. However, it is not possible for a Tivoli OMEGAMON XE for z/OS agent to gather all enclave and address space information in one atomic operation. Thus, when an address space owning an enclave has recently ceased to exist, it is possible for the agent to have enclave information and still be unable to locate information associated with the owning address space.

Owning Jobname The job name of the owning address space. This attribute is not available on all levels of z/OS.

Owning System Name of the system where the owning address space is executing. This attribute is not available on all levels of z/OS.

ResGroup The resource group name.

RptClass The report class name.

Status The enclave’s status. This value will be one of the following:

- Active—if the enclave is active
- Inactive—if the enclave is inactive. An enclave is considered inactive if there is no unit of work currently associated with the enclave.
- Not available—For older or unsupported releases of the system.

Subtype Type of subsystem to which the enclave belongs. This attribute is not available on all levels of z/OS. If the information is not available, the column contains Not available.

SvcClass The service class name.

SvcClass Period The service class period.

Token The 8-byte token name in hexadecimal format. The entire token is displayed including any leading zeroes. If you are relating information in this report to information provided by other products, you should be aware that some products do not display the leading zeroes. The token can be used to RESET the enclave. There is no z/OS command for this function, but it is supported by SDSF.

Total CPU The total CPU consumption in seconds. For a multisystem enclave, this value does not include CPU consumption on systems other than the system that produced the report.

Total IFA Total IFA consumption, in seconds. Valid value is an integer in the range 1 - 2147483647.

Total IFA on CP Total regular processor resource consumption performing work eligible for a zSeries Application Assist Processor (zAAp), in seconds. For a multisystem enclave, zAAP consumed on systems other than the system which produced the report is not included in this figure. Valid value is an integer in the range 1 - 2147483647.

Total zIIP Total System z9 Integrated Information Processor consumption in seconds.

Total zIIP on CP Total standard resource consumed performing work eligible for System z9 Integrated Information Processors (zIIPs), in seconds. For a multisystem enclave, zIIP resources consumed on systems other than the system which produced the report is not included in this figure.

Type The enclave type. This value will be one of the following:
- Dep — Dependent
- Indep — Independent
- ForeignDep — Foreign Dependent
- ForeignIndep — Foreign Independent
- Unknown — All other types

Workload The name of the workload.

zIIP on CP Percent The percentage of regular processor resource recently consumed by the enclave performing work eligible for a System z9 Integrated Information Processor (zIIP). The determination of this value is not possible for a very recently created enclave. Unusual circumstances involving the restarting of a transaction can make the value on which this number is based appear negative. In that case, NEG is displayed. The value can exceed 100%. The maximum value is *n*00%*zIIP normalization factor, where *n* is the number of zIIPs configured to the LPAR and the zIIP normalization factor is a multiplier that allows normalization of zIIP speeds to standard processor speeds on ″knee-capped″ processors.

zIIP Percent The percentage of System z9 Integrated Information Processor resources recently consumed by the enclave. The determination of this value is not possible for a very recently created enclave. Unusual circumstances involving the restarting of a transaction can make the value on which this number is based appear negative. In that case, NEG is displayed.

Enqueue Conflicts attributes

The Enqueue Conflicts attributes assist you in determining those system resources that have contention and those address spaces that are being delayed because they cannot acquire the needed resources.

ASID Hexadecimal address space ID.

Major Name This is the major name (qname) of the resource. Valid format is a simple text string of from 1 through 8 characters; for example, SYSDSN.

Managed System A z/OS operating system in your enterprise that a Tivoli OMEGAMON XE for z/OS agent is monitoring. Valid value is a character string with a maximum length of 32 bytes.

Maximum Wait Time Longest time, in seconds, that the task has been waiting for a resource. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 45.

Minor Name The minor name (rname) of the resource. Valid format is a simple text string of from 1 through 255 characters; for example, SYS1.PROCLIB.

Owning Address Space Name of the job containing the task that owns the resource. Valid format is a simple text string from 1 through 8 characters; for example, PRODJOB1. The column is blank for a waiting address space entry.

Owning Task Count Number of tasks owning the resource. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 1.

Swapped Swap status of the address space in which the task is executing. Valid values are SWAPPED or NOT SWAPPED.

System Name System name or SMF ID of the z/OS image where the task is executing. Valid format is a simple text string of from 1 through 8 characters; for example, SYSA.

Type Type of ENQ the task has issued for the resource. Valid values are:

- Shared - Capable of sharing with other resources
- Exclusive - No other processes have access to the resource

Wait Time Time in seconds that the task has been waiting for the resource. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 45.

Waiting Address Space Name of the address space currently waiting for the resource. Valid format is a simple text string from 1 through 8 characters; for example PRODJOB2. The column is blank for an owning address space entry.

Waiting Task Count Number of tasks waiting for the resource. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. An example is 1.

HiperDispatch Logical Processors attributes

The HiperDispatch Logical Processors attributes show the processor type and HiperDispatch statistics like priority, share percentage, and status for each processor configured in an LPAR.

CPU Type Type of processor. The possible values are:

Standard
A standard processor.

zAAP
zSeries Application Assist Processor, a special class of assist processor designed to run Java workloads. For reporting purposes, a zAAP is sometimes referred to as an integrated facility for applications (IFA).

zIIP
zSeries Integrated Information Processor, a special class of assist processor used mostly for DB2 workloads.

HiperDispatch Priority The HiperDispatch priority of work dispatched on the logical processor. The possible values are:

High
90% of the processor's execution time is spent running High priority work.

Medium
90% of the processor's execution time is spent running either High or Medium priority work.

Low
All other processors.

LCPU ID ID of the logical processor. This is a hexadecimal formatted value.

Managed System Name of the managed system where this data is collected. It has the format <*plexname*>:<*smfid*>:MVSSYS, where *plexname* is the sysplex this LPAR resides in and *smfid* is the name of the system management facility for this LPAR.

Management Percent The percentage of physical processor time spent in this logical processor managing this LPAR, formatted to three decimal places.

Physical CP Dispatch Pct The percentage of time that a physical processor was dispatched to this logical processor, formatted to three decimal places.

Share Percent The percentage of a physical processor that the logical processors are entitled to receive, formatted to 1 decimal place. For HiperDispatch, High-priority logical processors this should be 100%. Low priority logical processors should be 0%. Medium priority logical processors will have a share percentage proportional to their Polar Weight.

Status Status of the logical processor. The possible values are:

Online
The logical processor is configured and allowed to have work dispatched on it.

Offline
The logical processor is configured but is currently not available for work to be dispatched. Intelligent Resource Director has likely reduced the relative weight of this partition in favor of other partitions.

Parked
This logical processor is configured online but will not dispatch any work. The processor can be unparked under the right conditions and start dispatching work for this LPAR.

Park Pending
The logical processor is in the process of being parked.

Reserved
The logical processor is configured but has never been brought online.

HiperDispatch Management attributes

The HiperDispatchManagement attributes provide information about the weight and status of a logical processor when the LPAR is in vertical mode.

Current Weight Current weight for logical processors of this type in this LPAR. The range is 0 – 999 for shared processors, unless the LPAR uses dedicated CPs, in which case the value is 65,535. If the associated processor type is not used in the LPAR, this value is shown as **undefined**.

Logical Processor Type Type of logical processor. The possible values are:

Standard
A standard processor

zAAP
zSeries Application Assist Processor, a special class of assist processor designed to run Java workloads. For reporting purposes, a zAAP is sometimes referred to as an integrated facility for applications (IFA).

zIIP
zSeries Integrated Information Processor, a special class of assist processor used mostly for DB2 workloads.

HiperDispatch Management The status of the HiperDispatch feature. The possible values are: On, Off, Unavailable.

Unavailable
Either the required hardware is not installed or the operating system does not support it.

On
This LPAR is in HiperDispatch management mode.

Off
This LPAR is not in HiperDispatch management mode. Either HiperDispatch management mode has not been enabled in the SYS1.PARMLIB, or the LPAR does not have enough work to warrant its use.

LPAR Cluster Name of the cluster to which the LPAR belongs, as defined in the hardware console. LPARs that are members of the same cluster may share central processor and channel resources through the Intelligent Resource Director.

LPAR Group Name of the group to which the LPAR belongs, as defined in the hardware console. LPARs in the same group may have a group capacity defined. The sum of 4-hour rolling MSUs for the group may not exceed the group capacity. LPARs may also have individual capacity limits. An LPAR could be capped if either its individual limit is reached or the group total is reached.

LPAR Name Name of the logical processor, as defined in the hardware console.

Managed System Name of the managed system where this data is collected. It has the format <*plexname*>:<*smfid*>:MVSSYS, where *plexname* is the sysplex this LPAR resides in and *smfid* is the name of the system management facility for this LPAR.

Maximum Weight Maximum weight that can be assigned to the set of logical CPUs of the associated type within the logical partition. The weight is in the range 0-999; except that it is 65,535 if the logical partition uses dedicated CPUs, and 0 if no maximum weight was specified. If this processor type is not present in this LPAR the weight is shown as **undefined**.

Minimum Weight Minimum weight that can be assigned to the set of logical processors of the associated type within the logical partition. The weight is in the range 0-999, except that it is 65,535 if the logical partition uses dedicated CPUs, and 0 if no minimum weight was specified. If this processor type is not present in this LPAR the weight is shown as **undefined**.

Inspect Address Space CPU Use attributes

These attributes provide the Inspect data for the active task control blocks (TCBs) within the selected address space. You cannot use these attributes to create situations.

ASID IN Used to pass the ID of the address space to be inspected to the agent.

CSECT Address Hexadecimal address of the control section (CSECT) in storage within the owning load module. If Inspect is unable to determine the CSECT name, the value of this field will be the same as the load module address.

CPU % of CSECT CPU percentage usage for this granular section of the control section (CSECT). If the name of the CSECT is unknown, this value will be the same as Load Module. Valid value is a number in the range 0.0 - 100.0.

CSECT CPU % of Job CPU usage for this control section (CSECT) as a percentage of the CPU usage for all task control blocks (TCBS) in the address space. This value will be the same as Load Module CPU % of Job if CSECT name is unknown, since in that case the CSECT is mapped to the entire load module. Valid value is an integer in the range 0.0 - 100.0. This value will be 100% if the CSECT name is unknown, since the CSECT is then mapped to the entire load module.

CSECT CPU % of Load Module CPU usage for this control section (CSECT) as a percentage of the CPU usage for the owning load module. This value will be 100% if the CSECT name is unknown, since the CSECT is then mapped to the entire load module. Valid value is a number in the range 0.0 - 100.0.

CSECT CPU % of TCB CPU usage for this control section (CSECT) as a percentage of the CPU usage for the owning task control block (TCB). This value will be the same as Load Module CPU % of TCB if the CSECT name is unknown, since in this case the CSECT is mapped to the entire load module.Valid value is a number in the range 0.0 - 100.0.

CSECT Name Name of the control section (CSECT) within the load module in which Inspect detected the execution. This column is blank if Inspect is unable to determine the load library that the module was loaded from. This may occur if the module was loaded from a link list library. Valid value is a string of up to 8 characters.

CSECT Offset in Load Module Offset of the control section in the load module. Valid value is a hexadecimal number in the range 0 - 7FFFFFFF.

Initial Program Name of the top level program that was attached to create this task control block (TCB). If the target address space is a CICS TS 1.3 or higher region, this field may contain the CICS transaction ID and task number instead. Valid value is a string of up to 8 characters.

Interval The time, in milliseconds, between samples taken by the Inspect agent. Valid value is an integer in the range 0-2147483647.

INTERVAL IN Used to specify the sampling interval to the agent. Valid value is an integer in the range 0-2147483647.

Jobname Used by the Inspect agent to check that the same job is still running in the target address space. Valid value is a string of up to 8 characters.

Load Module Address Address of the load module in the target address space. A value of *-CSA-* indicates that execution was in code resident in the common storage area; a value of *-ECSA-* indicates that execution was in code resident in the extended common storage area. A zero indicates that Inspect could not determine the load module name.

Load Module ASID Decimal Decimal ID of the address space when execution was seen to be in an address space other than the one being inspected.

Load Module ASID Hex Hexadecimal ID of the address space when execution was seen to be in an address space other than the one being inspected.

Load Module ASID Jobname Name of the job running within the address space when execution was seen to be in an address space other than the one being inspected. Valid value is a string of up to 8 characters.

Load Module CPU % of Job CPU usage for this load module as a percentage of the total CPU time seen by Inspect for all task control blocks (TCBs) in the address space. Valid value is a nubmer in the range 0.0 - 100.0.

Load Module CPU % of TCB CPU usage for this load module as a percentage of the total CPU time seen by Inspect for the owning task control block (TCB). Valid value is a number in the range 0.0 - 100.0.

Load Module Name The name of the load module that Inspect saw was executing when it took one or more samples. This field may contain *-CSA-* or *-ECSA-* if execution was within resident code within the CSA or ECSA. It may contain *-UNKN-* if Inspect was unable to determine the load module name. Valid value is a string of up to 8 characters.

Managed System A z/OS operating system in your enterprise that is being monitored by Tivoli OMEGAMON XE on z/OS. Valid value is a string of up to 32 characters, in the format <plexname>:<smfid>:MVSSYS.

Message Text of messages issued by the agent. Valid value is a string of up to 80 characters.

Offset in CSECT Offset in the control section (CSECT) of the granular CPU usage. Valid value is an integer in the range 0-2147483647.

Row An incrementing number returned for each row of data, starting at 1. This value is used to restore the data order in the table if you have sorted the data display by another column. This attribute is for internal use only. It should not be used in situations or queries.

Samples Number of samples to be taken by the Inspect agent.

SAMPLES IN Used to pass to the agent the number of samples to be taken. Valid value is an integer in the range 0-2147483647.

Samples Taken Number of samples taken by the Inspect process. The Inspect process ends when the number of samples taken equals the number of samples specified to be taken by the Inspect process in the SAMPLES IN parameter, or when the job in the address space being inspected ends, in which case the number of samples used will be less than SAMPLES IN and Samples. Valid value is an integer in the range 0-2147483647.

Samples Used Number of samples used to create the Inspect data, that is, the number of samples where the Inspect agent saw the target address space actually using CPU. You can use this value to determine the statistical accuracy of the sampled Inspect data. Valid value is an integer in the range 0-2147483647.

TCB Address Address of the task control block (TCB) within the target address space, in hexadecimal.

TCB CPU % of Job The percentage of CPU usage for this task control block (TCB), as a percentage of all the CPU time seen by Inspect for all TCBs in the address space. Valid value is a number in the range 0.0 - 100.0.

TCB Ended Indicates if task control block (TCB) ended (Yes) or not (blank) while Inspect was running.

Integrated Cryptographic Services Facilities Subsystems (ICSF) attributes

One row emitted per cryptographic agent to display subsystem and coprocessor status.

ICSF is a z/OS subsystem that provides cryptographic services to system functions and application servers. It provides publicly-documented service call exits that you may use. You can specify exits for each callable cryptographic service and other administrative function of ICSF. The table below shows the entry points.

Note: If you need to define your own exits, use the ICSF security exits as alternatives to the two service call exits, CSFEXIT3 and CSFEXIT4. If the monitoring agent discovers a user-defined exit that conflicts with an IBM performance-monitoring exit, it replaces the user-defined exit, issues a warning message, and proceeds with data collection.

Cryptographic service or function	Entry Point
ANSI X9.17 EDC Generate	CSFAEGN
ANSI X9.17 Key Export	CSFAKEX
ANSI X9.17 Key Import	CSFAKIM
ANSI X9.17 Key Translate	CSFAKTR
ANSI X9.17 Transport Key Partial Notarize	CSFATKN
Clear Key Import	CSFCKI
Clear PIN Encrypt	CSFCPE
Clear PIN Generate	CSFPGN
Clear PIN Generate Alternate	CSFCPA
Cipher/Decipher	CSFEDC
Ciphertext Translate	CSFCTT
Ciphertext Translate (with ALET)	CSFCTT1
Control Vector Translate	CSFCVT
Cryptographic Variable Encipher	CSFCVE
Data Key Export	CSFDKX
Data Key Import	CSFDKM
Decipher	CSFDEC
Decipher (with ALET)	CSFDEC1
Decode	CSFDCO
Digital Signature Generate	CSFDSG
Digital Signature Verify	CSFDSV
Diversified Key Generate	CSFDKG
Encipher under Master Key	CSFEMK
Encipher	CSFENC
Encipher (with ALET)	CSFENC1
Encode	CSFECO
Encrypted PIN Generate	CSFEPG
Encrypted PIN Translate	CSFPTR
Encrypted PIN Verify	CSFPVR
Generate a key	CSFGKC
Import a key	CSFRTC
Key Export	CSFKEX
Key Generate	CSFKGN
Key Import	CSFKIM
Key Part Import	CSFKPI
Key Record Create	CSFKRC
Key Record Delete	CSFKRD

Key Record Read	CSFKRR
Key Record Write	CSFKRW
Key Test	CSFKYT
Key Test Extended	CSFKYTX
Key Translate	CSFKTR
MAC Generate	CSFMGN
MAC Generate (with ALET)	CSFMGN1
MAC Verify	CSFMVR
MAC Verify (with ALET)	CSFMVR1
MDC Generate	CSFMDG
MDC Generate (with ALET)	CSFMDG1
Multiple Clear Key Import	CSFCKM
Multiple Secure Key Import	CSFSKM
One Way Hash Generate	CSFOWH
One Way Hash Generate (with ALET)	CSFOWH1
PCI Interface	CSFPCI
PKA Decrypt	CSFPKD
PKA Encrypt	CSFPKE
PKA Key Generate	CSFPKG
PKA Key Import	CSFPKI
PKA Public Key Extract	CSFPKX
PKDS Record Create	CSFPKRC
PKDS Record Delete	CSFPKRD
PKDS Record Read	CSFPKRR
PKDS Record Write	CSFPKRW
PKSC Interface	CSFPKSC
Prohibit Export	CSFPEX
Prohibit Export Extended	CSFPEXX
Random Number Generate	CSFRNG
Retained Key Delete	CSFRKD
Retained Key List	CSFRKL
Secure Key Import	CSFSKI
SET Block Compose	CSFSBC
SET Block Decompose	CSFSBD
Symmetric Key Export	CSFSYX
Symmetric Key Generate	CSFSYG
Symmetric Key Import	CSFSYI
Transform CDMF Key	CSFTCK
User Derived Key	CSFUDK
VISA CVV Service Generate	CSFCSG
VISA CVV Service Verify	CSFCSV

1_CC Cryptographic Coprocessor Available Indicates whether at least one cryptographic coprocessor is available. The values are: Yes, No, or Unknown.

1_CMOS Indicates whether at least one CMOS cryptographic coprocessor is available. The values are: The values are: Yes, No, or Unknown.

1_PCI Indicates whether at least one PCI coprocessor is available. The values are: The values are: Yes, No, or Unknown.

ASID The address space ID of the ICSF subsystem.

AvgWait The average internal wait time in seconds per sample.

CCC A cryptographic configuration control bit hexadecimal string.

CCMKeyOK Indicates whether a valid master key has been loaded into a coprocessor. The values are: The values are: Yes, No, or Unknown.

CDMF Indicates whether Commercial Data Masking Facility is enabled. The values are: Enabled, Disabled, or Unknown.

CICSWAITL Indicates the address of the CICS wait list represented as a hexadecimal string. A value of 0 indicates the wait list is not configured.

CKDS_80Full Indicates 80% or more utilization of the Cryptographic Key Dataset space. The values are: Yes, No, or Unknown.

CKDSAccess Indicates whether dynamic Cryptographic Key Dataset access is enabled. The values are: Enabled, Disabled, or Unknown.

CKDSname The Cryptographic Key Dataset name.

CryptoSvcs Indicates the status of the cryptographic services. The values are: Active or Inactive.

DES Indicates whether DES is enabled. The values are: Enabled, Disabled, or Unknown.

DomainIdx Is the Domain Index used to access coprocessors from an LPAR. An LPAR is a Logical Partition in a PR/SM environment. See PR/SM for more information.

KMMK_CMOS0 Indicates the state of the Public Key Algorithm, Key Management Master Key in CMOS coprocessor C0. The values are: Valid, Reset, and Unknown.

KMMK_CMOS1 Indicates the state of the Public Key Algorithm, Key Management Master Key in CMOS coprocessor C1. The values are: Valid, Reset, and Unknown.

KMMKey The Public Key Algorithm Key Management Master Key hash pattern.

MKey The Master Key verification pattern and authentication pattern.

MKVer The current Master Key version.

MonStatus Indicates the internal monitor state. The values are: Enabled or Disabled, or Unknown.

Note: You can correct the Overrun condition by recycling the ICSF subsystem.

ORIGINNODE The z/OS operating system in your enterprise monitored by a Tivoli OMEGAMON XE forz/OS agent from which the data is derived.

PCIStatus Indicates the status of PCI coprocessors. The values are: Active, Online, Present, or None.

PKACall Indicates whether Public Key Algorithm callable services are enabled. The values are: Enabled, Disabled, Unknown.

PKAMKeys Indicates whether the Public Key Algorithm Master Keys are valid. The values are: Valid, Invalid, Unknown.

PKDSname The Public Key Dataset name.

PKDSRead Indicates whether Public Key Dataset read access is enabled. The values are: Enabled, Disabled, or Unknown.

PKDSWrite Indicates whether Public Key Dataset write access is enabled. The values are: Enabled, Disabled, or Unknown.

PRSM Indicates whether the coprocessors are operating in a PR/SM configuration. The values are: Yes, No, or Unknown. PR/SM stands for Processor Resource/System Manager and is a function that allows the processor unit to operate several system control programs simultaneously in LPAR mode.

SCEDisabled The number of service call exits disabled due to a KCGSEXIT ABEND. If this value is 0, all collector exits are operational.

SMFID The z/OS system associated with the ICSF subsystem executing.

SMK_CMOS0 Indicates the state of the Public Key Algorithm, Signature Master Key in CMOS coprocessor C0. The values are: Valid, Reset, or Unknown.

SMK_CMOS1 Indicates the state of the Public Key Algorithm, Signature Master Key in CMOS coprocessor C1. The values are: Valid, Reset, or Unknown.

SMKey Is the Public Key Authentication Signature Master Key hash pattern.

SSMODE Indicates whether Special Secure Mode is enabled. The values are: Enabled, Disabled, or Unknown.

Status Indicates the status of the ICSF subsystem. The values are: Active, Inactive, Not_Found, Initializing, or Terminating. .

Version Is the ICSF subsystem version and release level.

WLDSname Is the CICS wait list dataset name.

KM5 Spin Lock attributes

These attributes collect spin lock data retrieved from Resource Management Facility (RMF). For data to be available for these attributes, OMEGAMON XE on z/OS must be configured to use RMF lock data and RMF must be configured to collect lock data. For other dependencies, see the prerequisites section of Attributes topic.

ASID The unique system-assigned identifier for the address space in which the job is running. This value is displayed only when Mode is blank.

Held Percent The percentage of samples where the address space has been found holding the lock. Values are reported to two decimal places of precision.

Job Name The name of the address space spinning due to lock request. This value is displayed only when Mode is blank.

Managed System The name of the managed system from which this data is collected. The name has the format: `<plexname>:<smfid>:MVSSYS`, where `<plexname>` is the sysplex this LPAR resides in, and `<smfid>` is the system's SMF identifier.

MODE

Mode Indicates whether the lock held exclusively or shared. Valid values are:

- Exclusive
- Shared
- Warning

The Warning mode is used only when RMF Status is an exception state (2 or 3).

Requesting Address The instruction address where the lock was requested. This value is displayed only when Mode is blank.

Resource Name The resource name of the spin lock.

RMFSTAT

RMF Status Indicates the availability of Resource Management Facility (RMF) and lock data. The valid values are:

- 1= Spin_lock_active
- 2=No_spin_lock_activity or Monitor_III_set_to_ NOLOCK
- 3=RMF_DDS_not_available

Sample Period Sample period, in seconds.

Spin Percent The percentage of samples, to two decimal places, where the address space has been found spinning due to a lock request. This value is displayed only when Mode is blank.

KM5 Suspend Lock attributes

This attribute group collects suspend lock data retrieved from Resource Management Facility (RMF). For data to be available for these attributes, OMEGAMON XE on z/OS must be configured to use RMF lock data and RMF must be configured to collect lock data. For other dependencies, see the prerequisites section of Attributes topic.

ASID The unique system-assigned identifier for the address space in which the job is running.

Dispatchable Percent The percentage of samples where the address space has been found dispatchable while holding the lock. Values are reported to two decimal places of precision.

Held Percent The percentage of samples where the address space has been found holding the lock. Values are reported to two decimal places of precision.

INTERPCT

Interrupted Percent The percentage of samples where the address space has been found interrupted while holding the lock. Values are reported to two decimal places of precision.

Job Name The name of the address space that holds the lock.

Lock Type The type of lock: Global or Local.

Managed System The name of the managed system from which this data is collected. The name has the format: `<plexname>:<smfid>:MVSSYS`, where `<plexname>` is the sysplex this LPAR resides in, and `<smfid>` is the system's SMF identifier.

Requesting Address The instruction address where the lock was obtained.

Resource Name The resource name of the suspend lock. For Local locks, this is the name of the home address space, which is equal to the holder's jobname. For CML locks, this is the name of the primary address space, which is different from the holder's name.

RMF Status Indicates the availability of RMF and lock data. Possible values are:
- Suspend_lock_active
- No_suspend_lock_activity
- Monitor_III_set_to_NOLOC
- RMF_DDS_not_available

Sample Period Sample period, in seconds.

SUSPPCT

Suspend Percent The percentage of samples where the address space has been found suspended while holding the lock. Values are reported to two decimal places of precision.

KM5 zFS Directory Cache attributes

This attribute group is used to examine the numbers of requests, hits, and discards from the directory cache to determine how well the cache is performing.

Buffer Count Number of buffers in the cache.

Buffer Size Size of each buffer, in kilobytes.

Cache Discards Discards of data from the cache.

Cache Hits Hits in the cache.

RATIO

Cache Hits Ratio Cache hit ratio.

Cache Requests Requests to the cache.

Collection Interval Time between collections, in seconds.

Managed System The z/OS operating system in your enterprise monitored by a Tivoli OMEGAMON XE on z/OS agent from which data is derived.

SMFID z/OS system producing this data.

System Name System name reporting zFS statistics.

Timestamp Timestamp showing the date and time of day when data was recorded.

Total KBytes Total kilobytes in cache.

zFS Address Space Name of the address space running zFS. OMEGAMON XE on z/OS uses an address space name of ZFS, unless the parameter KM3KZFSASNM=*xxxxxxxx* (where *xxxxxxxx* is the started task (STC) name of the zFS address space) has been added to the &*rhilev*. &*rte*.RKANPARU(KDSENV). This field is blank if zFS is not implemented (that is, there is no FILESYSTYPE TYPE(ZFS) specified in SYS1.PARMLIB(BPXPRM*xx*).

KM5 zFS Kernel attributes

This attribute group is used to determine the number of kernel operations and average time for the operations. A longer average may indicate that a larger user, metadata, or directory cache is needed.

Average Wait Time Average wait time for this operation, in milliseconds.

Collection Interval The time between collections, in seconds.

Managed System A z/OS operating system in your enterprise that is being monitored by Tivoli OMEGAMON XE on z/OS. Valid value is a string of up to 32 characters, in the format <plexname>:<smfid>:MVSSYS.

Operation Count Number of times the kernel operation was called.

Operation Name The name of the kernel operation.

Operation Rate The operation rate over the collection interval.

Record Type The type of record, either detail or summary.

SMFID z/OS system producing this data.

System Name The name of the system that is reporting zFS statistics.

Timestamp Timestamp showing the date and time of day when data was recorded.

Total Operations Grand total of operations.

Total Wait Time Wait time for all operations, in milliseconds.

zFS Address Space The name of the address space running zFS.

KM5 zFS Metadata Cache attributes

This attribute group is used to monitor the performance of the metadata cache. Metadata includes control structures and directory contents and exists in the zFS address space.

Backing Buffer Size Size of each buffer in kilobytes.

Backing Buffers Number of buffers in backing cache.

Backing Discards Discards of data from the cache.

Backing Hits Total backing hits in the cache.

Backing KBytes Size of backing cache, in kilobytes.

Backing Hit Ratio Backing cache hit ratio.

Back Requests Requests to the cache.

Collection Interval Time between collections, in seconds.

Managed System The z/OS operating system in your enterprise monitored by a Tivoli OMEGAMON XE on z/OS agent from which data is derived.

Primary Buffer Size Size of each buffer in K bytes.

Primary Buffers Number of buffers in primary cache.

Primary Hit Ratio Primary hits ratio.

Primary Hits Hits in the cache.

Primary KBytes Size of primary cache in kilobytes.

Primary Requests Requests to the cache.

Primary Updates Updates to buffers in the cache.

SMFID z/OS system producing this data.

System Name System name reporting zFS statistics.

Timestamp Timestamp showing the date and time of day when data was recorded.

zFS Address Space Name of the address space running zFS.

KM5 zFS Storage attributes

This attribute group provides a breakdown of zFS virtual storage usage.

Collection Interval Time between collections, in seconds .

Component Alloc Req Allocation requests for this component.

Component Bytes Bytes allocated for this component.

Component Desc Component description.

Component Free Req Free requests for this component.

Component Pieces Pieces allocated for this component.

Managed System The z/OS operating system in your enterprise monitored by a Tivoli OMEGAMON XE on z/OS agent from which data is derived.

SMFID z/OS system producing this data.

System Name System name reporting zFS statistics.

Timestamp Timestamp showing the date and time of day when data was recorded.

Total Alloc Requests Total allocation requests across all components.

Total Bytes Total bytes allocated across all components.

Total Free Requests Total free requests across all components.

Total Pieces Total pieces allocated across all components.

zFS Address Space Name of the address space running zFS.

KM5 zFS User Cache attributes

This attribute group is used to monitor the performance of the User File Cache. The User File Cache is a Least Recently Used (LRU) cache which holds only the most recently accessed blocks of files which are cached. The KM5 zFS User Cache attributes return two sets of statistics, direct and client.

Async Reads Number of asynchronous reads.

Cache Page Number of pages in cache.

Collection Interval The time between collections, in seconds.

Dataspace Count Number of dataspaces.

```
ERRORWAIT
```

Error Waits Number of error waits.

Error Writes Number of error writes.

Flushes Number of cache flush calls.

Free Pages Number of free pages.

Fsync Waits Number of fsync waits.

Fsyncs Number of fsync calls.

Getattrs Number of getattr calls.

Local Segment Size Local segment size in kilobytes.

Managed System The z/OS operating system in your enterprise monitored by a Tivoli OMEGAMON XE on z/OS agent from which data is derived.

Page Size Size of a page in kilobytes.

Pages Per Dataspace Number of pages in a dataspace.

Read Fault Ratio Read fault ratio.

Read IOs Number of read I/Os.

Read Wait Ratio Read wait ratio.

Read Waits Number of read waits.

Reads Number of read calls.

Reads Faulted Number of read faults.

Reclaim Steals Number of reclaim steals.

Reclaim Waits Number of reclaim waits.

Reclaim Writes Number of reclaim writes.

Remote Segment Size Remote segment size in kilobytes.

Request Type The request type, direct or client.

Scheduled Deletes Number of scheduled deletes.

Scheduled Writes Number of scheduled writes.

Schedules Number of schedule calls.

Setattrs Number of setattr calls.

SMFID z/OS system producing this data.

System Name The system name reporting zFS statistics.

Timestamp Timestamp showing the date and time of day when data was recorded.

Unmaps Number of unmap calls.

VM Segtable Cachesize Table size of the VM segment.

Waits For Reclaim Number of waits for reclaim.

Write Fault Ratio Write fault ratio.

Write Wait Ratio Write wait ratio.

Write Waits Number of write waits.

Writes Number of write calls.

Writes Faults Number of write faults.

zFS Address Space Name of the address space running zFS.

KM5 zFS User Cache Dataspaces attributes

This attribute group contains metrics returned by a request for User Cache statistics and shows the number of dataspaces used to manage the User Cache instances.

Cache Pages Number of pages in cache.

Collection Interval Time between collections, in seconds.

Dataspace Alloc Segs Number of allocated segments.

Dataspace Count Number of dataspaces.

Dataspace Free Pages Number of free pages in cache.

Dataspace Name Dataspace name.

Free Pages Number of free pages.

Local Segment Size Local segment size in K.

Managed System The z/OS operating system in your enterprise monitored by a Tivoli OMEGAMON XE on z/OS agent from which data is derived.

Page Size Size of a page in kilobytes.

Pages Per Dataspace Number of pages per dataspace.

Remote Segment Size Remote segment size in kilobytes.

SMFID z/OS system producing this data.

System Name System name reporting zFS statistics.

Timestamp Timestamp showing the date and time of day when data was recorded.

VM Segtable Cachesize Table size of the VM segment.

zFS Address Space Name of the address space running zFS.

LPAR Clusters attributes

The LPAR Cluster attributes permit you to view information about, or create situations that report on, the status of various resources belonging to an LPAR Cluster.

Capping Percent The defined soft capping percent. Valid value is a 4-byte integer in the range 0.0 to 100.0.

Capping Status The status of LPAR capping for this LPAR as defined by the LPAR configuration. LPAR capping ensures that a logical processor's use of the physical CPs cannot exceed a specified amount. Valid values are:

- HARD - Traditional LPAR capping as specified in the HMC profile
- SOFT - Software license capping
- NONE - No capping

LPMDCBU

CBU Adjustments Indicates if Capacity Backup (CBU) adjustments have been made.

Cluster LPARs In a Cluster summary data row, the number of LPARs that are assigned to the LPAR Cluster. In an LPAR data row, this number is 0. Valid value is a 4-byte integer.

Cluster Name The name of the LPAR cluster. The format of this attribute is xxxxxxx.ssss.tttt, where:

xxxxxxx	specifies the sysplex name
ssss	are the last 4 characters of the Central Processing Complex (CPC) Serial Number
tttt	are the first 4 characters of the CPU Model number

CP Online Time Logical processor online time (in seconds). Valid value is a 4-byte integer.

CP Status LPAR processor(s) status flag. Valid values are:

- DED - Dedicated
- SHR - Shared
- MIX - Both shared and Dedicated

CPs Dispatch Time The CPC total CP-type processor dispatch time. Valid value is a 4-byte integer.

CPC CPU Overhead The total CP-type processor utilization for LPAR management overhead. Valid value is a 4-byte integer in the range 0.0 to 100.0.

CPC CPU Percent The total CPC CP-type processor utilization. Valid value is a 4-byte integer in the range 0.0 to 100.0.

CPC Model The model number of the CPU where the LPAR is running; for example, 2064-103. Valid value is an 8-byte character string.

CPC MSUs The total CPC CP-type capacity in millions of service units (MSUs) per hour. Valid value is a 4-byte integer.

CPC Serial # The 6-character CPC Serial Number; for example, 010567. Valid value is an 8-byte character string.

CPC Storage Total Central Storage (in megabytes) assigned to LPARs that are sharing or using CP-type processors. Valid value is an 8-byte character string.

CPC Total Weight Sum of the weights assigned to LPARs that are active and sharing CP-type processors on the CPC. Valid value is a 4-byte integer.

CPU Index Ratio of actual physical CPU percent (Physical%CPU) utilization to target (Physical%Weight). Valid value is a 4-byte integer in the range 0.0 to 100.0.

CPU Ready Percent Percent of time the LPAR had available (ready) work and was not dispatched. Valid value is a 4-byte integer in the range 0.0 through 100.0.

Current Weight Current shared processor weight assigned to the LPAR or Cluster. In a shared physical processor configuration, weight determines the relative importance of the LPAR for the allocation of processor resources. Current Weight is that weight currently assigned for an LPAR that is enabled for Workload Manager LPAR CPU Weight Management.

Dispatch Time Logical processor dispatch time (in seconds). Valid value is a 4-byte integer.

Effective Weight Index The ratio (that is, velocity index) of actual logical LPAR weight (Effective%Weight) to target logical LPAR weight (Logical%Weight). Valid value is a 4-byte integer.

Effective Weight Percent Actual Logical Weight Percent for the LPAR (that is, velocity), a measure of how well the LPAR is able to obtain processor resources. Valid value is a 4-byte integer in the range 0.0 through 100.0.

Host LPAR Flag An indicator that identifies the Host LPAR data row. Valid values are Y or N.

Host LPAR Name Name of the LPAR for this host system (that is, the system where the data was collected). Valid value is an 8-byte character string.

Initial Weight The shared processor weight assigned to the LPAR at startup. The attribute will not contain a value if the LPAR is assigned dedicated CPs.

Interval Time Data collection time interval duration (in seconds). Valid value is a 4-byte integer.

LCPs Offline The number of LCPs currently offline that were online to this LPAR and have been brought offline by IRD or the operator.

LCPs Online The number of logical processors (LCPs) currently online to this LPAR.

LCPs Reserved The number of configured logical processors that cannot be varied online to this LPAR. That is, a configuration change such as Capacity Upgrade on Demand would be required to make this processor available. Valid value is a 4-byte integer.

LCPs Standby The number of logical processors that are not online to the LPAR, but could be varied online via operator intervention. Valid value is a 4-byte integer.

Logical CPU Percent The percent of time that the LPAR was utilizing online logical processors. Valid value is a 4-byte integer in the range 0.0 through 100.0.

Logical Weight Percent The demand that the LPAR could place on its online logical processors relative to the LPAR's defined (physical) weight. Valid value is a 4-byte integer in the range 0.0 through 100.0.

LPAR Name Name of the LPAR for this row of data. A value of _CPTotal indicates that this is a CPC summary data row. A value of _CLTotal indicates that this is a cluster summary data row. Valid value is an 8-byte character string.

LPAR Number LPAR number. Valid value is a 4-byte integer.

LPAR SMFID LPAR SMF ID as defined in the IEASYM or IEASYS PARMLIB members. This field is defined for future use and currently will contain no value.

LPAR Status Status of the LPAR. Valid values are ACTIVE or DEACTIVATED.

Managed System The z/OS operating system in your enterprise that is being monitoring by a Tivoli OMEGAMON XE on z/OS agent. Valid value is a character string with a maximum length of 32 bytes.

Maximum Weight The defined maximum shared processor weight that Workload Manager LPAR Weight Management can assign to the LPAR. A value of 0 indicates that IRD is to determine the value. The attribute will not contain a value if the LPAR is assigned dedicated CPs.

Minimum Weight The defined minimum shared processor weight that Workload Manager LPAR Weight Management can assign to the LPAR. A value of 0 indicates that IRD is to determine the value. The attribute will not contain a value if the LPAR is assigned dedicated CPs.

Model Capacity ID The processor model capacity ID representing the service the CPC is currently capable of delivering.

`LPMDMSU`

Model Capacity Rating The processor model capacity, in millions of service units (MSUs), that the CPC is currently capable of delivering.

Model Permanent Capacity ID The processor model capacity ID representing the service the CPC is capable of delivering, exclusive of capacity temporarily added through On/Off Capacity on Demand (OOCoD) and capacity temporarily replaced through Capacity Backup (CBU). Valid value is a character string with a maximum length of 16 bytes. This value appears as Unavailable if the hardware does not support Capacity provisioning management (CPM).

Model Permanent Capacity Rating The processor model capacity, in millions of service units (MSUs), that the CPC is capable of delivering, exclusive of capacity temporarily added through On/Off Capacity on Demand (OOCoD) and capacity temporarily replaced through Capacity Backup (CBU). This value appears as Unavailable if the hardware does not support Capacity provisioning management (CPM).

Model Temporary Capacity ID The processor model capacity ID representing the service that the CPC is capable of delivering by adding the Model Permanent Capacity and capacity added through On/Off Capacity on Demand (OOCoD), but excluding any capacity added through Capacity Backup (CBU). This value appears as Unavailable if the hardware does not support Capacity provisioning management (CPM).

LPMDMSUT

Model Temporary Capacity Rating The processor model capacity, in millions of service units (MSUs), that the CPC is capable of delivering by adding the Model Permanent Capacity and capacity added through On/Off Capacity on Demand (OOCoD) but excluding any capacity added through Capacity Backup (CBU). This value appears as Unavailable if the hardware does not support Capacity provisioning management (CPM).

LPMDOCOD

OOCoD/CPM Adjustments Indicates whether On/Off Capacity on Demand (OOCoD) or Capacity Provisioning Manager (CPM) adjustments have been made (Yes or No).

Overhead CPU Percent Percent of total physical processor resource utilized by the LPAR or Cluster for LPAR Management. Valid value is a 4-byte integer in the range 0.0 through 100.0.

Overhead Time The amount of time utilized for LPAR management (in seconds). Valid value is a 4-byte integer.

Physical CPs The number of physical CP-type processors that are online for this Central Processing Complex (CPC).

Physical CPU Percent Percent of total physical processor resource utilized by the LPAR or cluster, excluding LPAR Management utilization. Valid value is a 4-byte integer in the range 0.0 through 100.0.

Physical Weight Percent Current shared processor weight for the LPAR or Cluster as a percentage of total CPC weight. Valid value is a 4-byte integer in the range 0.0 to 100.0.

Ready Time Logical processor ready time (in seconds). This value represents the amount of time that the LPAR had available (ready) work and was not dispatched. Valid value is a 4-byte integer.

Special CPs The number of ICF/IFL/IFA/zIIP-type physical processors that are online for this Central Processing Complex (CPC). This count includes processors of the following types:

- Integrated Coupling Facility (ICF)
- Integrated Facility for Linux (IFL)
- zSeries Application Assist Processor (zAAP), also known as Integrated Facility for Applications (IFA)
- System z9 Integrated Information Processor (zIIP)

Storage (Meg) Amount of central storage, in megabytes, assigned to this LPAR. Valid value is a 4-byte integer.

Wait Completion LPAR processor(s) wait completion flag. Valid values are NO, YES, or MIX.

Wait Time Logical processor wait time (in seconds). This value represents the amount of time the LPAR was idle (that is, having no work to do; PSW Wait bit on). Valid value is a 4-byte integer.

WLM CPU Flag Indicates whether the LPAR is enabled for Intelligent Resource Director (IRD) WLM CPU management. Valid values are YES or NO.

Operator Alerts attributes

The Operator Alerts® Attribute Group identifies the various types of operator alerts that can be issued by Tivoli OMEGAMON XE for z/OS.

ASVT Slot Utilization A percentage that represents the maximum number of address space vector table slots that are in use or unavailable. ASVT slots contain the list of address spaces and are used to allocate address space control blocks (ASCBs). The system cannot create new address spaces when it runs out of slots. Valid value is a 4-byte integer.

GRS Status Status of the GRS configuration and connection. Valid values are BROKEN, RING, or STAR.

- BROKEN indicates the configuration is broken.
- RING indicates the configuration of the GRS is a ring.
- STAR indicates the configuration of the GRS is a star.

GTF Active Whether or not GTF is active. Valid values are TRUE or FALSE.

HSM Recall Wait Time The wait time in seconds of the longest single HSM recall that is waiting. Valid value is a 4-byte integer.

Managed System The z/OS operating system in your enterprise that is being monitoring by an OMEGAMON XE on z/OS monitoring agent. Valid value is a character string with a maximum length of 32 bytes.

OLTEP Active Whether or not OLTEP is active. Valid values are TRUE or FALSE. OLTEP is used primarily by your hardware representative to run diagnostics on hardware devices. It can drastically affect the performance of your system because of its high resource utilization.

Outstanding Operator Replies Count of outstanding Write to Operator with Reply operations. Valid value is a 4-byte integer.

RMF Not Active Whether or not RMF is active. Valid values are TRUE or FALSE. The RMF Monitor I is a background subtask of the RMF address space. It collects performance information such as resource utilization and response time data and can be used to help identify performance problems.

SMF Not Recording Whether or not SMF is recording information. Valid values are TRUE or FALSE.

SYSLOG Not Recording Whether or not the SYSLOG is recording information. Valid values are TRUE or FALSE. The SYSLOG records all console activity, providing an information trail that can be used to research problems or verify that jobs have run in the proper sequence.

WTO Buffers Remaining The count of remaining WTO buffers. You can use the z/OS DISPLAY CONSOLES command to list all of the defined consoles and the number of WTO buffers each one is holding. Note: If you run out of WTO buffers, the Master console may become locked out, causing the system to hang. Valid value is a 4-byte integer.

Page Dataset Activity attributes

The Page Dataset Activity attribute group provides information about availability and response time for specific page dataset.

Address The hexadecimal device address for this page dataset.

Dataset Type The type of dataset. Valid values are:

- Unknown
- Local — Page datasets that back up the private area of all address spaces and the Virtual I/O (VIO) pages of the system
- Common — Page dataset containing CSA, MLPA, and SQA (that is, containing the common area that is other than PLPA)
- PLPA — Page dataset containing only frames of Pageable Link Pack Area
- Swap — Page datasets containing the swapped out part of an address space

Errors The number of errors for this dataset. Valid value is a 4-byte integer.

Managed System A z/OS operating system in your enterprise monitored by a Tivoli OMEGAMON XE on z/OS agent. Valid value is a character string with a maximum length of 32 bytes.

Page Rate The pages per second read or written for this dataset. Valid value is a 4-byte integer.

Percent Full The percentage of slots on this page (or swap) dataset that are in use. Valid value is a 4-byte integer in the range 0.0 to 100.0. In this field, the value 123 would represent 12.3%.

Response Time The average response time in milliseconds for an I/O operation to this dataset. Valid value is a 4-byte integer.

Volume The volume serial number of this page dataset. Valid value is a 6-byte character string.

Real Storage attributes

The Real Storage Attribute Group provides information about the use of real storage on a z/OS system in terms of various types of frame counts, slot counts, and paging rates.

Note: In version 4.1, several attributes used in previous versions have been deprecated. New values have been added to others. For information about upgrading situations and workspaces using these attributes, see the *IBM Tivoli OMEGAMON XE on z/OS: User's Guide.*

Area End Address The ending address of an area, in hexadecimal. The end address of the high private area (above the bar area) is shown as 16EB. For a Storage Type of Extended Nucleus, the Area End Address is the higher of the end address of the Extended Read Only Nucleus and the Extended Read Write Nucleus areas. For a Storage Type of Nucleus, Area End Address is the higher of the end address of the Read Only Nucleus and the Read Write Nucleus areas.

Area Start Address The starting address of an area, in hexadecimal. The starting address of the high private area (above the bar area) is shown as 2G. 6. For a Storage Type of Extended Nucleus, Area Start Address is the lower of the start address of the Extended Read Only Nucleus and the Extended Read Write Nucleus areas. For a Storage Type of Nucleus, Area Start Address is the lower of the start address of the Read Only Nucleus and the Read Write Nucleus areas.

Available Frames The number of pageable frames available to the system. Valid value is a 4-byte integer.

Managed System A z/OS operating system in your enterprise that Tivoli Enterprise Portal is monitoring. Valid value is a character string with a maximum length of 32 bytes.

Migration Age The time in seconds that has passed since the oldest frame of expanded storage was last referenced. This column is applicable to expanded storage, but not to central storage. Valid value is a 4-byte integer. This attribute is deprecated in version 4.1.

Migration Rate The number of pages per second that are being moved from expanded to auxiliary storage. This column is applicable to expanded storage, but not to central storage. Valid value is a 4-byte integer to one decimal precision. This attribute is deprecated in version 4.1.

Offline Frames The number of frames of this type that have been varied offline either by means of a z/OS command or by the system itself. Valid value is a 4-byte integer.

Online Frames The number of storage frames of this type that are online or accessible to this processor. Valid value is a 4-byte integer.

Pages Read From Expanded The number of pages per second that are being moved from expanded to central storage. Valid value is a 4-byte integer. This attribute is deprecated in version 4.1.

Pages Written To Expanded The number of pages per second that are being moved from central to expanded storage. Valid value is a 4-byte integer. This attribute is deprecated in version 4.1.

Storage Type Type of storage area. Valid values are:

- Summary
- Shared
- Data Spaces
- Available
- Other

  where "Other" is the sum of
  - Top double frames
  - Bottom double frames
  - SQA reserved frames
  - Sdump buffer frames
  - V=R waiting frames
  - Deferred freemain frames
  - IDA frames
  - HSA frames
  - Flawed frames
  - Unitialized frames
  - VIO frames
  - Local quad frames
  - Page table frames
  - Unknown frame types
  - PSA frames
  - Nucleus DAT off frames
- Extended private
- Extended CSA
- Extended PLPA
- Extended MLPA
- Extended FLPA
- Extended SQA
- Extended NUC

  where "Extended NUC" is the sum of
  - Extended Read Only nucleus

  - Extended Read/Write nucleus
- NUC

  where "NUC" is the sum of
  - Read Only nucleus
  - Read/Write nucleus
- SQA
- PLPA
- MLPA
- FLPA
- CSA
- Private
- High Private (above 2G)

Total Fixed Frames Total number of fixed frames for this image. The Total Fixed Frames value is the sum of the **Total Fixed Frames Allocated** column for all the storage areas listed in the Real Storage view. Valid value is a 4-byte integer.

Total Fixed Frames Allocated Number of fixed frames allocated to an area. Valid value is a 4-byte integer.

Total Frames Total number of frames for this image. The Total Frames value is the sum of the **Total Frames Allocated** column for all the storage areas listed in the Real Central Storage view. Valid value is a 4-byte integer.

Total Frames Allocated The total number of frames allocated to an area. Valid value is a 4-byte integer.

Total Kilobytes Allocated The size of the area, in kilobytes. Valid value is a 4-byte integer.

Unreferenced Interval Count Represents the time in seconds for a complete steal cycle. A complete steal cycle is the time the stealing routine needs to check all frames in the system. When there is a demand for storage, the stealing routine tests the reference bit of a frame, decides whether to steal the frame, and schedules the page-out. Valid value is a 4-byte integer.

Service Call Type Detail (SVCDET) attributes

The Service Call Type Detail monitor performance by service call type detail. The agent monitors performance by collecting data at the beginning and end of each service call.

ArrivalRate The average arrival rate (moving average) in service calls per minute. The averages are computed for the last ten minutes of observation. They are read three places to the right of the decimal. Valid values are 0-99999999.

Bytes The average number of bytes (moving average) processed per service call. The averages are computed for the last ten minutes of observation. They are read three places to the right of the decimal. Valid values are 0-99999999.

ORIGINNODE The managed system name of the z/OS system being monitored by a Tivoli OMEGAMON XE on z/OS agent.

Pending The average calls pending (rolling average) per service call. The averages are computed for the last ten minutes of observation. The calls are read three places to the right of the decimal. Valid values are: 0-10000.

SvcCall The service call name or entry point. Valid values are:

- CSFAEGN
- CSFAKEX
- CSFAKIM
- CSFAKTR
- CSFATKN
- CSFCKI
- CSFCPE
- CSFPGN
- CSFCPA
- CSFEDC
- CSFCTT
- CSFCTT1
- CSFCVT
- CSFCVE
- CSFDKX
- CSFDKM
- CSFDEC
- CSFDEC1
- CSFDCO
- CSFDSG
- CSFDSV
- CSFDKG
- CSFEMK
- CSFENC
- CSFENC1
- CSFECO
- CSFEPG
- CSFPTR
- CSFPVR
- CSFGKC
- CSFRTC
- CSFKEX
- CSFKGN
- CSFKIM
- CSFKPI
- CSFKRC
- CSFKRD
- CSFKRR
- CSFKRW
- CSFKYT
- CSFKYTX
- CSFKTR
- CSFMGN
- CSFMGN1
- CSFMVR
- CSFMVR1
- CSFMDG
- CSFMDG1
- CSFCKM
- CSFSKM
- CSFOWH
- CSFOWH1
- CSFPCI
- CSFPKD
- CSFPKE
- CSFPKG

- CSFPKI
- CSFPKX
- CSFPKRC
- CSFPKRD
- CSFPKRR
- CSFPKRW
- CSFPKSC
- CSFPEX
- CSFPEXX
- CSFRNG
- CSFRKD
- CSFRKL
- CSFSKI
- CSFSBC
- CSFSBD
- CSFSYX
- CSFSYG
- CSFSYI
- CSFTCK
- CSFUDK
- CSFCSG
- CSFCSV

SvcDesc The service call description. Valid values are:

- ANSI X9.17 EDC Generate
- ANSI X9.17 Key Export
- ANSI X9.17 Key Import
- ANSI X9.17 Key Translate
- ANSI X9.17 Transport Key Partial Notarize
- Clear Key Import
- Clear PIN Encrypt
- Clear PIN Generate
- Clear PIN Generate Alternate
- Cipher/Decipher
- Ciphertext Translate
- Ciphertext Translate (with ALET)
- Control Vector Translate
- Cryptographic Variable Encipher
- Data Key Export
- Data Key Import
- Decipher
- Decipher (with ALET)
- Decode
- Digital Signature Generate
- Digital Signature Verify
- Diversified Key Generate
- Encipher under Master Key
- Encipher
- Encipher (with ALET)

- Encode
- Encrypted PIN Generate
- Encrypted PIN Translate
- Encrypted PIN Verify
- Generate a key
- Import a key
- Key Export
- Key Generate
- Key Import
- Key Part Import
- Key Record Create
- Key Record Delete
- Key Record Read
- Key Record Write
- Key Test
- Key Test Extended
- Key Translate
- MAC Generate
- MAC Generate (with ALET)
- MAC Verify
- MAC Verify (with ALET)
- MDC Generate
- MDC Generate (with ALET)
- Multiple Clear Key Import
- Multiple Secure Key Import
- One Way Hash Generate (with ALET)
- One Way Hash Generate
- PCI Interface
- PKA Decrypt
- PKA Encrypt
- PKA Key Generate
- PKA Key Import
- PKA Public Key Extract
- PKDS Record Create
- PKDS Record Delete
- PKDS Record Read
- PKDS Record Write
- PKSC Interface
- Prohibit Export
- Prohibit Export Extended
- Random Number Generate
- Retained Key Delete
- Retained Key List
- Secure Key Import
- SET Block Compose

- SET Block Decompose
- Symmetric Key Export
- Symmetric Key Generate
- Symmetric Key Import
- User Derived Key
- Transform CDMF Key
- VISA CVV Service Generate
- VISA CVV Service Verify

SvcTime The average service call completion time (moving average) in milliseconds per call. The averages are computed for the last ten minutes of observation. They are read three places to the right of the decimal.

SMFID The z/OS system associated with the service call.

System CPU Utilization attributes

The System CPU Utilization Attribute Group provides information about CPU usage for the system. In an LPAR environment, this group also provides partition management statistics.

Note: All the percentage calculations in this workspace are unnormalized values. If an address space uses more than a single processor worth of CPU, the percentage calculation may exceed 100%. The simple TCB % is a likely candidate for exhibiting this, since its denominator is only about 2.3 seconds.

4 Hour MSUs Four-hour rolling average value of the millions of service units (MSUs) per hour. The value is updated every five minutes. It is the average of the previous 48 5-minute samples, which in turn are the average of 30 samples, taken every 10 secs. To avoid capping during system bringup (IPL), the calculation of the 4-hour rolling average assumes that the system has run for four hours without load, so for the first four hours after IPL the rolling average is not valid. Valid value is an integer in the range 0 to 2147483647.

Note: Data for this attribute is available only if you have opted to use a defined capacity limit, otherwise Not available is displayed.

Average CPU Percent The percentage of time that all processors available in this z/OS system were busy dispatching work. If an address space uses more than a single processor worth of CPU, the percentage may exceed 100%.

Average IFA Percent Average percentage of time the system consumes zSeries Application Assist Processor (zAAP) resources across all zAAPs configured to this LPAR. If an address space uses more than a single processor worth of CPU, the percentage may exceed 100%.

Average IFA Percent on CP Average percentage of time the system consumes regular processor resource executing zSeries Application Assist Processor (zAAP) work across all regular processors configured to this logical partition (LPAR). If an address space uses more than a single processor's worth of CPU, the percentage may exceed 100%.

Average zIIP Percent Average percentage of System z9 Integrated Information Processor (zIIP) time the system consumed for all zIIPs configured to this LPAR.

Average zIIP Percent on CP Average percentage of regular processor time the system consumes performing work eligible for System z9 Integrated Information Processor work for all regular processors configured to this LPAR.

CPU Flag A flag that indicates how the Average CPU Percent value is calculated. Valid values can be:

RMF	As reported by Resource Management Facility (RMF), this value is represented as a running average for this partition over the RMF interval.
Native	This value applies to systems running in native mode or in PR/SM or Amdahl Multiple domain Facility (MDF) environments, and is based on hardware timers. It is calculated by adding TCB, SRB, and z/OS overhead times.

CPUs Offline A string of CPU sequence numbers, separated by commas, that identifies those processors that are offline to this LPAR. Valid value is a 32-byte character string.

CPUs Online A string of processor sequence numbers, separated by commas, that identifies those processors that are online to this LPAR. Valid value is a 32-byte character string.

HiperDispatch Management Status of the HiperDispatch feature. The possible values are: On, Off, and Unavailable.

IFA Crossover Indicates whether IFA crossover of IFA work to regular processors is allowed (Yes) or not allowed (No).

IFA Honor Priority Indicates whether IFA work that has crossed over to a regular processor has its dispatch priority honored (Yes) or not (No).

IFA Relative Processor Speed Speed of IFAs relative to regular processors. Expressed as a ratio in the format *nn* .*nn* :1.

IFAs Offline IFAs offline to this LPAR, as a string of comma-delimited hexadecimal 2-digit processor IDs.

IFAs Online IFAs online to this LPAR as a string of comma-delimited hexadecimal 2-digit processor IDs.

Managed System The z/OS operating system in your enterprise that is being monitored by an OMEGAMON XE on z/OS monitoring agent. Valid value is a character string with a maximum length of 32 bytes.

MVS Overhead The CPU utilization percentage that is not attributable to any user or address space. It is calculated as the difference between the total software utilization times and the total hardware time ((TCB + SRB)-CPU) over the last reporting interval. Valid value is a 4-byte integer.

Partition LCPD% In an LPAR environment, the percentage of time on average that a logical processor was dispatched in the logical partition during the interval.

This value is based on dispatch time, not CPU busy time. The value is calculated by dividing dispatch time by elapsed time and then dividing that product by the number of logical processors. Thus, if an LPAR has been defined as having two logical processors and the sum of their dispatch times was 90 seconds over a 60 second interval, the logical CPU percent for the partition would be (90 seconds / 60 seconds) / 2 which equals 75%.

When the number of logical processors assigned to a partition is the same as the number of physical processors, the partition's physical and logical CPU utilization should be the same.

Partition Overhead% In an LPAR environment, the percentage of time the system spends managing a logical partition. This column is valid for systems at z/OS SP4.2 or later. The value is calculated as the difference between the total software times (TCB and SRB) and the total hardware time (CPU) over the last reporting interval. Typically, z/OS overhead increases as workloads make greater demands on resources. Valid value is a 4-byte integer in the range 0 to 100.

In a complex with more than one CPU, z/OS overhead can be computed based on the number of processors, or normalized to a maximum of 100%.

Partition PCPD% In an LPAR environment, the percentage of time on average that a physical processor was dispatched in the logical partition during the interval. This value is based on dispatch time, not CPU busy time.

The value is calculated by dividing dispatch time by elapsed time and then dividing that product by the number of physical processors in the LPAR. Thus, if an LPAR has been defined as having four physical processors and the sum of their dispatch times was 120 seconds over a 60 second interval, the physical CPU percent for the partition would be (120 seconds / 60 seconds) / 4 which equals 50%.

When the number of logical processors assigned to a partition is the same as the number of physical processors, the partition's physical and logical CPU utilization should be the same.

Physical CPU Count A count of the physical CPUs online to this LPAR. Valid value is a 4-byte integer.

RMF LPAR CPU Percent Average system CPU percentage as seen by the LPAR hosting the z/OS system, calculated over the most recent RMF interval. The first RMF interval after the agent is started contains partial values. Valid value is an integer in the range 0.0 to 100.0. If RMF is not active, this value will be zero. "NoRMF" indicates that this value is not available because the RMF API (ERBSMFI) is not active. "NoStorage" indicates that this value is not available because the CMS address space did not have sufficient storage for the RMF API answer area.

Note: The initial display will contain zeros for the RMF LPAR CPU Percent attribute. It is most likely that the collection agent will be started somewhere in the middle of an already running RMF interval. Therefore, at the completion of the first RMF interval after the collection agent was started, these attributes will contain partial values. Only after completion of the first full RMF interval will these attributes display complete and correct values. If Resource Measurement Facility (RMF) is not started, the value for this attribute will be 0.0.

RMF MVS CPU Percent Average system CPU percentage as seen by z/OS, calculated over the most recent RMF interval. If RMF is not active, this value will be zero. The first RMF interval after the agent is started contains partial values. Valid value is an integer in the range 0.0 to 100.0. "NoRMF" indicates that this value is not available because the RMF API (ERBSMFI) is not active. "NoStorage" indicates that this value is not available because the CMS address space did not have sufficient storage for the RMF API answer area.

Note: The initial display will contain zeros for the RMF MVS CPU Percent attribute. It is most likely that the collection agent will be started somewhere in the middle of an already running RMF interval. Therefore, at the completion of the first RMF interval after the collection agent was started, these attributes will contain partial values. Only after completion of the first full RMF interval will these attributes display complete and correct values. If Resource Measurement Facility (RMF) is not started, the value for this attribute will be 0.0.

Total Enclave Percent The percentage of total CPU time spent doing enclave work. If an address space uses more than a single processor's worth of CPU, the percentage may exceed 100%.

Total SRB% The percentage of total CPU time spent doing SRB work. If an address space uses more than a single processor's worth of CPU, the percentage may exceed 100%.

Total TCB% The percentage of total CPU time spent doing TCB work. If an address space uses more than a single processor's worth of CPU, the percentage may exceed 100%.

Undispatched Tasks The number of tasks or address spaces that have not been dispatched by the SRM due to constraints. Valid value is a 4-byte integer.

WLM Mode The workload manager mode. Valid values are Goal or Compatibility.

zIIP Honor Priority Indicates whether System z9 Integrated Information Processor work that has crossed over to a regular processor has its dispatch priority honored (Yes) or not (No).

zIIP Relative Processor Speed Speed of System z9 Integrated Information Processors relative to regular processors. Expressed as a ratio in the format *nn.nn*:1.

zIIPs Offline System z9 Integrated Information Processors offline to this LPAR, as a string of comma-delimited hexadecimal 2-digit processor IDs.

zIIPs Online System z9 Integrated Information Processors online to this LPAR, as a string of comma-delimited hexadecimal 2-digit processor IDs.

System Paging Activity attributes

The System Paging Activity attribute group provides information about factors taking place in your z/OS system that affect system paging rates, including the paging rate, page fault rate, and number of I/O requests waiting for processing.

ASM Queue Length The number of I/O requests waiting to be processed by the Auxiliary Storage Manager (ASM). Valid value is a 4-byte integer.

Datasets Not Operational The number of page datasets that are not operational. This condition can occur if there are a number of I/O errors to the device or the device has not responded to an I/O request. Valid value is a 4-byte integer.

Expanded Storage Pages Moved The total number of pages moved per second to and from expanded storage. Valid value is a 4-byte integer having one decimal position.

Managed System The z/OS operating system in your enterprise that is being monitoring. Valid value is a character string with a maximum length of 32 bytes.

Page Fault Rate The number of page faults per second. A page fault is an interrupt that occurs when an executing program references an address that is not in central storage. Valid value is a 4-byte integer.

System Page Rate The rate at which paging occurs over all storage areas (common, private, and system). Valid value is a 4-byte integer having one decimal position.

Unreferenced Interval Count Represents the time in seconds for a complete steal cycle. A complete steal cycle is the time the stealing routine needs to check all frames in the system. When there is a demand for storage, the stealing routine tests the reference bit of a frame, decides whether to steal the frame, and schedules the page-out. Valid value is a 4-byte integer.

Tape Drives attributes

The Tape Drives Attribute Group provides information about various factors that affect the performance and availability of the tape drives on your system.

Address The device address. Valid value is a 2-byte integer displayed as a hexadecimal value.

DDR Swaps in Progress The number of Dynamic Device Reconfiguration (DDR) swaps in progress. DDR permits tapes to be remounted on another drive without ABENDing a running job. However, if not acted upon promptly, DDR may create severe system I/O performance degradation, since all I/O allocations must wait until the DDR swap is completed. Valid value is a 4-byte integer.

Dropped Ready The number of devices in a dropped ready condition. Valid value is a 4-byte integer.

I/O Count The count of I/O operations for this device since it was mounted. Valid value is a 4-byte integer.

Managed System The z/OS operating system in your enterprise that is being monitored by an OMEGAMON XE on z/OS monitoring agent. Valid value is a character string with a maximum length of 32 bytes.

Not Responding The number of tape drives that have not responded to an I/O request. Valid value is a 4-byte integer.

Permanent Errors The number of permanent, nonrecoverable I/O errors that have occurred for this device. Valid value is a 4-byte integer.

Status The status of the device. Valid values are:

- Unknown
- Free
- Allocated
- Mount
- Dropped
- Not Responding
- DDR Swap

Tape Mount Wait Time The amount of time the task has been waiting for a tape volume to be mounted on this device, in seconds. For multitasking jobs that run multiple threads simultaneously, this may not reflect the total amount of time the tape mount has been pending.

Tape Mounts Pending The number of tape drives waiting for a tape to be mounted. Valid value is a 4-byte integer.

Temporary Errors The number of temporary, recoverable I/O errors that have occurred on this device. Valid value is a 4-byte integer.

User The address space name that has the device allocated. If this name is not available, the column is blank. Valid value is an 8-byte character string.

Volume The volume serial number in the mounted tape label. If this number is not available, the column is blank. Valid value is a 6-byte character string.

Top User Performance attributes

The Top User Performance attributes provide information on performance by top users of cryptographic services.

ArrivalRate The average arrival rate (moving average) in service calls per minute issued by the address space. The averages are computed for the last ten minutes of observation. They are read three places to the right of the decimal. Valid values are 0-99999999.

Bytes The average number of bytes (moving average) processed per service call issued by the address space. The averages are computed for the last ten minutes of observation. They are read three places to the right of the decimal. Valid values are 0-99999999.

JOBNAME The address space name issuing cryptographic service calls.

LastSvcCall The last service call name issued by the address space.

CSFAEGN
CSFAKEX
CSFAKIM
CSFAKTR
CSFATKN
CSFCKI
CSFCPE
CSFPGN
CSFCPA
CSFEDC
CSFCTT
CSFCTT1
CSFCVT
CSFCVE
CSFDKX
CSFDKM
CSFDEC
CSFDEC1
CSFDCO
CSFDSG
CSFDSV
CSFDKG
CSFEMK
CSFENC
CSFENC1
CSFECO
CSFEPG
CSFPTR
CSFPVR
CSFGKC
CSFRTC
CSFKEX
CSFKGN
CSFKIM
CSFKPI
CSFKRC

CSFKRD
CSFKRR
CSFKRW
CSFKYT
CSFKYTX
CSFKTR
CSFMGN
CSFMGN1
CSFMVR
CSFMVR1
CSFMDG
CSFMDG1
CSFCKM
CSFSKM
CSFOWH
CSFOWH1
CSFPCI
CSFPKD
CSFPKE
CSFPKG
CSFPKI
CSFPKX
CSFPKRC
CSFPKRD
CSFPKRR
CSFPKRW
CSFPKSC
CSFPEX
CSFPEXX
CSFRNG
CSFRKD
CSFRKL
CSFSKI
CSFSBC
CSFSBD
CSFSYX
CSFSYG
CSFSYI
CSFTCK
CSFUDK
CSFCSG
CSFCSV

LastSvcDesc The service call description.

ANSI X9.17 EDC Generate
ANSI X9.17 Key Export
ANSI X9.17 Key Import
ANSI X9.17 Key Translate
ANSI X9.17 Transport Key Partial Notarize
Clear Key Import
Clear PIN Encrypt
Clear PIN Generate
Clear PIN Generate Alternate
Cipher/Decipher
Ciphertext Translate
Ciphertext Translate (with ALET)
Control Vector Translate
Cryptographic Variable Encipher
Data Key Export
Data Key Import
Decipher
Decipher (with ALET)
Decode
Digital Signature Generate
Digital Signature Verify
Diversified Key Generate
Encipher under Master Key
Encipher
Encipher (with ALET)
Encode
Encrypted PIN Generate
Encrypted PIN Translate
Encrypted PIN Verify
Generate a key
Import a key
Key Export
Key Generate
Key Import
Key Part Import
Key Record Create
Key Record Delete
Key Record Read
Key Record Write

Key Test
Key Test Extended
Key Translate
MAC Generate
MAC Generate (with ALET)
MAC Verify
MAC Verify (with ALET)
MDC Generate
MDC Generate (with ALET)
Multiple Clear Key Import
Multiple Secure Key Import
One Way Hash Generate
One Way Hash Generate (with ALET)
PCI Interface
PKA Decrypt
PKA Encrypt
PKA Key Generate
PKA Key Import
PKA Public Key Extract
PKDS Record Create
PKDS Record Delete
PKDS Record Read
PKDS Record Write
PKSC Interface
Prohibit Export
Prohibit Export Extended
Random Number Generate
Retained Key Delete
Retained Key List
Secure Key Import
SET Block Compose
SET Block Decompose
Symmetric Key Export
Symmetric Key Generate
Symmetric Key Import
Transform CDMF Key
User Derived Key
VISA CVV Service Generate
VISA CVV Service Verify

ORIGINNODE The z/OS operating system in your enterprise monitored by a Tivoli OMEGAMON XE on z/OS agent from which the data is derived.

Pending The average calls pending (moving average) per service call issued by the address space. The averages are computed for the last ten minutes of observation. They are read three places to the right of the decimal.
0-10000.

SvcTime The average service call completion time (moving average) in milliseconds per call issued by the address space. The averages are computed for the last ten minutes of observation. They are read three places to the right of the decimal.

SMFID The z/OS system associated with the address space issuing cryptographic services.

User Response Time attributes

The User Response Time Attribute Group provides an overview of the activities of your system users, as well as the average response time a user is experiencing.

Host Response The average time between the receipt of input by VTAM and the receipt of output by VTAM from the application. Valid value is a 4-byte integer representing time in seconds, to two decimal places.

Managed System An z/OS operating system in your enterprise that is being monitored by a Tivoli OMEGAMON XE on z/OS monitoring agent. Valid value is a character string with a maximum length of 32 bytes.

Network Response The average time between VTAM's sending output to the terminal and VTAM's receiving the response. Valid value is a 4-byte integer representing time in seconds, to two decimal places.

Total Response The sum of the host and network response times. Valid value is a 4-byte integer representing time in seconds, to two decimal places.

Transaction Count The total number of group transactions that ended during the RMF interval. Valid value is a 4-byte integer.

User ID The TSO user name. Valid value is an 8-byte character string.

USS Address Spaces attributes

The USS Address Spaces attributes provide information about all z/OS address spaces (so-called dubbed address spaces) that have issued a call to the UNIX System Services application programming interface (API).

A/S Name The name of the address space. For a batch job, this is the jobname from the job statement. For a TSO address space, this is the userid of the logged-on user. Tivoli OMEGAMON XE on z/OS may also generate address space names. An address space name consists of alphanumeric characters and follows the z/OS rules for address space names.

A/S Type Identifies the type of address space that is using UNIX System Services services. Valid values are JOB (batch job), STC (started task), or TSO (TSO user login).

ASID The numeric address space identifier. Its form is a positive integer in the range 0 through 65535. (This value is generally not useful when defining situations.)

Central Storage Frames Number of frames currently in use by this address space. Valid value is a 4-byte integer.

Client Seconds Service request block (SRB) time for this address space's preemptable SRBs and for the client; or related SRBs for which this address space is the client. Valid value is a 4-byte integer.

Client Time% Percent of service request block (SRB) time for this address space's preemptable SRBs and for client-related SRBs for which this address space is the client. Valid value is a 4-byte integer.

Cold Frames Number of frames in UIC interval 4 as set by the System Resources Manager (SRM) via the RCEFRV# fields. Valid value is a 4-byte integer.

Cool Frames Number of frames in UIC interval 3 as set by the System Resources Manager (SRM) via the RCEFRV# fields. Valid value is a 4-byte integer.

CPU Seconds The total CPU time of the application using services provided by UNIX System Services. (Note that this attribute is not just UNIX System Services CPU Time.) Valid values are positive integers in the range 0 through 2147483647 and can include the use of the *AVG, *MIN, *MAX, or *SUM functions.

CPU Time% Overall CPU percent (a value that represents TCB time + SRB time + Client time). Valid value is a 4-byte integer in the range 0-100.

Enclave Active Time The time the enclave has been active. Valid value is a 4-byte integer.

Enclave CPU Seconds CPU seconds used by the enclave. Valid value is a 4-byte integer.

Enclave CPU Service CPU service units received by the enclave. Valid value is a 4-byte integer.

Enclave TX Count Enclave transaction count. Valid value is a 4-byte integer.

Expanded Storage Frames Number of pages on expanded storage. Valid value is a 4-byte integer.

Home SRB Seconds CPU time for all types of preemptable SRBs (PSRB, CRSRB, ESRB) that are executing with this address space as their home address space. Valid value is a 4-byte integer.

Home SRB Time% Percent of home CPU time for all types of preemptable SRBs (PSRB, CRSRB, ESRB). Valid value is a 4-byte integer in the range 0 - 100.

Hot Frames Number of frames in UIC interval 1 as set by SRM via the RCEFRV# fields. Valid value is a 4-byte integer.

ORIGINNODE The monitored z/OS operating system on which UNIX System Services is hosted. Valid value is a character string with a maximum length of 32 bytes.

Performance Group The performance group of the address space that is using services provided by z/OS UNIX System Services. A performance group is a class specified at installation that regulates the turnaround time of a user's jobs, jobsteps, and system interactions. This attribute is valid only when the system is in Workload Manager compatibility mode. Valid value is a positive integer in the range 0 through 65535.

Service Class The service class of the address space that is using services provided by z/OS UNIX System Services. This attribute is valid only when the system is in Workload Manager goal mode. Valid value is an 8-character alphanumeric name.

SRB Seconds Accumulated SRB time. Valid value is a 4-byte integer.

SRB Time% The percentage of CPU time used by the SRB. Valid value is a 4-byte integer in the range 0 - 100.

Step Name The name used on the execute statement for the program that is using z/OS UNIX System Services. For a batch job, for example, this is the name of a cataloged procedure specified by an EXEC statement. A stepname consists of alphanumeric characters and follows the z/OS rules for job name.

TCB Seconds Elapsed seconds of TCB CPU usage. Valid value is a 4-byte integer.

TCB Time% Percent of TCB CPU usage. Valid value is a 4-byte integer in the range 0 - 100.

Total Fixed Pages Total number of fixed pages in this address space. This value does not include shared pages. Valid value is a 4-byte integer.

Total Shared Pages Total number of shared page views in this address space. Valid value is a 4-byte integer.

UNIX System Time System CPU time for UNIX work. Valid value is a 4-byte integer.

UNIX System Time% Percent of system CPU time for UNIX work. Valid value is a 4-byte integer in the range 0 - 100.

UNIX User Time User CPU time for UNIX work. Valid value is a 4-byte integer.

UNIX User Time% Percent of user CPU time for UNIX work. Valid value is a 4-byte integer in the range 0 - 100.

Userid The user ID associated with the address space. For a batch job, the value is taken from the USER= specification on the job statement. Userid consists of alphanumeric characters and follows the z/OS rules for job name.

Warm Frames Number of frames in UIC interval 2 as set by SRM via the RCEFRV# fields. Valid value is a 4-byte integer.

Working Set Total of hot plus warm plus fixed pages. Valid value is a 4-byte integer.

USS BPXPRMxx Value attributes

The USS BPXPRMxx Value attributes provide information about the BPXPRMxx members that establish defaults for UNIX System Service. These members reside in the SYS1.PARMLIB data set.

ORIGINNODE The monitored z/OS operating system on which UNIX System Services is hosted. Valid value is a character string with a maximum length of 32 bytes.

Parameter Keyword identifying a BPXPRMxx parameter. Valid value is an alphanumeric string of maximum length 64 bytes.

Source Identifies the source of a BPXPRMxx parameter. Valid value:

Member	Source is a current BPXPRM*xx* member
Active	Source is active system value
System	Source is system-generated

Value Value or values associated with a BPXPRMxx parameter. Valid value is an alphanumeric string of maximum length 1024 bytes.

USS Files attributes

The USS File Information attributes provide file, path, and time information for the files in your monitored systems.

Access Defines the access permissions associated with the file. Valid value is a positive integer that has a 3-digit octal number. Access is the numeric form of the Permissions attribute. From left to right, the digits have the following meanings:

1st	Set user ID of the file upon execution
2nd	Permissions for the file owner
3rd	Permissions for the group
4th	Permissions for other users

Blocks The number of blocks of size Blocksize used to store the data for the file. Valid value is a positive integer in the range 0 through 2147483647, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions.

Blocksize The size of a physical block of data in bytes. Valid value is a positive integer in the range 0 through 2147483647, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions.

Creation Time The date and time the directory or file was created. Valid value is a timestamp in the form mm/dd/yy hh:mm:ss.

Extended Attributes The file's extended attributes. Each character position represents one of the extended attributes as follows:

a	In the first position, indicates that the program runs as APF-authorized if linked with AC-=1.
p	In the second position, indicates the program is "program-controlled".
s	In the third position, indicates the program runs in a shared address space.
-	A dash in any of these positions indicates that the attribute is not set.

Valid value is an 8-byte character string, of which only the first three positions are used.

File The name of the file or directory. Use file names to identify a regular file, directory, special character file, pipe, link, symbolic link, or socket. Valid value is a string of 1 through 255 characters; the string may not include the slash (/) or null characters.

GID The numeric group ID of the owner of the file. Valid value is a positive integer in the range 0 through 2147483647.

Group The logical group to which the owner of the file belongs. Valid value consists of up to 8 alphanumeric characters and follows the rules for z/OS userid.

ORIGINNODE The monitored z/OS operating system on which UNIX System Services is hosted. Valid value is a character string with a maximum length of 32 bytes.

Last-Accessed Time The date and time the file was last accessed for any purpose by any user or process. Valid value is a timestamp in the form mm/dd/yy hh:mm:ss.

Last-Changed Time The date and time the file was last changed by any user or process. Valid value is a timestamp in the form mm/dd/yy hh:mm:ss.

Links The number of files that are linked to the same data. Use links information to understand how many references are made to the same file data. In order to remove the data pointed to by a file, all links to the file must be deleted as well. Valid value is a positive integer in the range 0 through 2147483647, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions.

Owner The user name or number (if the name is unknown) of the userid that created or has been given ownership of the file. Valid value consists of up to 8 alphanumeric characters and follows the rules for z/OS userid.

Path Specifies the file(s) of directory type that are selected to access a file. Use paths to organize files to more easily find information; for example, by project, type, purpose, or access permission. Valid value is a sequence of directory files, separated by the slash (/), pointing to a set of files. The length of the path is limited to 1024 characters including the slash character. An example path is `/u/IBM/Candle/OE/src`

Permissions Defines the permissions that are set for a given user or for a given file. This is a 9-character field. Permissions is the character form of the Access attribute.

To analyze permissions for a file, divide the 9 characters that comprise this attribute into three sets of 3. From left to right, these sets show permissions for:

rwx	the owner of the file
rwx	the owner's group
rwx	persons other than those above (others)

For files:

r	the file can be read
w	the file can be written to
x	the file can be executed
blank	the corresponding permission is not granted

For directories:

r	the directory can be read
w	entries in the directory can be created, moved, copied, or removed
x	the directory can be searched
blank	the corresponding permission is not granted

Note: for both files and directories, the following characters can appear only in the execution permission (x) position:

s	In the owner permissions section indicates that the set-userid bit is on. In the group permissions section indicates that the set group ID bit is on.
S	provides the same indication as s, but the execute bit is turned off.
t	indicates the sticky bit is on.
T	indicates the sticky bit is on but the execute bit is turned off.

Size An alias for SizeLo.

SizeHi For files greater than 4 gigabytes, the value that, when multiplied by 4GB and added to the value in the Size field, represents the size of the file. Valid value is a positive integer in the range 0 through 2147483647, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions.

SizeLo The value contained in the first 32 bits of the 64-bit value z/OS UNIX System Services uses to contain the file size. This value will represent the true file size for any but exceptionally large files. Valid value is a positive integer in the range 0 through 2147483647, and includes the use of the *AVG, *MAX, *MIN, or *SUM functions.

Type The type of file. A file can be one of the following types:

Dir	Directory File
Char	Character Special File
File	Regular File
Pipe	Named Pipe (FIFO) File
Link	Symbolic Link
Block	Reserved for Block
Socket	Socket File

UID Represents the numeric form of the userid that created or has been given ownership of the file. Valid value is a positive integer in the range 0 through 2147483647.

USS HFS ENQ Contention attributes

The USS HFS ENQ Contention attributes describe contention for a given hierarchical file system (HFS). The attributes are only available when enqueue contention exists. Since the HFS uses a scope of Systems to serialize HFS datasets, these attributes contain Scope Systems enqueue contention information only.

ASID ID of the address space issuing the enqueue. Valid format is a 4-byte integer.

Jobname Jobname of the address space issuing the enqueue. Valid format is an 8-character string.

ORIGINNODE The monitored z/OS operating system on which UNIX System Services is hosted. Valid value is a character string with a maximum length of 32 bytes.

Status Indicates whether the issuer of the enqueue is waiting for the resource or currently owns the resource. Valid values may be OWNER or WAITER.

System The system in which the issuer of the enqueue is executing. Valid format is an 8-character string

Time (Secs) Indicates how long, in seconds, OMEGAMON XE on z/OS has observed the issuer of the enqueue waiting for or owning the resource. The time displayed does not represent true ownership. It represents the time the issuer of the enqueue spent waiting for the resource to become available. If issuer has been waiting longer than the time the monitoring agent has been watching, the time represents the amount of time the monitoring agent has been watching. Valid format is a 4-byte integer.

Type Indicates whether the issuer of the enqueue requested exclusive (EXC) or shared (SHR) control. Valid values may be EXC or SHR.

USS Kernel attributes

The USS Kernel attributes provide information about the kernel address space.

CPU% The percent of CPU time used by the Kernel. Valid value is a 4-byte integer.

CPU Seconds The total CPU time of the application using services provided by UNIX System Services. (Note that this attribute is not just UNIX System Services CPU Time.) Valid values are positive integers in the range 0 through 2147483647.

I/Os Rate The number of I/Os per second. Valid value is a 4-byte integer.

I6Sock Current The number of Internet V6 sockets currently in use.

I6Sock Curr Pct Percentage of maximum sockets allowed that are Internet V6.

I6Sock HW Mark Most Internet V6 sockets in use at one time (high water mark).

I6Sock HW PCT Percentage of maximum sockets in use at one time (high water mark) that are Internet V6.

I6Sock Maximum Maximum Internet V6 sockets allowed.

ISock Current Number of Internet sockets currently in use.

ISock Curr Pct Percentage of maximum sockets in use which are Internet sockets.

ISock HW Mark Most Internet sockets in use at one time (high water mark).

ISock HW Pct Percentage of maximum sockets in use at one time (high water mark) that are Internet.

ISock Maximum Maximum Internet sockets allowed.

Kernel Address Space The name of the address space containing the z/OS support for the services of z/OS UNIX System Services. This address space can also be called the kernel. Valid value is an 8-byte alphanumeric character string and follows the z/OS rules for job name.

ORIGINNODE The monitored z/OS operating system on which UNIX System Services is hosted. Valid value is a character string with a maximum length of 32 bytes.

Max Processes The maximum [permissible number of processes in the system. Valid value is a 2-byte integer in the range 5 - 32767.

Message Queue ID Limit The maximum permissible number of message queue IDs. Valid value is a 4-byte integer.

Message Queue IDs Current number of message queue IDs. Valid value is a 4-byte integer.

Message Queue IDs% Ratio of used message queue IDs to maximum message queue IDs. Valid value is a 4-byte integer in the range 0 - 100.

mmap Storage Pages Current number of mmap storage pages. Valid value is a 4-byte integer.

mmap Storage Pages Limit Maximum number of mmap storage pages. Valid value is a 4-byte integer.

mmap Storage Pages% Ratio of used mmap storage pages to the maximum shared mmap storage pages. Valid value is a 4-byte integer in the range 0 - 100.

Number of Active Userids All active users, whether logged on or not. A started task or a batch job could be included in this count. The number of logged-on z/OS UNIX System Services users. Valid value is a positive integer in the range 0 through 32768.

Number of I/Os The total number of I/O operations performed by the kernel address space since it was started. When the value in this field reaches the field maximum, the value wraps to 0 again. Valid value is a positive integer in the range 0 through 2147483647.

Number of Processes The number of processes currently executing (active) in the system. Valid value is a positive integer in the range 0 through 32768.

Number of Syscalls The number of system calls made to the kernel address space since the address space was started. Some trivial system calls are not included. When the value in this field reaches the field maximum, the value wraps to 0 again. Valid value consists of a positive integer in the range 0 through 2147483647.

Release UNIX system's release number in standard format. Valid value is a 16-byte character string.

Semaphore ID Limit Maximum number of semaphore IDs. Valid value is a 4-byte integer.

Semaphore IDs Current number of semaphore IDs. Valid value is a 4-byte integer.

Semaphore ID% Ratio of used semaphore IDs to the maximum number of semaphore IDs. Valid value is a 4-byte integer in the range 0 - 100.

Shared Memory ID Limit Maximum number of shared memory IDs. Valid value is a 4-byte integer.

Shared Memory IDs Current number of shared memory IDs. Valid value is a 4-byte integer.

Shared Memory IDs% Ratio of used shared memory IDs to the maximum number of shared memory IDs. Valid value is a 4-byte integer in the range 0 - 100.

Shared Memory Pages The current total number of shared memory pages in use by all processes in the system. Valid value is a 4-byte integer.

Shared Memory Pages Limit Maximum number of shared memory pages for all segments. Valid value is a 4-byte integer.

Shared Memory Pages% Ratio of used shared memory pages for all segments to the maximum number of shared memory pages for all segments. Valid value is a 4-byte integer in the range 0 - 100.

Shared Storage Pages The current total number of shared pages in use by the system and all of the functions and services that share storage. Valid value is a 4-byte integer.

Shared Storage Pages Limit Maximum number of shared storage pages. Valid value is a 4-byte integer.

Shared Storage Pages% Ratio of used shared storage pages to the maximum number of shared storage pages. Valid value is a 4-byte integer in the range 0 - 100.

Syscall Rate The number of syscalls per second. Valid value is a 4-byte integer.

Sysname Standard format UNIX system name. Valid value is a 16-byte character string.

Total Shared CSTOR Pages Total number of groups in central storage including shared segment pages. Valid value is a 4-byte integer.

Total Shared ESTOR Pages Total number of shared page groups in expanded storage including shared segment pages. Valid value is a 4-byte integer.

Total Shared Fixed Pages Total number of shared page groups fixed in the system, including shared segments pages. Valid value is a 4-byte integer.

Total Shared Pages SYS Total number of shared page groups in the system including shared segments. Valid value is a 4-byte integer.

Used Processes Ratio of the current number of processes to the maximum number of processes that can be supported in the system. Valid value is a 4-byte integer in the range 0 - 100.

USock Current UNIX system's current number of sockets in use.

USock Curr Pct UNIX system's percentage of maximum sockets in use.

USock HW Mark UNIX system's most sockets in use (high water mark) at one time.

USock Maximum UNIX system's maximum of sockets allowed.

USock HW Pct UNIX system's percentage of maximum sockets in use at one time (high water mark).

Version UNIX system version in standard format. Valid value is an 8-byte character string.

zFS Address Space The name of the address space running zFS.

USS Logged on Users attributes

The USS Logged on Users attributes supply terminal identification and time data for each user logged into your monitored systems.

Idle Time (Mins) The elapsed time, in minutes, since the user last caused a system interrupt. If the idle time for a remote user is high, you may want to take action to prevent the misuse of system resources. Valid value is a numeric value in the range 0 through 2147483647 and can include the use of *AVG, *MIN, *MAX, or *SUM functions.

Login Name The name recognized by the system to permit a user to log in. This is a character form of the userid; generally a TSO userid. Valid value is comprised of alphanumeric characters and follows the z/OS rules for userid. An example is JDOE.

Login Time The day and time the user logged into the system. Valid value is a timestamp in the form mm/dd/yy hh:mm:ss.

Name The user’s real name specified in alphanumeric characters; for example, John Doe or Doe, John.

ORIGINNODE The monitored z/OS operating system on which UNIX System Services is hosted. Valid value is a character string with a maximum length of 32 bytes.

Process ID Name of the executing process associated with this logged-on user. Valid value is a 4-byte integer.

Remote Host The name of the system host from which a login (rlogin) connection was established. When the connection is local or established via telnet, the remote host is displayed as a blank. Only the first 16 characters of the host name are displayed.

Terminal The identification code of the terminal associated with the logged-in user. Valid value is comprised of alphanumeric characters and follows the z/OS rules for userid.

UID The numerical id of the user. Valid value is a positive integer in the range 0 through 2147483647.

USS Mounted File Systems attributes

The USS Mounted File Systems attributes provide basic status and activity information for the Mounted File Systems.

Bytes Read/S The number of bytes read per second. A value may not be available for recently mounted file systems. If a value is not available, Not available is displayed. Valid value is a 4-byte integer.

Bytes Written/S The number of bytes written per second. A value may not be available for recently mounted file systems. If a value is not available, Not available is displayed. Valid value is a 4-byte integer.

DDNAME The DDNAME of the hierarchical file system. This attribute is not applicable to Temporary File Systems, automount directories, or file systems owned by another system (such as in the case of shared HFSs). Valid value is an 8-character string.

Dir I/Os/S The number of directory I/Os per second. A value may not be available for recently mounted file systems. If a value is not available, Not available is displayed. Valid value is a 4-byte integer.

File System Name z/OS name of the HFS dataset that contains the file system. *AMD indicates an automount mount point. /TMP indicates a Temporary File System (TFS). Valid value is a string with maximum length of 44.

ORIGINNODE The monitored z/OS operating system on which UNIX System Services is hosted. Valid value is a character string with a maximum length of 32 bytes.

Members The number of members in the file system (number of files plus number of subdirectories). For Temporary File Systems (TFSs), Not available is displayed. Valid value is a 4-byte integer.

Mode Indicates whether the file is read or write. Valid values are READ or RDWR.

Mount Point Mount point or file system root. Sometimes referred to as the mount point path name or mount point directory. Valid value is a string with a maximum length of 1024.

Own Sys Identifies the system that ’owns’ the mount.

Percent Used Percentage of allocated space in the file system that has been used. For automount file systems, Not available is displayed. Valid value is a 4-byte integer.

Quiesce Job Name of the job that issued a quiesce for this file system. Valid value is an 8-byte character string.

Quiesce PID Process ID of the process that issued a quiesce for this file system. Valid value is a 4-byte integer.

Quiesce Sys Name of the system that issued a quiesce for this file system. Valid value is an 8-byte character string.

Read I/Os/S Number of data blocks read per second. One or more data blocks may be required to satisfy an application’s call to the z/OS UNIX System Services read API. The data block may have been read from DASD or from a cache. The value shown here is the upper bound with respect to actual device I/O caused by read requests to this file system. A value may not be available for recently mounted file systems. If a value is not available, Not available is displayed. Valid value is a 4-byte integer.

Reads/S Number of reads per second. In this context, a read is an application’s call to the z/OS UNIX System Services read API. A value may not be available for recently mounted file systems. If a value is not available, Not available is displayed. Valid value is a 4-byte integer.

Security Specifies the security available for this file system. Valid values are SECURITY and NOSECURITY.

SETUID Specifies whether SETUID and SETGID mode bits are honored when a program in the file system is executed. Valid values are SETUID and NOSETUID.

Status Status of the file system. Valid values are:

- Active
- Dead
- Resetting
- Unmounting (Drain)
- Unmounting (Force)
- Unmounting (Immediate)
- Unmounting
- Unmount (immediate) failed
- Quiesced
- Mounting
- Mounting (asynchronous)

Status Duration The length of time in seconds that the file system has been observed in its current status Valid value is a 4-byte integer or Not available.

SYSZDSN ENQ Wait Ct The number of contenders for the SYSZDSN resource. The HFS data sets are protected by enqueues using major names SYSZDSN and SYSDSN. Both enqueues are for SHARED control unless the HFS is mounted as Read/Write (RDWR). In that case, the SYSZDSN enqueue is for EXCLUSIVE control.

Normally, the OMVS address space should hold this resource and there should be no contention; that is, this column normally contains the value 0. If contention exists (that is, this column contains a value of 2 or greater), right-click on a value in the row to navigate to a report showing all contenders for the resource. Valid value is a 4-byte integer.

SYSDSN ENQ Wait Ct The number of contenders for the SYSDSN resource. The HFS data sets are protected by enqueues using major names SYSZDSN and SYSDSN. Both enqueues are for SHARED control unless the HFS is mounted as Read/Write (RDWR). In that case, the SYSDSN enqueue is for EXCLUSIVE control. Valid value is a 4-byte integer or Not available.

Normally, the OMVS address space should hold this resource and there should be no contention; that is, this column normally contains the value 0. If contention exists (that is, this column contains a value of 2 or greater), right-click on a value in the row to navigate to a report showing all contenders for the resource. Valid value is a 4-byte integer or Not available.

SYSZDSN Wait Seconds The longest time, in seconds, that any one waiter has been waiting for the SYSZDSN resource. The enqueue contention status of the HFS data sets is observed periodically. A waiter that is waiting for a resource for two consecutive observations is assumed to have been waiting for the resource for the entire observation interval. If ERR appears in this column, the wait time could not be determined. This value occurs rarely and could be the result of too many resource conflicts. If you are seeing this value occur regularly, you may want to write a situation triggered by this condition to determine the cause of the problem.

(See also SYSZDSN Wait Ct, above, for general information about the enqueues the OMVS address space uses for HFS data sets.) Valid value is a 4-byte integer or Not available.

SYSDSN Wait Seconds The longest time, in seconds, that any one waiter has been waiting for the SYSDSN resource. The enqueue contention status of the Hierarchical File System (HFS) data sets is observed periodically. A waiter that is waiting for a resource for two consecutive observations is assumed to have been waiting for the resource for the entire observation interval. If ERR appears in this column, the wait time could not be determined. This value occurs rarely and could be the result of too many resource conflicts. If you are seeing this value occur regularly, you may want to write a situation triggered by this condition to determine the cause of the problem.

(See also SYSDSN Wait Ct, above, for general information about the enqueues the OMVS address space uses for HFS data sets.) Valid value is a 4-byte integer or Not available.

Total Size The total size of the file system in units of 1K (1024) bytes. Not available may be displayed if an error was encountered when determining this value. Valid value is a 4-byte integer.

Total Used Total used space in units of 1K (1024) bytes. Not available may be displayed if an error was encountered when determining this value. Valid value is a 4-byte integer.

VOLSER The volume serial number of the volume or volumes on which the HFS resides. Up to 20 volsers can be displayed. Valid value is a string of length 139.

Write I/Os/S Number of data blocks written per second. One or more data blocks may be required to satisfy an application's call to the z/OS UNIX System Services write API. The data block may have been written to DASD or to a cache. The value shown here is the upper bound with respect to actual device I/O caused by write requests to this file system. A value may not be available for recently mounted file systems. If a value is not available, Not available is displayed. Valid value is a 4-byte integer.

Writes/S The number of writes per second. In this context, a write is an application's call to the z/OS UNIX System Services write API. A value may not be available for recently mounted file systems. If a value is not available, Not available is displayed. Valid value is a 4-byte integer.

USS Processes attributes

The USS Processes attributes provide detailed information about each process your system is currently executing.

ASID The numeric identifier of the address space containing this process. Valid value is a positive integer in the range 0 through 65535. (This value is generally not useful when defining situations.)

Command The individual command that initiated the process, together with any specified arguments. The column contains the command as entered; path information might not have been specified. The **Path** column always contains the path to the command, but does not contain any arguments, even if arguments were entered. This attribute can be any valid command.

z/OS UNIX System Services permits the total length of a command and its arguments to be up to 1024 bytes. If the command and its arguments are longer than the width of the column, the information is truncated and a plus sign (+) is displayed in the last position of the column.

If this is a traditional z/OS load module residing in a partitioned data set (PDS) or PDS/E (as opposed to an z/OS UNIX executable residing in an HFS file), this column will contain the 8-character load module name.

Command Name The command name alone, stripped of any parameters, to be used in situations.

Current VNode Tokens The current number of VNode tokens. A VNode token is an internal identifier used to refer to a file or a directory. Valid value is a 4-byte integer.

Effective Group ID The current group ID, but not necessarily the user's own ID. For example, a user logged in under a particular group ID may be able to change to another group ID. The ID to which the user changes becomes the effective group ID. Generally, this attribute contains a copy of the effective group identifier at the time a program is executed. The effective group ID, along with the effective userid, is used to determine file access permissions. Valid value is a positive integer in the range 0 through 2147483647.

Effective User A character representation of the effective userid. If the effective user ID is 0 (superuser), and if more than one user is assigned superuser (or root) authority, then the information in this column is unpredictable. A valid value consists of alphanumeric characters and follows the z/OS rules for user IDs.

Effective Userid The current user ID, but not necessarily the user's login ID. For example, if a user changes to another user's ID (assuming the proper authority), the ID to which the user changes becomes the effective user ID until the user switches back to the original login ID. The effective userid, together with the effective group ID, is used to determine file access permissions. Initially, this value is the same as the numeric userid (UID), but can be changed as the result of issuing an EXEC or SETUID. Valid value is a positive integer in the range 0 through 2147483647.

Execution State The execution state of the process. Valid values are:
- Unknown
- IPC_MSGRCV_WAIT
- IPC_MSGSND_WAIT
- Communication_kernel_wait
- IPC_SEMOP_WAIT
- Quiesce_freeze
- File_system_kernel_wait
- MVSPAUSE
- Process_terminated_still_group_leader
- Other_kernel_wait
- Quiesce_termination_wait
- Running_not_in_kernel_wait
- Sleep()_issued
- Waiting_for_child_process
- Fork()_a_new_process
- Process_terminated_parent_not

Foreground Process Group A process group whose member processes have certain privileges when accessing the controlling terminal. These privileges are denied to processes in background process groups. Valid value is a positive integer in the range 0 through 2147483647.

Jobname The z/OS name of the batch job, started task, or TSO address space containing this process.

Leader Session ID The process ID of the process that created the session to which this process belongs. (A session is a collection of one or more process groups.) Valid value is a positive integer in the range 0 through 2147483647.

Maximum VNode Tokens The maximum number of VNode tokens. A VNode token is an internal identifier used to refer to a file or a directory. Valid value is a 4-byte integer.

MVS Status The status of the z/OS subsystem process. Valid values are Normal, Swapped_Out, or Not_Used.

ORIGINNODE The monitored z/OS operating system on which UNIX System Services is hosted. Valid value is a character string with a maximum length of 32 bytes.

Parent Process ID The process ID of the parent process; that is, the process that created this process. When a process terminates, its children will be adopted by PID 1. This attribute is not useful in defining a situation. Valid value is a positive integer in the range 0 through 2147483647.

Path The full path to the command. This differs from the **Command** column in that the **Command** column contains the command as entered (path might not have been specified) and any specified options. The information in this column may be truncated if it contains too many characters. If truncation occurs, a plus sign (+) is displayed in the last column position.

Process Group A collection of processes that permits the signalling of related processes. Each process in the system is a member of a process group that is identified by a process group ID. Valid value is a positive integer in the range 0 through 2147483647.

Process ID Numerical identifier that z/OS UNIX System Services assigns to a process. Valid value is a positive integer in the range 0 through 2147483647.

Process Status The current state of the process. This information is generally meaningful only to persons having knowledge of z/OS UNIX System Services internals. Valid values are:

- One_Regular_Task_in_One_Process_in_Addr_Space
- Stopped_Process
- PTrace_Active_+_Multiple_Processes_In_Addr_Space
- PTrace_Active
- Multiple_Tasks_in_Process_+_Pthread_Task_in_Process
- Multiple_Tasks_in_Process_+_Multiple_Processes_in_Addr_Space
- More_Than_One_Open_Task_in_Process
- Pthread_Task_in_Process
- More_Than_One_Process_in_Addr_Space
- Not_Used

Real Group ID The group ID of the user who created the process and is defined in the password file. Valid value is a positive integer in the range 0 through 2147483647.

Real User A character representation of the real userid. If the real userid is 0 (superuser), and if more than one user is assigned superuser (or root) authority, then the information in this column is unpredictable. A valid value consists of alphanumeric characters and follows the z/OS rules for userid.

Real Userid The user ID specified in the /etc/password file. The attribute of a process that, at the time of process creation, identifies the user who created the process. Valid value is a positive integer in the range 0 through 2147483647.

Saved-Set Group ID The effective group ID at the time the setgid function was invoked (usually by exec) to change the effective group id; that is, the effective group ID before it was last changed. Valid value is a positive integer in the range 0 through 2147483647.

Saved-Set User A character representation of the saved-set userid. If the saved set userid is 0 (superuser), and if more than one user is assigned superuser (or root) authority, then the information in this column is unpredictable. A valid value consists of alphanumeric characters and follows the z/OS rules for userid.

Saved-Set Userid The effective userid at the time the setuid function was invoked (usually by exec) to change the effective userid; that is, the effective userid before it was last changed. Valid value is a positive integer in the range 0 through 2147483647.

Server Flags Server Flags.

Server Name The name of the server. Valid value is a string of length 32.

Server Type The type of server. Valid values are Not available, File, or Lock.

ServTypeH The Server Type in hexadecimal.

Start Time The time the process started executing. Valid value is a timestamp in the form mm/dd/yy hh:mm:ss.

Starting Time The time the process started executing. Valid value is a 4-byte integer.

System CPU Time The CPU time devoted to executing z/OS UNIX System Services system kernel code. System CPU time includes time spent executing system calls and performing administrative functions. It does not include some trivial calls. The value displayed is in seconds to the nearest hundredth of a second. For example, a displayed value of 1.23 would be 1 23/100 seconds. Valid values are numeric in the range 0 through 2147483647 and include the use of the *AVG, *MAX, *MIN, or *SUM functions.

Terminal Device The name of the terminal device that started the process. A valid value consists of alphanumeric characters and follows the z/OS rules for userid.

Total Size Represents the amount of virtual storage (in bytes) allocated to the process or to all processes running in a given address space. This value applies to extended (above the line) private storage only. It is calculated by subtracting the lowest allocated address in extended private storage from the highest allocated address.

UNIX Run Time CPU time for UNIX work. Valid value is a 4-byte integer.

UNIX Run Time% Percent of CPU execution for UNIX work. Valid value is a 4-byte integer in the range 0 - 100.

User CPU Time The CPU time that has been used by this process (excluding execution of kernel code). User CPU time includes time spent executing both user programs and library functions. It does not include CPU time spent executing system calls. The value displayed is in seconds to the nearest hundredth of a second. For example, a displayed value of 1.23 would be 1 23/100 seconds. Valid values are numeric in the range 0 through 2147483647 and include the use of the *AVG, *MAX, *MIN, or *SUM functions.

USS Threads attributes

The USS Threads attribute group shows information about all threads for a given process.

ASID ID of the address space containing this thread. Valid value is a 4-byte integer.

ORIGINNODE The monitored z/OS operating system on which UNIX System Services is hosted. Valid value is a character string with a maximum length of 32 bytes.

Process ID ID of the process this thread belongs to. Valid value is a 4-byte integer.

Process ID Hex Hexadecimal representation of the ID of the process this thread belongs to. Valid value is a 4-byte integer.

Thread ID The ID of the thread. Valid value is a character string of maximum length 20.

Thread Sequence Sequence portion of the thread ID. Valid value is a 4-byte integer.

UNIX Run Time CPU time for UNIX work. Valid value is a 4-byte integer.

UNIX Run Time% Percentage of CPU time for UNIX work. Valid value is a 4-byte integer in the range 0 - 100.

Weight Weight of the thread. Valid value may be NORMAL or LIGHT.

WLM Service Class Resources attributes

The WLM Service Class Resources Attribute Group monitors service class period performance relative to goals.

Actual Host Host actual value, either a response time in milliseconds or a percentage, based on the goal type. The value in this column is 0 if the goal type is Discret or Sysgoal . Valid value is 4-byte integer having 1 decimal place.

Average Response Time The average host response time of transactions in this service class. Valid value is an integer in the range of 0 through 2147483647, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions. If used as a threshold in a situation, this value should be expressed in milliseconds; for example 1200.

Average Storage The average storage for this service class, excluding private virtual storage allocated in this service class through cross-memory services on behalf of other service classes. Valid value is a 4-byte integer.

CSS Priority Average Channel Subsystem (CSS) DASD I/O priority observed for address spaces with I/O Activity for a given Service Class Period (SCP). For SCPs with no observed I/O activity, or if CSS I/O priority management is disabled in the CPC Reset Profile, a value of Not available will be reported. Valid value is a 4-byte integer in the range 0 -255.

Class Flag Categorizes the service class based on the general types of transactions managed by it. Valid values:

- Transaction — Service class manages subsystem transactions such as CICS or IMS transactions.
- Address Space — Service class handles workloads such as batch jobs and started tasks.
- TSO — Service class handles TSO users.

Duration The number of processor service units the period may use before work is passed to the next period. Valid value is an integer in the range of 0 through 2147483647 and can include the use of the *AVG, *MAX, *MIN, or *SUM functions.

Goal Importance The importance of this goal for the service class period. Goal importance for the service class is set as part of the Workload Manager (WLM) policy. When WLM cannot satisfy all goals, it tries to meet goals for more important service class periods first. Valid values:

- Not available
- Highest
- High
- Medium
- Low
- Lowest

Goal Percentile The percentage of work in the service class period (that is, the percentage of transactions) that completed within the expected response time (or goal value). Goal percentile is valid for percentile response time goals only. This value is set by the system administrator for this service class when the WLM policy is defined. It is not a measured result. Valid value is an integer in the range 0 to 100. This value can include the use of the *AVG, *MAX, *MIN, or *SUM functions.

Goal Type The type of goal for the service class period. Valid values:

- Not available
- SysGoal–Indicates that z/OS is to assign as high a priority as necessary to complete the work. It is reserved for use by the operating system to help prioritize system-related tasks.
- Discret (Discretionary)–Indicates that no specific service goal was defined for this service policy for the period. Work should run when resources are available. If multiple periods exist for a service class, this goal type applies only to the last period.
- Velocity–A percentage value indicating how fast work should execute when ready, without being delayed for processor or storage access. Velocity goals reflect the ratio of time when a unit of work was ready to use the CPU to the time work was actually using the CPU. This type of goal is useful for managing long-running non-transaction servers such as job entry subsystem (JES) or hierarchical storage management (HSM). A high percentage indicates that work should process quickly; a low percentage indicates that a greater amount of delay is acceptable.
- AvgResp–Specifies the average amount of time in which all work should complete.
- PctResp–%Resp mmm < nnnnuu, where mmm is the percentage of work to be completed, nnnn is the response time to be achieved, and uu is a unit of time (milliseconds, seconds, minutes, or hours).
- Velocio–A percentage value indicating how fast work should execute when ready, without being delayed for processor or storage access. Velocity goals reflect the ratio of time when a unit of work was ready to use the CPU to the time work was actually using the CPU or I/O. This type of goal is useful for managing long-running non-transaction servers such as job entry subsystem (JES) or or hierarchical storage management HSM. A high percentage indicates that work should process quickly; a low percentage indicates that a greater amount of delay is acceptable.

Goal Value A value representing the goal assigned to the service class period. Values corresponding to goal types appear on the workspace as follows:

- SysGoal—Actual goal values display as Not available. This type of goal indicates that z/OS is to assign as high a priority as necessary to complete the work, and is reserved for use by the operating system to help prioritize system-related tasks.
- Discret (Discretionary)—Actual goal values display as Not available. No service goal was defined for this service policy for the period.
- Velocit —Velocity > *nnn*%, where *nnn* is a numeric value indicating how fast work should complete.
- AvgResp—AvgResp < *nnnnuu*, where *nnn.n* is the average amount of time in which all work should complete and *uu* is a unit of time (milliseconds, seconds, minutes, or hours).
- PctResp—%Resp *mmm* < *nnnnuu*, where *mmm* is the percentage of work to be completed, *nnnn* is the response time to be achieved, and *uu* is a unit of time (milliseconds, seconds, minutes, or hours).
- Velocio—Velocio > *nnn*%, where *nnn* is a numeric value indicating how fast work should complete.

I/O Priority Average I/O Subsystem - UCB/Control Unit (IOS) priority observed for address spaces with I/O activity for a given Service Class Period (SCP). For Service Class Periods with no observed I/O activity, or if I/O priority management has been set to NO in the Workload Manager (WLM) policy, a value of Not available will be reported. Valid value is a 4-byte integer in the range 0-255.

I/O Rate The number of I/Os per second as a measurement of data set and JES spool I/O activity for all data sets associated with this service class period. Valid value is an number to one decimal place accuracy.

IFA Service Units zSeries Application Assist Processor (zAAP) service units as standard CP service units in floating point. This column is valid only on z/OS Version 1.6 and later. Valid value is a string consisting of a floating number followed by a decimal point and up to 22 numbers. You cannot use any numerical functions such as *AVG, *MAX, *MIN with this attribute.

IFA on CP Service Units zSeries Application Assist Processor (zAAP) work on regular processor service units as standard CP service units. This column is only valid on z/OS Version 1.6 and later. Valid value is a string consisting of a floating number followed by a decimal point and up to 22 numbers. You cannot use any numerical functions such as *AVG, *MAX, *MIN with this attribute.

Managed System z/OS operating system in your enterprise that a Tivoli OMEGAMON XE for z/OS agent is monitoring. Valid value is a character string with a maximum length of 32 bytes.

Percent CPU The percentage of CPU time consumed by this service class period. Valid value is a 4-byte integer in the range 0.00 - 100.00.

Percent IFA Percentage of Integrated Facility for Applications (IFA) resource used by the service class period. The value is normalized to reflect the faster speed of an IFA compared to a standard processor if LPAR is dispatched on a “knee-capped” processor. This column is only valid on z/OS Version 1.6 and later. Valid range is 0.00 - 100.00.

Percent IFA on CP Percentage of standard processor resource used by the service class period performing Integrated Facility for Applications (IFA) related work. This occurs when IFACrossOver=YES is specified in IEAOPTxx. This column is valid only on z/OS Version 1.6 and later. Valid range is 0.00 - 100.00.

Percent zIIP Percentage of System z9 Integrated Information Processors (zIIP) resource used by the service class period. The value is normalized to reflect the faster speed of a zIIP compared to a standard processor if the LPAR is dispatched on a ″knee-capped″ processor. This column is only valid on z/OS Version 1.8 and later or z/OS Version 1.6 or 1.7 with maintenance applied. Valid range is 0.00 - 100.00.

Percent zIIP On CP Percentage of standard processor resource used by the service class period performing System z9 Integrated Information Processors (zIIP) related work. This occurs when zIIP processors need help and zIIP-eligible work is place on the standard processor dispatch queue. This column is only valid on z/OS Version 1.8 and later or z/OS Version 1.6 and z/OS Version 1.7 with maintenance. Valid range is 0.00 - 100.00.

Performance Index Performance index for the host; a measure of how well a service class period is performing relative to its goal. Valid value is a 4-byte integer in the range 0.00 - 100.00. A performance index of less than 1 indicates that the period is performing better than its goal. A performance index of 1 indicates the period is meeting its goal. A performance index greater than 1 indicates the period is performing worse than its goal.

Average and percentile response time performance indexes are calculated by dividing the actual values by the goal values. The velocity performance index is calculated by dividing the goal value by the actual value.

Period The service class period number. If specifying a situation where the values should represent the service class as a whole, then specify Period=“SUM”. Period is used when building a situation for workflow thresholds at the service class level. Valid value is an integer in the range 1 through 8, and can include the use of the *AVG, *MAX, *MIN, or *SUM functions.

Promoted Percent The percentage of time the service class workload ran at a promoted dispatching priority.

Resource Group A group that specifies how much processor capacity is available to one or more service classes. Resource groups are used to limit or guarantee some amount of processing capacity to a service class or set of service classes. Valid value is a 4-byte integer.

Service Class The name of the service class being reported on. Valid value is a simple text string of from 1 through 8 characters; for example, TSO. A service class represents a uniquely-named group of work having common performance requirements such as similar performance goals, availability requirements, resource requirements, or business importance.

Service Class Description The text describing this service class as specified in the active service policy.

Workload The name of the workload that contains the service class. Valid value is a simple text string of from 1 through 8 characters; for example, BATCH.

Workload Description The text describing this workload as specified in the active service policy.

zIIP Service Units Number of service units consumed by the service class period performing work on a System z9 Integrated Information Processor. This column is valid on on z/OS version 1.6 and above.

zIIP on CP Service Units Number of service units consumed by the service class period performing work eligible for System z9 Integrated Information Processors on a standard processor. This column is valid on on z/OS version 1.6 and above.

About BPXPRM*xx* members

z/OS UNIX System Services is set up, in part, using members in the PARMLIB concatenation. These members are prefixed BPXPRM and have user-specified suffixes. One or more BPXPRM*xx* members are selected at initial program load (IPL). MVS operator commands allow the members in use to be changed at any time and allow many individual parameters to be overridden at any time.

The active values for parameters may not match those in the active members for two reasons: members being edited but not activated, and overrides by operator command.

Chapter 11. Workspaces

OMEGAMON XE on z/OS provides two sets of predefined workspaces: sysplex-level and system-level. These workspaces provide data in tabular and graphic form and enable you to effectively monitor the availability, resource consumption and performance of your sysplexes and their component z/OS systems.

The Sysplex Enterprise Overview workspace allows you to monitor the status and performance of your sysplexes and shared resources at a glance. Additional sysplex-level workspaces provide data on coupling facilities, global enqueues, GRS ring systems, report classes, service classes, resource groups, shared DASD groups, and cross-system coupling facilities (XCFs) used by each sysplex.

The system-level workspaces report data on address space CPU utilization, storage usage and bottlenecks, channel activity, common storage usage, single image DASD device usage, Workload Manager service class resource usage, and LPAR cluster activity for the z/OS images that participate in the sysplex. There are also workspaces that provide information about the status and configuration of any installed IBM cryptographic coprocessors and information about the use of UNIX System Services.

In addition to table and graph views, a workspace modified to contain other views, such as a notepad view, a browser session, an event console, a 3270 console, or a Take Action view that gives you the ability to send commands to the operator console.

The descriptions of each workspace in the help apply to the default settings (the components of the workspace in its original configuration). Any changes or updates that you make to a workspace are not reflected in the description of the workspace. For additional information about workspaces, as well as information about sorting and filtering table view data, refer to *IBM Tivoli OMEGAMON XE on z/OS: User's Guide* and the Tivoli Enterprise Portal online help.

Related topics:
  "Organization of OMEGAMON XE on z/OS workspaces"
  "Reporting historical data" on page 228
  "Customizing workspaces" on page 229
  "Accessing OMEGAMON XE on z/OS workspaces" on page 229
  "System workspaces" on page 247
  "Sysplex workspaces" on page 229

,

Organization of OMEGAMON XE on z/OS workspaces

The physical view of the Tivoli Enterprise Portal Navigator shows an enterprise as a mapping of platforms, systems, agents, and monitored resources. In a sysplex environment, monitored sysplexes appear between the platform and system levels of the Navigator tree, listed by their managed system names.

- 🌐 ENTERPRISE
  - ▣ z/OS Systems
    - ▣ SYSPLEX1:MVS:SYSPLEX
    - ▣ SYSPLEX2:MVS:SYSPLEX

Below each sysplex name are ▭ entries that represent a set of common objects shared by the sysplex that are monitored by Tivoli OMEGAMON XE on z/OS, and 🖥 entries that represent z/OS systems or LPARs being monitoring by Tivoli OMEGAMON XE on z/OS.

- ▣ SYSPLEX1:MVS:SYSPLEX
  - ▭ Coupling Facility Policy Data for Sysplex

- Coupling Facility Structures Data for Sysplex
- ...
- System1
- System2

Each entry is associated with one or more workspaces which provide information on those monitored objects. For example, the Shared DASD Groups Data for Sysplex entry is associated with the Shared DASD for Groups workspace, which provides information about sysplex-wide activity of shared DASD devices. From the Shared DASD for Groups workspace you can link to the Shared DASD Devices workspace, which displays information about the activity of the shared devices for a selected group, averaged over all systems in the sysplex.

Under each system name is an entry for each type of resource monitored by a Tivoli OMEGAMON XE monitoring agent on that LPAR. For example, if you have installed both Tivoli OMEGAMON XE on z/OS and Tivoli OMEGAMON XE for IMS, you will see two entries at this level:

- System1
  - MVS Operating System
    - SYSPLEX:SYSTEM1:MVSSYS
  - IMS

Under each agent entry is an entry for the managed system.

Sysplex managed system names take the form *plexname* :MVS:SYSPLEX, where *plexname* is normally the true name of the sysplex, but could be configured to be an alias name for the sysplex.

LPAR managed system names take the form *plexname* : *smfid* :MVSSYS, where *plexname* is normally the true name of the sysplex but could be configured to be an alias name for the sysplex. (This part of the LPAR managed system name typically matches the *plexname* component of its parent sysplex in the navigation tree.) The *smfid* component is the true SMF ID for the LPAR being monitored.

Organization of the sysplex workspaces

The following table shows the organization of the predefined workspaces, beginning with the “Sysplex Enterprise Overview workspace” on page 239. From the Sysplex Enterprise Overview workspace you can link to all of the sysplex-level workspaces listed in the Navigator.

Primary workspaces, listed in alphabetical order below, can be accessed directly from the Navigator. Other, secondary workspaces can be accessed only from another workspace. These workspaces are shown nested beneath the primary or secondary workspaces from which they can be linked.

- “Sysplex Level Overview workspace” on page 241
- “Coupling Facility Policy Data for Sysplex workspace” on page 230
- “Coupling Facility Structures Data for Sysplex workspace” on page 231
  - “MVS Systems Workspace for CF Structure” on page 234
  - “Users Workspace for CF Structure” on page 242
  - “Statistics for CF List or Lock Structures workspace” on page 238
  - “Statistics for CF Cache Structure workspace” on page 238
- “Coupling Facility Systems Data for Sysplex workspace” on page 231
  - “Paths Workspace for CF System” on page 234
  - “Statistics Workspace for CF System” on page 239

- “Global Enqueue Data for Sysplex workspace” on page 233
  - “Global Enqueue and Reserve workspace” on page 232
- “GRS Ring Systems Data for Sysplex workspace” on page 233
- “Report Classes Data for Sysplex workspace” on page 235
  - “Address Spaces Workspace for Report Classes” on page 229
- “Resource Groups Data for Sysplex workspace” on page 236
  - “Service Classes Workspace for Resource Group” on page 237
- “Service Classes Data for Sysplex workspace” on page 236
  - “Workflow Analysis Workspace for Service Class” on page 244
    - “Workflow Analysis Enqueue Workspace for Service Class Sysplex” on page 243
    - “Workflow Analysis I/O Workspace for Service Class” on page 243
  - “Periods Workspace for Service Class” on page 235
    - “Workflow Analysis Workspace for Service Class Period” on page 244
      - “Workflow Analysis Enqueue Workspace for Service Class Period” on page 242
      - “Workflow Analysis I/O Workspace for Service Class Period” on page 244
    - “Address Spaces Workspace for Service Class Period” on page 230
  - “Address Spaces Workspace for Service Class” on page 230
  - “Systems Workspace for Service Class” on page 242
    - “Workflow Analysis Workspace for Service Class System” on page 245
      - “Workflow Analysis Enqueue Workspace for Service Class System” on page 243
      - “Workflow Analysis I/O Workspace for Service Class System” on page 244
    - “Periods Workspace for Service Class System” on page 235
      - “Workflow Analysis Workspace for Service Class Period System” on page 245
        - “Workflow Analysis Enqueue Workspace for Service Class Period System” on page 243
        - “Workflow Analysis I/O Workspace for Service Class Period System” on page 244
  - “Subsystem Workflow Analysis for Service Class workspace” on page 239
- “Service Definition Data for Sysplex workspace” on page 237
- “Shared DASD Groups Data for Sysplex workspace” on page 237
  - “Shared DASD Devices workspace” on page 237
    - “Shared DASD Systems workspace” on page 238
- “XCF Systems Data for Sysplex workspace” on page 246
  - “XCF System Statistics workspace” on page 247
  - “XCF Paths Workspace from System Device To” on page 246
- “XCF Groups Data for Sysplex workspace” on page 245
  - “Members Workspace for XCF Group” on page 234
- “XCF Paths Data for Sysplex workspace” on page 246
  - “XCF Paths Workspace from System Device To” on page 246

Attribute groups used by sysplex-level predefined workspaces

In most cases, a workspace contains data or columns that have similar attributes in an attribute group. The table shows the relationships between the predefined workspaces and the attribute groups. (The workspaces are listed in alphabetical order.)

Workspace	Related Attribute Group
“Address Spaces Workspace for Report Classes” on page 229	“Service Class Address Spaces attributes” on page 113

“Address Spaces Workspace for Service Class” on page 230	“Service Class Address Spaces attributes” on page 113
“Address Spaces Workspace for Service Class Period” on page 230	“Service Class Address Spaces attributes” on page 113
“Coupling Facility Policy Data for Sysplex workspace” on page 230	“CF Policy attributes” on page 99
“Coupling Facility Structures Data for Sysplex workspace” on page 231	“CF Structures attributes” on page 99
“Coupling Facility Systems Data for Sysplex workspace” on page 231	“CF Systems attributes” on page 104
“Cross-system Cryptographic Coprocessor Overview workspace” on page 231	“Integrated Cryptographic Services Facilities Subsystems (ICSF) attributes” on page 168 “Service Call Type Detail (SVCDET) attributes” on page 186 “Top User Performance attributes” on page 194
“Enterprise Enqueue and Reserve workspace” on page 232	“Global Enqueues attributes” on page 110
“Global Enqueue Data for Sysplex workspace” on page 233	“Global Enqueues attributes” on page 110
“Global Enqueues attributes” on page 110	“Global Enqueues attributes” on page 110
“GRS Ring Systems Data for Sysplex workspace” on page 233	“GRS attributes” on page 111
“Members Workspace for XCF Group” on page 234	“XCF Members attributes” on page 128
“MVS Systems Workspace for CF Structure” on page 234	“CF Structure to MVS attributes” on page 103
“Paths Workspace for CF System” on page 234	“CF Path attributes” on page 98
“Periods Workspace for Service Class” on page 235	“Sysplex WLM Service Class Period attributes” on page 122
“Periods Workspace for Service Class System” on page 235	“Sysplex WLM Service Class Period attributes” on page 122
“Report Classes Data for Sysplex workspace” on page 235	“Report Classes attributes” on page 112
“Resource Groups Data for Sysplex workspace” on page 236	“Resource Groups attributes” on page 112
“Service Classes Data for Sysplex workspace” on page 236	“Sysplex WLM Service Class Period attributes” on page 122
“Service Classes Workspace for Resource Group” on page 237	“Resource Groups attributes” on page 112
“Service Definition Data for Sysplex workspace” on page 237	“Service Definition attributes” on page 118
“Shared DASD Devices workspace” on page 237	“Sysplex DASD Device attributes” on page 119
“Shared DASD Groups Data for Sysplex workspace” on page 237	“Sysplex DASD Group attributes” on page 121
“Shared DASD Systems workspace” on page 238	“Sysplex DASD attributes” on page 119
“Statistics for CF Cache Structure workspace” on page 238	“CF Structures attributes” on page 99
“Statistics for CF List or Lock Structures workspace” on page 238	“CF Structures attributes” on page 99
“Statistics Workspace for CF System” on page 239	“CF Systems attributes” on page 104

“Subsystem Workflow Analysis for Service Class workspace” on page 239	“Service Class Subsystem Workflow Analysis attributes” on page 116
“Sysplex Enterprise Overview workspace” on page 239	The overview workspace does not have a single associated attribute group. For information about the columns in the table, see: “Sysplex Coupling Facility Columns” on page 132 “Sysplexes Global Enqueue Columns” on page 133 “Sysplexes GRS Columns” on page 134 “Sysplexes Shared DASD Columns” on page 133 “Sysplexes Workloads Columns” on page 134 “Sysplexes XCF Columns” on page 135
“Sysplex Level Overview workspace” on page 241	The views in the Sysplex Level Overview workspace are associated with these attribute groups: “Enqueue Conflicts attributes” on page 163 “GRS attributes” on page 111 “Sysplex WLM Service Class Period attributes” on page 122 “CF Structures attributes” on page 99 “Sysplex DASD Group attributes” on page 121 “XCF Paths attributes” on page 129
“Systems Workspace for Service Class” on page 242	“Sysplex WLM Service Class Period attributes” on page 122
“Users Workspace for CF Structure” on page 242	“CF Clients attributes” on page 97
“Workflow Analysis Enqueue Workspace for Service Class Period” on page 242	“Service Class Enqueue Workflow Analysis attributes” on page 115
“Workflow Analysis Enqueue Workspace for Service Class Period System” on page 243	“Service Class Enqueue Workflow Analysis attributes” on page 115
“Workflow Analysis Enqueue Workspace for Service Class System” on page 243	“Service Class Enqueue Workflow Analysis attributes” on page 115
“Workflow Analysis Enqueue Workspace for Service Class Sysplex” on page 243	“Service Class Enqueue Workflow Analysis attributes” on page 115
“Workflow Analysis I/O Workspace for Service Class” on page 243	“Service Class I/O Workflow Analysis attributes” on page 115
“Workflow Analysis I/O Workspace for Service Class Period” on page 244	“Service Class I/O Workflow Analysis attributes” on page 115
“Workflow Analysis I/O Workspace for Service Class Period System” on page 244	“Service Class I/O Workflow Analysis attributes” on page 115
“Workflow Analysis I/O Workspace for Service Class System” on page 244	“Service Class I/O Workflow Analysis attributes” on page 115
“Workflow Analysis Workspace for Service Class” on page 244	“Sysplex WLM Service Class Period attributes” on page 122
“Workflow Analysis Workspace for Service Class Period” on page 244	“Sysplex WLM Service Class Period attributes” on page 122
“Workflow Analysis Workspace for Service Class Period System” on page 245	“Sysplex WLM Service Class Period attributes” on page 122
“Workflow Analysis Workspace for Service Class System” on page 245	“Sysplex WLM Service Class Period attributes” on page 122
“XCF Groups Data for Sysplex workspace” on page 245	“XCF Group attributes” on page 127
“XCF Paths Data for Sysplex workspace” on page 246	“XCF Paths attributes” on page 129

“XCF Paths Workspace from System Device To” on page 246	“XCF Paths attributes” on page 129
“XCF System Statistics workspace” on page 247	“XCF System Statistics attributes” on page 131
“XCF Systems Data for Sysplex workspace” on page 246	“XCF System attributes” on page 130

Organization of system-level predefined workspaces

OMEGAMON XE on z/OS provides a set of predefined workspaces that appear in the Navigator in the Business view. The hierarchical tree below is an alphabetical listing of these *primary* workspaces.

Some workspaces are accessible only as links from other workspaces. These are called either secondary or subsidiary workspaces and are shown below nested beneath their parent, primary workspace. For example, if you right-click on either a row or the link icon in a table view in the Address Space Overview workspace and choose Link To, you can link to the Address Space Bottlenecks and Impact Analysis workspace. Some workspaces can be accessed from multiple places.

- “Address Space Overview workspace” on page 255
  - “Address Space Bottlenecks and Impact Analysis workspace” on page 248
  - “Address Space Bottlenecks Detail workspace” on page 248
  - “Address Space Bottlenecks Summary workspace” on page 250
    - “Address Space Bottlenecks and Impact Analysis workspace” on page 248
    - “Address Space Bottlenecks Detail workspace” on page 248
    - “OMEGAMON for MVS - Job Details workspace” on page 267
  - “Address Space Common Storage - Active Users workspace” on page 250
    - “Address Space Common Storage - Trend Details workspace” on page 251
    - “Address Space Common Storage - Allocation Details workspace” on page 251
    - “Address Space Common Storage - Orphaned Elements workspace” on page 251
      - “OMEGAMON for MVS - CSA Analyzer workspace” on page 266
  - “Address Space CPU Usage Details workspace” on page 252
  - “Address Space CPU Usage Enclaves workspace” on page 252
  - “Address Space CPU Utilization workspace” on page 253
    - “Address Space CPU Usage Details workspace” on page 252
    - “Address Space CPU Usage Enclaves workspace” on page 252
    - “Inspect Address Space CPU Use workspace” on page 262
    - “Enclaves Owned by Selected Address Space workspace” on page 261
    - “OMEGAMON for MVS - Job Details workspace” on page 267
    - “WLM Service Class Information for Selected Address Space workspace” on page 276
  - “Address Space Details for Job workspace” on page 254
    - “Inspect Address Space CPU Use workspace” on page 262
    - “Address Space CPU Usage Enclaves workspace” on page 252
    - “Address Space Bottlenecks and Impact Analysis workspace” on page 248
    - “Dubbed Address Spaces workspace” on page 259
    - “Address Space Common Storage - Allocation Details workspace” on page 251
    - “Global Enqueue and Reserve workspace” on page 232
    - “WLM Service Class Information for Selected Address Space workspace” on page 276
  - “Address Space Storage workspace” on page 256
    - “OMEGAMON for MVS - Job Details workspace” on page 267
  - “Address Space Storage for Job workspace” on page 256
  - “Enclaves Owned by Selected Address Space workspace” on page 261
  - “OMEGAMON for MVS - Job Details workspace” on page 267
  - “WLM Service Class Information for Selected Address Space workspace” on page 276
    - “DASD MVS Devices workspace” on page 259
    - “Enqueue, Reserve, and Lock Summary workspace” on page 261
- “Channel Path Activity workspace” on page 257
- “Common Storage workspace” on page 258
- “Cryptographic Services workspace” on page 258

  - “Service Call Performance workspace” on page 269
  - “Top Users Performance workspace” on page 271
  - “Cross-system Cryptographic Coprocessor Overview workspace” on page 231
- “DASD MVS workspace” on page 258
- “DASD MVS Devices workspace” on page 259
- “Enclave Information workspace” on page 260
  - “Enclave Details workspace” on page 260
  - “Address Space Owning Selected Enclave workspace” on page 255
  - “WLM Service Class Resources workspace” on page 275
- “Enqueue, Reserve, and Lock Summary workspace” on page 261
  - “Enqueue and Reserve Detail workspace” on page 261
- “LPAR Clusters workspace” on page 265
  - “LPARs Assigned to a Cluster workspace” on page 266
  - “OMEGAMON for MVS – LPAR PR/SM Processor Statistics workspace” on page 267
- “Operator Alerts workspace” on page 268
- “Page Dataset Activity workspace” on page 268
- “Real Storage workspace” on page 268
- “System CPU Utilization workspace” on page 269
  - “HiperDispatch Details workspace” on page 264
- “System Paging Activity workspace” on page 270
- “Tape Drives workspace” on page 271
- “User Response Time workspace” on page 275
- “WLM Service Class Resources workspace” on page 275
  - “Address Space Bottlenecks in Service Class Period workspace” on page 249
  - “Enclaves in Selected Service Class and Period workspace” on page 261
  - “Address Space CPU Usage Class and Period workspace” on page 252
- “z/OS System Overview workspace” on page 277
  - “Address Space Common Storage - Active Users workspace” on page 250
  - “Address Space CPU Utilization workspace” on page 253
  - “Common Storage workspace” on page 258
  - “Enqueue, Reserve, and Lock Summary workspace” on page 261
  - “System CPU Utilization workspace” on page 269
- “z/OS UNIX System Services Overview workspace” on page 278
  - “Address Space CPU Usage Details workspace” on page 252 (filtered on the OMVS address space name)
  - “Address Space CPU Usage Details workspace” on page 252 (filtered on the zFS address space)
  - “Dubbed Address Spaces workspace” on page 259 (filtered for OMVS address space name)
    - “UNIX Processes workspace” on page 274 (for processes in the address space)
  - “Dubbed Address Spaces workspace” on page 259 (filtered for zFS address space name)
    - “UNIX Processes workspace” on page 274 (for processes in the address space)
  - “UNIX BPXPRMxx Values workspace” on page 271
  - “UNIX Files workspace” on page 272
    - “UNIX Files workspace” on page 272 (for a selected subdirectory)
    - “UNIX Mounted File Systems workspace” on page 273 (for the mounted file system containing the directory or file)
  - “UNIX Kernel workspace” on page 272
  - “UNIX Logged-on Users workspace” on page 273
    - “UNIX Processes workspace” on page 274 (for UID and PID)
  - “UNIX Mounted File Systems workspace” on page 273
    - “UNIX Files workspace” on page 272 (for the mount point)
    - “UNIX HFS ENQ Contention workspace” on page 272 (for SYSDSN and SYSZDSN Enqueues)
    - “UNIX Processes workspace” on page 274 (for processes using a selected mounted file system)
      - “UNIX Mounted File Systems workspace” on page 273 (for a selected process)
  - “UNIX Processes workspace” on page 274
    - “Dubbed Address Spaces workspace” on page 259 (for address spaces containing processes)
    - “UNIX Mounted File Systems workspace” on page 273 (for files systems in use by a process

        - “UNIX Threads workspace” on page 274 (for threads in a process)
        - “UNIX Processes workspace” on page 274 (for child processes)
    - “UNIX Threads workspace” on page 274
    - “zFS Overview workspace” on page 276
        - “zFS User Cache workspace” on page 277
- “z/OS System Overview workspace” on page 277
    - “Address Space CPU Utilization workspace” on page 253
    - “Address Space Common Storage - Active Users workspace” on page 250
    - “Common Storage workspace” on page 258
    - “Enqueue, Reserve, and Lock Summary workspace” on page 261
    - “System CPU Utilization workspace” on page 269

Attribute groups used by the system-level predefined workspaces

This table shows the relationships between the system-level predefined workspaces and attribute groups. (The workspaces are listed in alphabetical order.)

Workspace	Related Attribute Group
“Address Space Bottlenecks and Impact Analysis workspace” on page 248	“Address Space IA Impact attributes” on page 144
“Address Space Bottlenecks Detail workspace” on page 248	“Address Space Bottleneck attributes” on page 135
“Address Space Bottlenecks in Service Class Period workspace” on page 249	“Address Space Bottleneck attributes” on page 135
“Address Space Bottlenecks Summary workspace” on page 250	“Address Space Bottleneck attributes” on page 135
“Address Space Common Storage - Active Users workspace” on page 250	“Address Space ComStor Owned attributes” on page 140
“Address Space Common Storage - Allocation Details workspace” on page 251	“Address Space ComStor Owned Detail attributes” on page 142
“Address Space Common Storage - Orphaned Elements workspace” on page 251	“Address Space ComStor Unowned attributes” on page 143
“Address Space Common Storage - Trend Details workspace” on page 251	“Address Space ComStor Trends attributes” on page 143
“Address Space CPU Usage Class and Period workspace” on page 252	“Address Space CPU Utilization attributes” on page 146
“Address Space CPU Usage Details workspace” on page 252	“Address Space CPU Utilization attributes” on page 146
“Address Space CPU Usage Enclaves workspace” on page 252	“Address Space CPU Utilization attributes” on page 146
“Address Space CPU Utilization workspace” on page 253	“Address Space CPU Utilization attributes” on page 146
“Address Space Overview workspace” on page 255	“Address Space Summary attributes” on page 151
“Address Space Owning Selected Enclave workspace” on page 255	“Address Space CPU Utilization attributes” on page 146
“Address Space Storage workspace” on page 256	“Address Space Real Storage attributes” on page 149 “Address Space Virtual Storage attributes” on page 152
“Address Space Storage for Job workspace” on page 256	“Address Space Real Storage attributes” on page 149 “Address Space Virtual Storage attributes” on page 152
“Channel Path Activity workspace” on page 257	“Channel Paths attributes” on page 153
“Common Storage workspace” on page 258	“Common Storage attributes” on page 155

“Cross-system Cryptographic Coprocessor Overview workspace” on page 231	“Integrated Cryptographic Services Facilities Subsystems (ICSF) attributes” on page 168 “Service Call Type Detail (SVCDET) attributes” on page 186 “Top User Performance attributes” on page 194
“Cryptographic Services workspace” on page 258	“Integrated Cryptographic Services Facilities Subsystems (ICSF) attributes” on page 168
“DASD MVS workspace” on page 258	“DASD MVS attributes” on page 156
“DASD MVS Devices workspace” on page 259	“DASD MVS Devices attributes” on page 157
“Dubbed Address Spaces workspace” on page 259	“USS Address Spaces attributes” on page 199
“Enclave Information workspace” on page 260	“Enclave Table attributes” on page 161 “Address Space CPU Utilization attributes” on page 146
“Enclave Details workspace” on page 260	“Enclave Detail attributes” on page 158
“Enclaves in Selected Service Class and Period workspace” on page 261	“Enclave Table attributes” on page 161
“Enclaves Owned by Selected Address Space workspace” on page 261	“Enclave Table attributes” on page 161
“Enqueue and Reserve Detail workspace” on page 261	“Enqueue Conflicts attributes” on page 163
“Enqueue, Reserve, and Lock Summary workspace” on page 261	“Enqueue Conflicts attributes” on page 163 “KM5 Suspend Lock attributes” on page 173 “KM5 Spin Lock attributes” on page 172
“Inspect Address Space CPU Use workspace” on page 262	“Inspect Address Space CPU Use attributes” on page 166
“HiperDispatch Details workspace” on page 264	“HiperDispatch Logical Processors attributes” on page 164 “HiperDispatch Management attributes” on page 165
“LPAR Clusters workspace” on page 265	“LPAR Clusters attributes” on page 179
“LPARs Assigned to a Cluster workspace” on page 266	“LPAR Clusters attributes” on page 179
“OMEGAMON for MVS - CSA Analyzer workspace” on page 266	OMEGAMON for MVS
“OMEGAMON for MVS - Job Details workspace” on page 267	OMEGAMON for MVS
“OMEGAMON for MVS – License Manager MSU and WLM Capping workspace” on page 267	OMEGAMON for MVS
“OMEGAMON for MVS – LPAR PR/SM Processor Statistics workspace” on page 267	OMEGAMON for MVS
“Operator Alerts workspace” on page 268	“Operator Alerts attributes” on page 183
“Page Dataset Activity workspace” on page 268	“Page Dataset Activity attributes” on page 183
“Real Storage workspace” on page 268	“Real Storage attributes” on page 184
“Service Call Performance workspace” on page 269	“Service Call Type Detail (SVCDET) attributes” on page 186
“System CPU Utilization workspace” on page 269	“System CPU Utilization attributes” on page 190
“System Paging Activity workspace” on page 270	“System Paging Activity attributes” on page 193
“Tape Drives workspace” on page 271	“Tape Drives attributes” on page 193
“Top Users Performance workspace” on page 271	“Top User Performance attributes” on page 194
“UNIX BPXPRMxx Values workspace” on page 271	“USS BPXPRMxx Value attributes” on page 201
“UNIX Files workspace” on page 272	“USS Files attributes” on page 201

“UNIX HFS ENQ Contention workspace” on page 272	“USS HFS ENQ Contention attributes” on page 204
“UNIX Kernel workspace” on page 272	“USS Kernel attributes” on page 204
“UNIX Logged-on Users workspace” on page 273	“USS Logged on Users attributes” on page 207
“UNIX Mounted File Systems workspace” on page 273	“USS Mounted File Systems attributes” on page 208
“UNIX Processes workspace” on page 274	“USS Processes attributes” on page 210
“UNIX Threads workspace” on page 274	“USS Threads attributes” on page 213
“User Response Time workspace” on page 275	“User Response Time attributes” on page 199
“WLM Service Class Resources workspace” on page 275	“WLM Service Class Resources attributes” on page 214
“WLM Service Class Information for Selected Address Space workspace” on page 276	“WLM Service Class Resources attributes” on page 214
“Workloads for Classes workspace” on page 276	
“z/OS System Overview workspace” on page 277	“Address Space CPU Utilization attributes” on page 146 “Address Space ComStor Owned attributes” on page 140 “Global Enqueues attributes” on page 110 “System CPU Utilization attributes” on page 190 “Common Storage attributes” on page 155
“z/OS UNIX System Services Overview workspace” on page 278	“USS Kernel attributes” on page 204 “USS Logged on Users attributes” on page 207 “USS Mounted File Systems attributes” on page 208 “USS Address Spaces attributes” on page 199 “USS Processes attributes” on page 210
“zFS Overview workspace” on page 276	“KM5 zFS Directory Cache attributes” on page 174 “KM5 zFS Metadata Cache attributes” on page 175 “KM5 zFS Kernel attributes” on page 175 “KM5 zFS Storage attributes” on page 176
“zFS User Cache workspace” on page 277	“KM5 zFS User Cache attributes” on page 177 “KM5 zFS User Cache Dataspaces attributes” on page 178

Reporting historical data

All of the system-level workspaces, except the Address Space Bottlenecks workspace, and all of the sysplex-level workspaces, except the Users Data for CF Structure and GRS Ring Systems Data for Sysplex workspaces, can record and report historical data.

Workspace views that have history recording enabled have the history icon in the upper left corner.

For historical data to be available, historical data collection must be configured and enabled on all Tivoli Enterprise Monitoring Server in a sysplex. (Because the identity of sysplex proxy monitoring server can vary, the Tivoli Enterprise Portal cannot know which address space is currently serving as the proxy.)

For information on configuring Tivoli OMEGAMON XE on z/OS data storage, see the *IBM Tivoli OMEGAMON XE on z/OS: Configuration Guide*.

For information on configuring and reporting of historical data collection with the Tivoli Enterprise Portal, see the topics Historical Reporting Overview, Historical Reporting, and Configure History Data Collection in the online Help, and *OMEGAMON XE on z/OS: User's Guide*.

Customizing workspaces

With the appropriate authority, you can modify the predefined workspaces to better meet site-specific monitoring requirements, or create your own workspaces for specific monitoring purposes. (If you modify the predefined workspaces, you must save them with a new name.)

For help creating custom workspaces, see the topics Edit a Workspace, Create a Workspace, and Workspace Customization Options in theTivoli Enterprise Portal online help.

Accessing OMEGAMON XE on z/OS workspaces

To access Tivoli OMEGAMON XE on z/OS workspaces from the Navigator, expand the ⊞ 📁 z/OS Systems.

Under z/OS Systems you will see Navigator items for each of the sysplexes you are monitoring with IBM Tivoli OMEGAMON XE on z/OS agents.

When you expand the ⊞ plus icon beside the sysplex managed system name, you will see the Navigator items associated with the sysplex-level workspaces and an item for each system in the sysplex. When you expand a system entry, you will see a Navigator item associated with the OMEGAMON XE on z/OS agent (MVS Operating System), as well as an item for every Tivoli OMEGAMON XE agent that you have installed on that system (for example, you might see Storage Subsystem if you have a Tivoli OMEGAMON XE for Storage agent installed on the system).

When you select the managed system name of a sysplex or an LPAR (for example, LPAR320M:SYSG:MVSSYS), a list of available workspaces for the sysplex or LPAR is displayed below the name.

Sysplex workspaces

The topics in this section provide descriptions of the predefined sysplex-level workspaces and the navigational links between them. The predefined sysplex workspaces help you to determine if there is a problem within a sysplex and, if so, which component of the sysplex is in trouble.

To access the description of a workspace, select its name in the table of contents or select a link in the list of workspaces in “Organization of the sysplex workspaces” on page 220.

Related topics:

“Historical data collection for sysplex workspaces”

“Organization of the sysplex workspaces” on page 220

“Attribute groups used by sysplex-level predefined workspaces” on page 221

Historical data collection for sysplex workspaces

All the sysplex-level workspaces can record history, except the Users Data for CF Structure, GRS Ring Systems Data for Sysplex, and the Members Data for XCF Group workspaces.

For information on historical data collection and reporting, see the Historical Reporting Overview, Historical Reporting, Configure History Data Collection topics in the Tivoli Enterprise Portal help.

Address Spaces Workspace for Report Classes

This workspace provides performance and resource usage information for all workspaces in a selected report class. Use this information to help determine which address spaces may be having performance problems.

The Address Spaces for Report Class table view provides the system name, address space ID, and service class associated with each address space, as well as the number of working set frames in central and extended storage for the address space, velocity, page and I/O rate, and CPU percentage. The latter statistics are presented graphically in bar chart views.

You can link to this workspace from the Report Class Data for Sysplex workspace.

This workspace can record history.

Related topics: “Report Classes Data for Sysplex workspace” on page 235, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Address Spaces Workspace for Service Class

The Address Spaces workspace for Service Class displays information about the performance and resource utilization of all address spaces in the sysplex that are currently running in a given service class.

The Address Spaces for Service Class table view provides the system name, address space ID, and report class associated with each address space, as well as the number of working set frames in central and extended storage for the address space, velocity, page and I/O rate, and CPU percentage. The latter statistics are presented graphically in bar chart views.

You can link to this workspace from the Service Class Data for Sysplex workspace.

This workspace can record history.

Related topics: “Service Classes Data for Sysplex workspace” on page 236, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Address Spaces Workspace for Service Class Period

The Address Spaces workspace for Service Class Period displays information about the performance and resource utilization of all address spaces in the sysplex that are currently running in a selected service class period.

For each address space, the Address Spaces for Service Class table view shows the system name, address space ID, report class, velocity, page rate, the number of working set frames in central and expanded storage, number of I/Os per second experienced in the current interval, and the percentage of CPU used in the current interval.

Bar chart views graphically display the velocity and CPU usage percentages, page rates and I/O per second, and number of frames in central and expanded storage.

You can link to this workspace from the Period workspace for Service Class workspace. This workspace can record history.

Related topics: “Periods Workspace for Service Class” on page 235, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Coupling Facility Policy Data for Sysplex workspace

The Coupling Facility Policy Data for Sysplex workspace displays Coupling Facility Resource Management (CFRM) policy information. It lists the policy name and the date the policy was activated, and indicates whether a reformat of the CFRM is required.

Related topics: “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Coupling Facility Structures Data for Sysplex workspace

Coupling facilities provide communication between members of a sysplex to coordinate locking and buffer pool information. The Coupling Facility Structures Data for Sysplex workspace, accessed from the Navigator, presents information for each coupling facility structure throughout the sysplex.

The User Count bar chart graphically displays the total number of connected users and the number of users connected in an exception state for all structures with a problem user count greater than zero.

The Percentage of CF System Total Storage bar chart show the amount of the total allocated storage currently in use.

The Coupling Facility Structures table view provides the name, type, and status of each structure and statistics such as synchronous and asynchronous requests, allocated storage and percentage of storage in use, number of maximum possible and current users.

From the Coupling Facility Structures table you can link to the following workspaces, which provide detailed information on the selected structure:

“Users Workspace for CF Structure” on page 242
“MVS Systems Workspace for CF Structure” on page 234
“Statistics for CF List or Lock Structures workspace” on page 238
“Statistics for CF Cache Structure workspace” on page 238

This workspace can record history.

Related topics: “Service Classes Data for Sysplex workspace” on page 236, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Coupling Facility Systems Data for Sysplex workspace

The Coupling Facility Systems Data for Sysplex workspace displays status, performance, and resource usage information on each system for all coupling facilities defined to the sysplex.

The Resource Usage bar chart graphically displays the percentage of CPU and percentage of storage currently being utilized by each facility.

The Coupling Facilities Systems table provides such information as the number of connected z/OS systems, the number of structures in and out of policy, and the amount of allocated storage currently in use.

From the Coupling Facilities Systems table view you can link to workspaces that provide path information or detailed statistics for a selected coupling facility.

Related topics:“Paths Workspace for CF System” on page 234, “Statistics Workspace for CF System” on page 239, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Cross-system Cryptographic Coprocessor Overview workspace

The Cross-system Cryptographic Coprocessor Overview workspace shows summary information for Integrated Cryptographic Service Facility (ICSF) subsystems on each z/OS image, service call performance information by system, and top users by system.

The Cross-System ICSF report contains a row for each system which is being monitored. Each row indicates the number of active coprocessors (CMOS and PCI). Because coprocessors can be shared by LPARs, the total number of physical coprocessors cannot easily be determined from the the ICSF view. For example, if three systems are monitored and they are all LPARs of the same complex, they may show that they have two coprocessors each. But there may only be two physical coprocessors being shared by the three LPARs. (This is true of all LPAR hardware resources, like CECs and real storage.)

This workspace contains three table views:

- ICSF Subsystem by System table view

  From this view you can link to the Cryptographic Services workspace for the system you select.
- Service Call Performance by System table view

  From this view you can link to the Service Call Performance workspace for the system you select.
- Top Users by System table view

  From this view you can link to the Top Users Performance workspace for the system you select.

To access this workspace:

1. Select (click) the z/OS Systems Navigator item.
2. When the default workspace appears, right-click the Navigator item and select Workspace > Cross-system Cryptographic Coprocessor Overview from the pop-up menu.

Related topics: “Cryptographic Services workspace” on page 258, “Service Call Performance workspace” on page 269, “Top Users Performance workspace” on page 271, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Enterprise Enqueue and Reserve workspace

In configurations in which enqueue management spans two or more sysplexes, each z/OS image can be defined as belonging to an Enqplex (a group of z/OS images under common enqueue management). A resource in one Enqplex is distinct from a resource having the same name in another enqplex. Systems (that is, z/OS images) that do not define the enqplex to which they belong are assigned to the *DEFAULT Enqplex and are assumed to share resources and enqueue management. Two or more sysplexes having the same enqplex name share qname/rname resources.

The Enterprise Enqueue and Reserve workspace displays information about resource owners and waiters over the enterprise for a selected major.minor name. Since resources are unique to an enqplex, all owners and waiters will be from the same enqplex. The system name, sysplex name, and address space are provided to uniquely identify contenders using the same job name.

You can link to this workspace from the Enterprise Global Enqueue table of the Sysplex Enterprise Overview workspace.

This workspace cannot record history.

Related topics: “Sysplex Enterprise Overview workspace” on page 239, “Global Enqueue and Reserve workspace,” “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Global Enqueue and Reserve workspace

The Global Enqueue and Reserve workspace displays information on the global enqueues and reserves in the sysplex. It shows all of the owners and waiters for a particular resource within the sysplex. You can link to this workspace from the Global Enqueue Data for Sysplex workspace.

The Global Enqueue and Reserve table view provides:

- the name of the job containing the task that owns the resource
- the name of the address space currently waiting on the resource
- the hexadecimal ID of the address space
- the SYSNAME or SMF ID of the z/OS image where the task is executing

- the swap status of the address space in which the task is executing
- the type of enqueue the task has issued for the resource
- the amount of time (in seconds) that the task has been waiting for the resource

A bar graph graphically presents the task wait times.

You can link to this workspace from the Global Enqueue Data for Sysplex workspace.

Related topics: “Global Enqueue Data for Sysplex workspace,” “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Global Enqueue Data for Sysplex workspace

The Global Enqueue Data for Sysplex workspace displays information on all global enqueues for the sysplex. It provides data on enqueue conflicts in the sysplex, including its contribution to conflicts between sysplexes in configurations in which enqueue management spans two or more sysplexes.

In this workspace, two bar charts graphically illustrate the number of address spaces waiting for a resource and the amount of time each has been waiting.

You can link to this workspace from the Sysplexes Global Enqueues table of the Sysplex Enterprise Overview workspace.

From this workspace you can link to the Global Enqueue and Reserve workspace.

This workspace can record history.

Related topics: “Sysplex Enterprise Overview workspace” on page 239, “Global Enqueue and Reserve workspace” on page 232, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

GRS Ring Systems Data for Sysplex workspace

A global resource serialization (GRS) ring consists of one or more systems connected to each other by communication links. The links are used to pass information about requests for global resources from one system to another in the complex. Requests are made by passing a message or token, called the ring system authority (RSA) message, between systems in a round-robin or ring fashion.

The GRS Ring Systems Data for Sysplex workspace enables you to monitor the condition of the ring connections, the traffic on the rings, and the rings relevant to specific z/OS systems and their activities and response times. You can use the view to understand when GRS response time is elongating or if there are disruptions on the ring. This information is important for analyzing enqueue delays in workloads.

The Response Time bar chart graphically displays the actual measured ring response time, the anticipated or average time a requestor must wait for access to a global resource, and the minimum amount of time the RSA message is delayed in each z/OS image in the GRS complex.

The Global GRS Ring Systems table view displays information about the RSA messages issued by the z/OS image and the time required to allocate a resource.

Note: If the GRS complex is not in ring mode, this workspace will display data in only two fields: GRS Mode and Ring Acceleration. In such instances, use the GRS structure information in the coupling facility workspaces.

This workspace cannot record historical data.

Related topics: “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Members Workspace for XCF Group

This workspace provides information on the status and performance of members of a selected cross-system coupling facility (XCF) group.

The information provided by the Members of XCF Group table includes the job name of the task using this member’s name, the name of the z/OS system on which the member is executing, the member’s status and the status checking interval, and the number of signals sent and received by the member. The last two numbers are illustrated graphically in a bar chart.

You can link to this workspace from the XCF Groups Data for Sysplex workspace.

This workspace can record history.

Related topics: “XCF Groups Data for Sysplex workspace” on page 245, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

MVS Systems Workspace for CF Structure

The MVS Systems workspace for CF Structure displays information for each system or image that currently has users connected to the selected coupling facility structure.

The I/O Activity bar chart shows the average time (in milliseconds) synchronous and asynchronous requests take to complete on each image, as well as the average time requests are queued to the image. The User Count and Request Counts bar chart shows the number of users on each image connected to the structure, the number of requests to the structure from the system over the last interval, the number of requests per second from the system, and the number of synchronous requests over the last interval that were converted to asynchronous.

The MVS Systems view shows the same data in tabular form.

You can link to this workspace from the Coupling Facilities Structures Data for Sysplex workspace.

This workspace can record history.

Related topics: “Coupling Facility Structures Data for Sysplex workspace” on page 231, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Paths Workspace for CF System

The Paths workspace for CF System displays information about the XCF signaling used by the z/OS systems attached to the coupling facility. (A path can be an IBM 3088 MCCU data link, a parallel channel-to-channel (CTC) adapter, an Enterprise Systems Connection (ESCON) channel operating in CTC or basic mode or, if appropriate, a list structure in a coupling facility.)

Two bar charts display graphically the I/O rate for each channel and the percentage of requests that were delayed.

The Path Details for CF System Details view shows the same information in tabular form, as well as the status of each path.

You can link to this workspace from the Coupling Facility Systems Data for Sysplex workspace. This workspace can record history.

Note: When Resource Management Facility (RMF) is used as the data source for coupling facility data, this workspace has no data.

Related topics: “Coupling Facility Systems Data for Sysplex workspace” on page 231, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Periods Workspace for Service Class

The Periods Workspace for Service Class displays performance information for periods in a service class, such as transaction rate, average response time, and velocity. Use this view to determine whether a problem is occurring in a particular service class.

The Performance Index bar chart graphically displays how the service class and periods are performing in relation to goal. Other bar charts illustrate the average host response time of transactions in this service class, the velocity of the service class periods, and the numbers of transactions per second being processed by each service class period.

You can link to this workspace from the Service Classes Data for Sysplex workspace.

From this workspace you can link to the following types of workspaces:

- A workflow analysis workspace for a selected service class period
- An address spaces workspace for a selected service class period

Related topics: “Service Classes Data for Sysplex workspace” on page 236, “Workflow Analysis Workspace for Service Class Period” on page 244, “Address Spaces Workspace for Service Class Period” on page 230, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Periods Workspace for Service Class System

Periods Workspace for Service Class System displays performance information about the various periods in a selected service class for a selected system.

The Periods Details for Service Class System table provides information about transaction rate, average host response time, velocity, goal type and importance, the number of CPU service units the period may use before work is passed to the next period, and the performance index.

Four bar charts graphically display the performance index, average response time, velocity, and transaction rate per second.

You can link to this workspace from the Systems Workspace for Service Class. You can link from this workspace to Workflow Analysis Workspace for Service Class Period System.

Related topics: “Service Classes Data for Sysplex workspace” on page 236, “Systems Workspace for Service Class” on page 242, “Workflow Analysis Workspace for Service Class Period System” on page 245, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Report Classes Data for Sysplex workspace

Incoming work can be assigned a report class as well as a service class. Even though the workload manager uses the service class to manage the workloads, report classes can be used for additional reporting across workloads or service classes.

The Report Classes Data for Sysplex workspace provides an overview of all the report classes in a sysplex. The Report Classes table view presents data on number of address spaces in each report class, the number of page faults per second experienced by the address spaces, the percentage of processor used by the address spaces (velocity), and the number of I/O per second experienced by the report class during the current interval. This data is also presented graphically in three bar charts.

From this workspace you can link to an Address Spaces for Report Class workspace for more information about a particular address space in the report class.

Related topics: “Address Spaces Workspace for Report Classes” on page 229, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Resource Groups Data for Sysplex workspace

The Resource Groups Data for Sysplex workspace displays information about individual resource groups.

The Resource Group table provides information about minimum and maximum number of service units per second that can be consumed by the service classes in the resource group, the number of service units currently being consumed by the services in the group, and the percentage of the maximum consumable service units currently in use. You can also select a specific resource group in this table and link to service class information for that group.

A resource group can be used to limit the processing capability available to a service class or to guarantee some minimum processing capability to a service class. This information can be used when specifying a minimum and maximum amount of capacity, in service units, to a resource group. You can also assign multiple service classes to the same resource group.

Related topics: “Service Classes Workspace for Resource Group” on page 237, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Service Classes Data for Sysplex workspace

The Service Classes Data for Sysplex workspace displays performance information for the service classes defined for the sysplex.

Use the information in this workspace to determine which workloads are not meeting their goals. Open a service class to obtain more detailed performance information.

The Service Classes table view provides goal, performance, and workload information for each service class in the sysplex. Three bar charts graphically present performance indices, measured results for the service class period, and transactions per second.

From the Service Classes tables you can select a service class and link to one of the following types of workspaces for the service class:

- Workflow analysis
- Address spaces
- Periods
- Systems
- Subsytem workflow analysis (This link is active only when the selected service class has a class flag value of “Transaction”.)

Related topics: “Workflow Analysis Workspace for Service Class” on page 244, “Periods Workspace for Service Class” on page 235, “Address Spaces Workspace for Service Class” on page 230, “Systems Workspace for Service Class” on page 242, “Subsystem Workflow Analysis for Service Class workspace” on page 239, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Service Classes Workspace for Resource Group

This workspace provides capacity information for the service classes assigned to a selected resource group. You can link to this workspace from the Resource Groups Data for Sysplex.

The Service Classes for Resource Groups table shows the percentage of the maximum consumable service units currently being consumed by all services class in the group and the number of unweighted service units currently being consumed by the group. Two bar charts present this information graphically.

Related topics: “Resource Groups Data for Sysplex workspace” on page 236, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Service Definition Data for Sysplex workspace

The Service Definition Data for Sysplex workspace displays information about the service definition currently active for the sysplex, such as its name and when it was activated. Use this view to obtain the information about the installation needed for workload management processing.

Related topics: “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Shared DASD Devices workspace

The Shared DASD Devices workspace displays information about the activity of the shared devices for a group, averaged over all systems in the sysplex. Use the information to help determine how equitably a device is serving all systems in the sysplex.

The Shared DASD Devices table shows information for each device, such as control unit and device type, volume serial number, and statistics such as the device contention index or system with the worst response time or worst disconnect time.

Bar charts graphically display the true busy value and the device contention index for all devices.

You link to this workspace from the Shared DASD Groups Data for Sysplex workspace. From this workspace you can link to a workspace that displays information about the systems that share a selected device.

Note: No data is displayed in this workspace unless a DASD filter situation is enabled.

Related topics: “Shared DASD Groups Data for Sysplex workspace,” “Shared DASD Systems workspace” on page 238, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Shared DASD Groups Data for Sysplex workspace

The Shared DASD Groups Data for Sysplex workspace displays information to help identify the devices and datasets responsible for significant I/O time or delays for important workloads. Use the information to determine how equitably a device is serving all systems in the sysplex. Select a DASD group in the table view to display details for that particular group of DASD devices.

Note: The group name is the system managed storage (SMS) managed group name.

The Shared DASD Groups table provides statistics on device contention and percentage busy for all the groups in the sysplex. The Group True Busy bar chart graphically displays the average and highest true busy percentages for all groups, and the Group Contention Index chart shows the average and highest device contention indices for each group.

From the Shared DASD Groups table, you can link to a workspace displaying statistics for the individual devices in a selected group.

Note: No data is displayed in this workspace unless a DASD filter situation is enabled.

Related topics: “Shared DASD Devices workspace” on page 237, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Shared DASD Systems workspace

The Shared DASD Systems workspace displays information about the systems that share a device. Use the information to see the systems that share a particular device and to help measure the performance and exceptions from the perspective of each system.

The Shared DASD Systems table provides the system name and device address for each system; vary, cache, mount, and reconfiguration status for the devices; and statistics for the device as seen from the system such as average response, average pending, connect, and disconnect time for average I/O.

Average response time and average I/O rate for the device are presented graphically in bar charts.

You link to the Shared DASD Systems workspace form the Shared DASD Devices workspace.

Note: No data is displayed in this workspace unless a DASD filter situation is enabled.

Related topics: “Shared DASD Devices workspace” on page 237, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Statistics for CF Cache Structure workspace

The Statistics for CF Cache Structure workspace displays additional information about a cache structure selected in the Coupling Facility Structures table of the Coupling Facility Structures Data for Sysplex workspace.

This workspace offers three views:

- The Storage Usage bar chart graphically displays the amount of storage currently allocated to the structure and the minimum and maximum size of the structure.
- The Request Rates bar chart illustrates the number of synchronous and asynchronous operations per minute for this structure in the last monitoring interval.
- The Statistics Report for CF Cache Structure table view provides information about the configuration, size, and operation of the structure.

Related topics: “Coupling Facility Structures Data for Sysplex workspace” on page 231, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Statistics for CF List or Lock Structures workspace

The Statistics for CF List or Lock Structures workspace displays additional information about a list or lock structure selected in the Coupling Facility Structures table of the Coupling Facility Structures Data for Sysplex workspace.

This workspace offers three views:

- The Storage Usage bar chart graphically displays the amount of storage currently allocated to the structure and the minimum and maximum size of the structure.
- The Request Rates bar chart illustrates the number of synchronous and asynchronous operations per minute for this structure in the last monitoring interval.

- The Statistics Report for CF List or Lock Structure table view provides information about the configuration, size, storage usage, and operation of the structure.

Related topics: "Coupling Facility Structures Data for Sysplex workspace" on page 231, "Organization of the sysplex workspaces" on page 220, "Attribute groups used by sysplex-level predefined workspaces" on page 221

Statistics Workspace for CF System

The Statistics Workspace for CF System displays detailed statistics about a selected coupling facility. Use this workspace to determine what structures are receiving many requests, what applications are using the structures, and whether resources need to be adjusted.

You access this workspace from the Coupling Facilities Systems Data for Sysplex workspace.

Related topics: "Coupling Facility Systems Data for Sysplex workspace" on page 231, "Organization of the sysplex workspaces" on page 220, "Attribute groups used by sysplex-level predefined workspaces" on page 221

Subsystem Workflow Analysis for Service Class workspace

This workspace shows the execution states of the work processed by a selected service class. This view is displayed for transaction-type service classes only. Use the information in this workspace to see where the work in this service class is spending its time.

You access this workspace from the Service Classes Data for Sysplex workspace.

Related topics: "Service Classes Data for Sysplex workspace" on page 236, "Organization of the sysplex workspaces" on page 220, "Attribute groups used by sysplex-level predefined workspaces" on page 221

Sysplex Enterprise Overview workspace

This workspace provides an overview of all the sysplexes in your enterprise. You can use it to view at a glance information about the sysplexes in your network and select a specific sysplex to investigate in more detail. The Sysplex Enterprise Overview workspace is the default workspace for the z/OS Systems Navigator entry.

The Sysplex Enterprise Overview workspace presents seven views, one for every area of information provided by Tivoli OMEGAMON XE on z/OS, as well as information about enterprise-wide global enqueues.

Sysplexes Global Enqueue

The Sysplexes Global Enqueue table view provides information on all global enqueues conflicts in a sysplex. In terms of business goals, this view enables you to expedite the release of a particular resource.

From this table you can link to a workspace providing details for a selected workspace, such as a list of contenders for resources, as well as the amount of time a user has been waiting for a resource. See "Global Enqueue Data for Sysplex workspace" on page 233.

Sysplexes GRS

The Sysplexes GRS (Global Resource Serialization) table provides information about the condition of ring connections, the traffic on the rings, and the rings relevant to specific z/OS systems and their activities and response times. You can use the view to understand when GRS response time is elongating or if there are disruptions on the ring. The information is important for analyzing enqueue delays in workloads.

From this view, you can link to a workspace displaying GRS ring systems data for a selected sysplex. See “GRS Ring Systems Data for Sysplex workspace” on page 233.

Enterprise Global Enqueue

The Enterprise Global Enqueue table view displays detailed information on all global enqueues in the enterprise. It contains data on all enqueue conflicts in the enterprise, including conflicts between sysplexes in those configurations in which enqueue management spans two or more sysplexes. In such configurations, each z/OS image can be defined as belonging to an Enqplex.

An Enqplex is a group of z/OS images under common enqueue management. A resource in one Enqplex is distinct from a resource having the same name in another Enqplex. Systems (that is, z/OS images) that do not define the Enqplex to which they belong are assigned to the *DEFAULT Enqplex and are assumed to share resources and enqueue management. Two or more sysplexes having the same Enqplex name share qname/rname resources.

From this workspace, you can link to a workspace displaying global enqueue data for a selected sysplex. See “Enterprise Enqueue and Reserve workspace” on page 232.

Crypto Coprocessor Overview

The Crypto Coprocessor Overview table displays the name and status of Integrated Cryptographic Service Facility (ICSF) subsystems on all monitored z/OS images and the status of cryptographic services in those subsystems.

From this workspace, you can link to a workspace displaying summary information about Integrated Cryptographic Service Facility (ICSF) subsystems on each z/OS image, service call performance information by system, and top users by system. See “Cross-system Cryptographic Coprocessor Overview workspace” on page 231.

Sysplexes Workloads

The Sysplexes Workloads table displays information that enables you to determine whether your workloads are meeting their performance goals. Through detailed views, you are able to examine how a particular business workload is performing without regard to the number of address spaces or z/OS images involved in processing the work. You are also able to identify specific address spaces and systems hindering the performance of a workload.

From this table, you can link to workspaces that provide information about the services classes, service definitions, resource groups, and report classes associated with a selected sysplex. See “Service Classes Data for Sysplex workspace” on page 236, “Service Definition Data for Sysplex workspace” on page 237, “Resource Groups Data for Sysplex workspace” on page 236, “Report Classes Data for Sysplex workspace” on page 235.

Sysplexes Coupling Facility

The Sysplexes Coupling Facility view displays coupling facility information for all your monitored sysplexes. The coupling facility acts as a shared storage device, allowing address spaces on separate z/OS images to share data buffers and messages. This table provides an overview of coupling facility resources, which allows you to ensure that connections are working, that an acceptable number of structures are defined within the coupling facility, that CPU and storage are used properly, and that I/O rates are maximized.

From this table you can link to workspaces displaying data for the coupling facility policy, structures, and systems associated with a selected sysplex. See “Coupling Facility Systems Data for Sysplex workspace” on page 231, “Coupling Facility Structures Data for Sysplex workspace” on page 231, “Coupling Facility Policy Data for Sysplex workspace” on page 230.

Sysplex Shared DASD

The Sysplexes Shared DASD table enables you to view information regarding DASD shared by systems in a monitored sysplex.

From this view, you can link to a workspace which provides detailed information about the DASD groups shared by a selected sysplex. See “Shared DASD Groups Data for Sysplex workspace” on page 237.

Sysplexes XCF

The Sysplexes XCF view provides information on cross-system coupling facilities (XCF) use by all the sysplexes. XCF allows address spaces on different systems of a sysplex to communicate with one another in order to process work cooperatively. You can use the Sysplexes XCF view and related workspaces to perform the following tasks:

- Examine all XCF systems together and scrutinize individual systems
- Track defined XCF clusters
- Identify the status of each system and each signaling path in the sysplex
- Identify the groups, members of each group, and the status of the members
- Be alerted when a member or signaling path is not available

This report makes it easy to uncover problems that may require a path redefinition or adjustments to hardware (such as the failure of a link).

From this table, you can link to workspaces that display information about the XCF groups, systems, and paths associated with a selected sysplex. See “XCF Systems Data for Sysplex workspace” on page 246, “XCF Groups Data for Sysplex workspace” on page 245, and “XCF Paths Data for Sysplex workspace” on page 246.

Related topics: “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Sysplex Level Overview workspace

The Sysplex Level Overview workspace provides an overview of the sysplex you selected. The six table views in this workspace display information about global enqueues, GRS ring systems, services classes, shared DASD groups, coupling facilities structures, and cross-system coupling facilities (XCF) groups.

- From the Service Classes table, you can navigate to the following workspaces:
  “Report Classes Data for Sysplex workspace” on page 235
  “Resource Groups Data for Sysplex workspace” on page 236
  “Service Definition Data for Sysplex workspace” on page 237
  “Workflow Analysis Workspace for Service Class Period” on page 244
  “Periods Workspace for Service Class” on page 235
  “Address Spaces Workspace for Service Class” on page 230
  “Systems Workspace for Service Class” on page 242

  (If you are linking from a CICS service class, you can also link to a “Subsystem Workflow Analysis for Service Class workspace” on page 239.)
- From the Shared DASD Groups table, you can link to the “Shared DASD Devices workspace” on page 237 workspace.
- From the Coupling Facility Structures table, you can link to:
  “Statistics for CF Cache Structure workspace” on page 238 (if structure type selected is cache)
  “Statistics for CF List or Lock Structures workspace” on page 238 (if structure type selected is list or lock)
  “Coupling Facility Policy Data for Sysplex workspace” on page 230
  “Coupling Facility Systems Data for Sysplex workspace” on page 231

“Users Workspace for CF Structure”
“MVS Systems Workspace for CF Structure” on page 234

- From the XCF Groups table, you can link to the “Members Workspace for XCF Group” on page 234.
- From the Global Enqueue table, you can link to the “Global Enqueue and Reserve workspace” on page 232.

Related topics: “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Systems Workspace for Service Class

The Systems Workspace for Service Class displays performance information for systems on which the selected service class runs.

The Systems Details for Service Class table view provides goal, goal importance, actual host, performance index, transaction rate, actual network, actual total, and workload information for each system.

The performance index, actual host, and transactions per second for all systems are shown in graphic views.

You access this workspace from the Service Classes Data for Sysplex workspace.

From this workspace you can link to workspaces displaying workflow analysis or period information for a selected system.

Related topics: “Service Classes Data for Sysplex workspace” on page 236, “Workflow Analysis Workspace for Service Class System” on page 245, “Periods Workspace for Service Class System” on page 235, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Users Workspace for CF Structure

Users Workspace for CF Structure displays information for each user currently connected to the selected coupling facility structure. Use this workspace to determine which address spaces are connected to and using the structure and whether or not there is a problem with the connection. For example, there may be failed persistent connections that need to be cleaned up if the number of connections available is limited.

You access this workspace from the Coupling Facilities Structures Data for Sysplex.

Related topics: “Coupling Facility Structures Data for Sysplex workspace” on page 231, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Workflow Analysis Enqueue Workspace for Service Class Period

The Workflow Analysis Enqueue Workspace for Service Class Period workspace displays the major names of the enqueues which are affecting the selected service class period. For each major name, information is provided about the percent of wait time compared to all wait reasons. This information helps you locate where the service class is spending its time.

You link to this workspace from the Workflow Analysis Workspace for Service Class Period.

Related topics: “Workflow Analysis Workspace for Service Class Period” on page 244, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Workflow Analysis Enqueue Workspace for Service Class Period System

The Workflow Analysis Enqueue Workspace for Service Class Period System workspace displays enqueue details for a selected service class period system. The workspace shows the percentage of time a service class spent queued for a resource.

You access this workspace from the Workflow Analysis Workspace for Service Class Period System.

Related topics: “Workflow Analysis Workspace for Service Class Period System” on page 245, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Workflow Analysis Enqueue Workspace for Service Class System

The Workflow Analysis Enqueue Workspace for Service Class System displays the major names of enqueues which are affecting the selected service class. Information about each major name is the percentage of wait time compared to all wait reasons.

You access this workspace from the Workflow Analysis Workspace for Service Class System.

Related topics: “Workflow Analysis Workspace for Service Class System” on page 245, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Workflow Analysis Enqueue Workspace for Service Class Sysplex

This workspace provides information on what major name enqueue conflicts are outstanding over the entire sysplex and what percentage of the time a selected service class is spending waiting for that enqueue.

Queued Percent Distribution bar chart shows the percentage of time a service class spent queued for a major resource.

You access this workspace from the Workflow Analysis Workspace for Service Class.

This workspace can record history.

Related topics: “Workflow Analysis Workspace for Service Class” on page 244, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Workflow Analysis I/O Workspace for Service Class

Workflow Analysis I/O Workspace for Service Class displays devices which are active or queued for the selected service class. The Workflow Analysis I/O Details for Service Class table view provides the device number, type, and serial number, and the I/O active and I/O wait percentages for each device.

A bar chart graphically shows the distribution of I/O active and wait percentages by device.

You access this workspace from the Workflow Analysis Workspace for Service Class.

Related topics: “Workflow Analysis Workspace for Service Class” on page 244, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Workflow Analysis I/O Workspace for Service Class Period

Workflow Analysis I/O Workspace for Service Class Period displays devices which are active or queued for the selected service class period.

You access this workspace from the Workflow Analysis Workspace for Service Class Period.

Related topics: “Workflow Analysis Workspace for Service Class Period,” “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Workflow Analysis I/O Workspace for Service Class Period System

Workflow Analysis I/O Workspace for Service Class Period System displays devices which are active or queued for the selected service class period system.

You access this workspace from the Workflow Analysis Workspace for Service Class Period System.

Related topics: “Workflow Analysis Workspace for Service Class Period System” on page 245, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Workflow Analysis I/O Workspace for Service Class System

Workflow Analysis I/O Workspace for Service Class displays devices which are active or queued for the selected service class on a selected system. The Workflow Analysis I/O Details for Service Class table view provides the device number, type, and serial number, and the I/O active and I/O wait percentages for each device.

A bar chart graphically shows the distribution of I/O active and wait percentages by device.

You access this workspace from the Workflow Analysis Workspace for Service Class System.

Related topics: “Workflow Analysis Workspace for Service Class System” on page 245, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Workflow Analysis Workspace for Service Class

The Workflow Analysis Workspace for Service Class shows the execution states of the work processed by this service class. Use the information in this workspace to understand where the work in this service class is spending its time.

You access this workspace from the Service Classes Data for Sysplex workspace.

From this workspace you can link to

- Workflow Analysis Enqueue Workspace for Service Class Sysplex
- Workflow Analysis I/O Workspace for Service Class

Related topics: “Service Classes Data for Sysplex workspace” on page 236, “Workflow Analysis Enqueue Workspace for Service Class Sysplex” on page 243, “Workflow Analysis I/O Workspace for Service Class” on page 243, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Workflow Analysis Workspace for Service Class Period

The Workflow Analysis Workspace for Service Class shows the execution states of the work processed by this service class period. Use the information in this view to see where the work in this service class is spending its time.

You access this workspace from the Periods Workspace for Service Class.

From this workspace you can link to

- Workflow Analysis Enqueue Workspace for Service Class Period
- Workflow Analysis I/O Workspace for Service Class Period

Related topics: “Periods Workspace for Service Class” on page 235, “Workflow Analysis Enqueue Workspace for Service Class Period” on page 242, “Workflow Analysis I/O Workspace for Service Class Period” on page 244, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Workflow Analysis Workspace for Service Class Period System

The Workflow Analysis Workspace for Service Class Period System shows the execution states of the work processed by the service class period on the selected system. Use the information in this view to see where the work in this service class period is spending its time.

You access this workspace from the Periods Workspace for Service Class System.

From this workspace you can link to

- Workflow Analysis Enqueue Workspace for Service Class Period System
- Workflow Analysis I/O Workspace for Service Class Period System

Related topics: “Periods Workspace for Service Class System” on page 235, “Workflow Analysis Enqueue Workspace for Service Class Period System” on page 243, “Workflow Analysis I/O Workspace for Service Class Period System” on page 244, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

Workflow Analysis Workspace for Service Class System

The Workflow Analysis Workspace for Service Class System shows the execution states of the work processed by the service class on the selected system. Use the information in this view to see where the work in this service class is spending its time.

You access this workspace from the Systems Workspace for Service Class.

From this workspace you can link to

- Workflow Analysis Enqueue Workspace for Service Class System
- Workflow Analysis I/O Workspace for Service Class System

Related topics: “Systems Workspace for Service Class” on page 242, “Workflow Analysis Enqueue Workspace for Service Class System” on page 243, “Workflow Analysis I/O Workspace for Service Class System” on page 244, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

XCF Groups Data for Sysplex workspace

The XCF Groups Data for Sysplex workspace displays information about the cross-system coupling facility (XCF) groups defined in the sysplex. Use the information in this view to determine the number of members in a group that are in a problem state.

The XCF Groups table view provides the member and problem count for each group. A bar chart presents the same data graphically.

From this workspace you can link to the Members Workspace for XCF to see information about the members in a selected group.

This workspace can record history.

Related topics: "Members Workspace for XCF Group" on page 234, "Organization of the sysplex workspaces" on page 220, "Attribute groups used by sysplex-level predefined workspaces" on page 221

XCF Paths Data for Sysplex workspace

The XCF Paths Data for Sysplex workspace displays the status of each of the cross-system coupling facility (XCF) paths that connect systems to one another throughout the sysplex. You access this workspace from the Navigator.

The workspace contains two table views:

- The XCF Systems table shows the status of z/OS images in a sysplex as they relate to the XCF communication feature. It provides the name and status of the system and the release version of the operating system, as well as the length of time it takes for XCF to detect a failure in the sysplex and the time it takes for XCF to report a failure to the operator.
- The XCF Paths table provides the origin and destination systems of each signaling path, the device from which the path originates and its destination device, its status, the name of an associated transport class, and the percentage of its retry limit it has reached.

Performance information such as the Status and Retry Percent indicates whether any of the XCF paths are causing performance problems for the Sysplex. If a path becomes overloaded or unavailable, or the retry percentage is consistently high, you might need to adjust the parameters in the COUPLExx member of SYS1.PARMLIB.

From the XCF Paths table you can link to the XCF Paths Workspace from System Device To.

This workspace can collect history data.

Related topics: "XCF Paths Workspace from System Device To," "Organization of the sysplex workspaces" on page 220, "Attribute groups used by sysplex-level predefined workspaces" on page 221

XCF Paths Workspace from System Device To

XCF Paths Workspace from System Device To displays statistics for the paths from a selected system and device to a selected system and device. You can link to this workspace from the XCF Paths Data for Sysplex workspace.

A bar chart graphically displays signal activity for the path, including signals sent, signals pending transfer, signals received, and times receipt of message was delayed because no buffer was available.

The XCF Path table also includes the retry and message limits for the path, the number of sends issued when the path was busy, the amount of storage in use by the path, and the number of times the XCF restarted the path.

Related topics: "XCF Paths Data for Sysplex workspace," "Organization of the sysplex workspaces" on page 220, "Attribute groups used by sysplex-level predefined workspaces" on page 221

XCF Systems Data for Sysplex workspace

The XCF Systems Data for Sysplex shows path information for all cross-system coupling facilities (XCF) defined to the sysplex, as well as the status of each system in the sysplex.

The workspace contains two table views:

- The XCF Systems table shows the status of z/OS images in a sysplex as they relate to the XCF communication feature. It provides the name and status of the system and the release version of the operating system, as well as the length of time it takes for XCF to detect a failure in the sysplex and the time it takes for XCF to report a failure to the operator.

- The XCF Paths table provides the origin and destination systems of each signaling path, the device from which the path originates and its destination device, its status, the name of an associated transport class, and the percentage of its retry limit it has reached.

Performance information such as the Status and Retry Percent indicates whether any of the XCF paths are causing performance problems for the Sysplex. If a path becomes overloaded or unavailable, or the retry percentage is consistently high, you might need to adjust the parameters in the COUPLExx member of SYS1.PARMLIB.

From the XCF Systems table you can link to the XCF System Statistics workspace, which provides information on message traffic to or from the selected system.

This workspace can collect history data.

Related topics: “XCF System Statistics workspace,” “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

XCF System Statistics workspace

The XCF System Statistics workspace displays information about cross system coupling facilities (XCF) signals (messages) sent between z/OS systems in the sysplex. The workspace lets you identify where changes to transport class definitions can be made to optimize XCF performance and resource utilization. You can link to this workspace from the XCF Systems table of the XCF Systems Data for Sysplex workspace.

The XCF System Statistics table provides the system from and the system to which the signals are sent, the transport class, the number of signals sent and received, the times the path or transport buffer was unavailable, the buffer length, the number of messages that fit or were larger or smaller than the buffer, and the percentage of all messages which were degraded because the length was larger than the buffer.

By default, the **Times Path Unavailable**, **Times Buffer Unavailable**, and **Percent Degraded** columns have predefined thresholds that trigger a 🟨 warning indicator when the value is greater than 0.

This workspace also has a bar chart that graphically shows XCF outbound buffer utilization for the sysplex.

Related topics: “XCF Systems Data for Sysplex workspace” on page 246, “Organization of the sysplex workspaces” on page 220, “Attribute groups used by sysplex-level predefined workspaces” on page 221

System workspaces

The topics in this section provide descriptions of the system-level workspaces and the navigational links between them. To access the description of a specific workspace, select its name in the table of contents or the list of workspaces below.

“Historical data collection for system-level predefined workspaces”
“Organization of system-level predefined workspaces” on page 224
“Attribute groups used by the system-level predefined workspaces” on page 226

Historical data collection for system-level predefined workspaces

All of the system-level workspaces, with the exception of the Enqueue, Reserve, and Lock Summary workspace, the Address Space Bottlenecks workspaces, the Cross-system workspace, and the Integrated Cryptographic Service Facility (ICSF) workspace, can record history.

For information on historical data collection and reporting, see the topics on Historical Reporting Overview, Historical Reporting, Configure History Data Collection in the Tivoli Enterprise Portal online help and the Tivoli Enterprise Portal administrator's guide.

Address Space Bottlenecks and Impact Analysis workspace

Bottleneck analysis is a performance-monitoring technique that identifies execution states of a workload and the frequency of each state. When the results of this analysis are averaged over time, it is possible to find which states (such as waiting for CPU) prevent the workload from achieving its service goal. Identification and easing of bottlenecks is a key part of performance management.

This workspace displays resource contention information for the address space you selected from the Address Space CPU Utilization Summary table in the Address Space Overview workspace. It shows which workloads are using the resources that the impacted workload needs, which allows you to determine how various workloads are interfering with each other. This helps you to reduce degradation of the monitored workloads.

At the top of the workspace is the Execution States table, which provides a snapshot of any non-zero execution states that this particular job or address space is in for this period. This table contains three columns: Attribute, Percent, and Resource. Each attribute listed under the Attribute heading is one whose percentage value is greater than 0.0 in the Address Space Bottlenecks Detail workspace. The Percent column shows the percentage of time associated with an attribute for the specific address space in question. The Resource column only applies to enqueues and I/O waits. For an I/O Wait, it lists the device type, the volume serial number, and the device address. For an enqueue, it lists the enqueue major name.

The Active I/O Wait row contains a link to the "Shared DASD Devices workspace" on page 237. The Enqueue Wait attribute row contains a link to the "Enqueue, Reserve, and Lock Summary workspace" on page 261.

The workspace also contains an Impactors for *Job Name* bar chart, which shows which address spaces are impacting the favored address space, how they are impacting it, and to what extent they are impacting it.

The Address Space Bottlenecks Detail - Contention (%) by Resource table shows detailed information about the waits the address space is experiencing for various resources. By default, the **Tape Mount** column has a preset threshold (GT 0.0) that provides an ℹ informational alert when the job is being delayed for a tape mount.

The Impact Analysis for *Job Name* table view provides the same information as the Impactors for *Job Name* bar chart, along with the Job Name, address space ID, and the percentage of the total delay for which each impactor is responsible.

You can link to this workspace from the Address Space Overview workspace.

Related topics: "Address Space Overview workspace" on page 255, "Organization of system-level predefined workspaces" on page 224, "Attribute groups used by the system-level predefined workspaces" on page 226

Address Space Bottlenecks Detail workspace

This workspace provides wait state (bottleneck) information, showing contention percentages by resource for the address space you have chosen to investigate on the Address Space Bottlenecks Summary workspace or the Address Space CPU Utilization Summary table of the Address Space Overview workspace.

This workspace displays the Address Space Bottlenecks Detail - Contention (%) by Resource table view, which provides lower-level information about the address space you selected. By default, the Tape Mount column has a preset threshold (GT 0.0) that provides an informational alert when the job is being delayed for a tape mount.

The Execution States table at the top of the Address Space Bottlenecks Detail workspace is a portrait view of the single row of data presented in the table at the bottom of the workspace. This table contains three columns: Attribute, Percent, and Resource. Each attribute listed under the Attribute heading is one whose percentage value is greater than 0.0 in the Bottlenecks Detail workspace. The Percent column shows the percentage of time associated with an attribute for the specific address space in question. The Resource column only applies to Enqueue Wait and Active I/O. For Active I/O wait, it lists the device type, the volume serial number, and the device address. For Enqueue Wait, it contains the enqueue major name.

The Active I/O Wait row contains a link to the “DASD MVS Devices workspace” on page 259. The Enqueue Wait attribute row contains a link to the “Enqueue, Reserve, and Lock Summary workspace” on page 261.

Related topics: “Address Space Bottlenecks Summary workspace” on page 250, “Address Space Overview workspace” on page 255, “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

Address Space Bottlenecks in Service Class Period workspace

This workspace provides address space-level wait state (bottleneck) information, showing contention percentages by resource for a selected service class period. Various types of wait states are identified, including CPU wait percent, I/O wait percent, and enqueue wait percent. (For the full list of address space wait states, see the “Address Space Bottleneck attributes” on page 135.)

To use this workspace effectively, review the Address Space Bottlenecks Summary - Contention (%) by Resource table view to determine whether any address space is exceeding the wait threshold you have specified. This can easily be seen by locating the cells of this table that show either red or yellow alerts to indicate Critical or Warning conditions.

By default, the Tape Mount column has a preset threshold (GT 0.0) that provides an informational alert when the job is being delayed for a tape mount.

This workspace also contains the Selected Execution States bar chart. The bar chart provides data for address spaces whose jobs have been in one or more of five execution states for at least 5% of time.

The five execution states are:

- Using CPU
- CPU Wait
- Active I/O
- Queued I/O
- Enqueue Wait

You can link to this workspace from the WLM Service Class Resources workspace.

Note: This workspace cannot collect historical data.

Related topics: “Address Space Overview workspace” on page 255, “Address Space Bottlenecks Detail workspace” on page 248, “Address Space Bottlenecks and Impact Analysis workspace” on page 248, “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

Address Space Bottlenecks Summary workspace

This workspace provides address space-level wait state (bottleneck) information, showing contention percentages by resource. Various types of wait states are identified, including CPU wait percent, I/O wait percent, and enqueue wait percent. (For the full list of address space wait states, see the “Address Space Bottleneck attributes” on page 135.)

To use this workspace effectively, review the **Address Space Bottlenecks Summary - Contention (%) by Resource** table view to determine whether any address space is exceeding the wait threshold you have specified. This can easily be seen by locating the cells of this table that show either red or yellow alerts to indicate Critical or Warning conditions. For any such address spaces, click the link to drill down to the “Address Space Bottlenecks Detail workspace” on page 248 workspace and review the lower-level data presented there. From this view, you can also link to the “OMEGAMON for MVS - Job Details workspace” on page 267 for the job in the selected row.

By default, the Tape Mount column has a preset threshold (GT 0.0) that provides an informational alert when the job is being delayed for a tape mount.

This workspace also contains the **Selected Execution States** bar chart. The bar chart provides data for address spaces whose jobs have been in one or more of five execution states for at least 5% of time.

The five execution states are:
- Using CPU
- CPU Wait
- Active I/O
- Queued I/O
- Enqueue Wait

You can link to this workspace from the Address Space Counts table in the “Address Space Overview workspace” on page 255.

From this workspace, you can link to the “Address Space Bottlenecks and Impact Analysis workspace” on page 248.

This workspace cannot collect historical data.

Related topics:

“Organization of system-level predefined workspaces” on page 224

“Attribute groups used by the system-level predefined workspaces” on page 226

Address Space Common Storage - Active Users workspace

This workspace displays the name and ASID of current users of common storage (CSA, SQA, ECSA, or ESQA), as well as the amount and percentage of total common storage in use. This information can be used to identify address spaces that have allocated a higher than normal percentage of common storage. Details of common storage allocation or trends for an address space can be obtained via links.

Note: The CSA Analyzer must be running for data to be available in this workspace. For information on configuring the CSA Analyzer, see Configuring IBM Tivoli OMEGAMON XE for z/OS .

You can link to this workspace from the Address Space Counts view of the Address Space Overview workspace.

From this workspace, you can link to workspaces that provide allocation and trend details and data on orphaned elements.

This workspace can collect historical data.

Related topics: “Address Space Common Storage - Allocation Details workspace,” “Address Space Common Storage - Orphaned Elements workspace,” “Address Space Common Storage - Trend Details workspace,” “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

Address Space Common Storage - Allocation Details workspace

This workspace displays the attributes, including the size and beginning address, of each common storage area (CSA, SQA, ECSA, or ESQA) allocated by the selected address space. It also shows the return address (address of the next instruction after the GETMAIN request). This information can help identify problems in cases where there is an extremely large allocation or repeated allocations made by the same GETMAIN request.

Note: The CSA Analyzer must be running for data to be available in this workspace. For information on configuring the CSA Analyzer, see Configuring IBM Tivoli OMEGAMON XE for z/OS.

You can link to this workspace from the Address Space Common Storage - Active Users workspace.

This workspace cannot collect historical data.

Related topics: “Address Space Common Storage - Active Users workspace” on page 250, “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

Address Space Common Storage - Orphaned Elements workspace

This workspace displays storage attributes and age of orphaned (unowned) common storage areas (CSA, SQA, ECSA, or ESQA), as well as the name and ID of the address space that allocated the storage. This information can be used to help identify and isolate problems with a program’s common storage management. From each of the four table views you can link to the “OMEGAMON for MVS - CSA Analyzer workspace” on page 266 for the selected job in that area.

Note: The CSA Analyzer must be running for data to be available in this workspace. For information on configuring the CSA Analyzer, see *IBM Tivoli OMEGAMON XE on z/OS: Planning and Configuration Guide*.

You can link to this workspace from the “Common Storage workspace” on page 258.

This workspace cannot record history data.

Related topics:
- “Organization of system-level predefined workspaces” on page 224
- “Attribute groups used by the system-level predefined workspaces” on page 226

Address Space Common Storage - Trend Details workspace

This workspace displays the trend in common storage (CSA, SQA, ECSA, and ESQA) use by the selected address space, measured at five-minute intervals over the period specified in the parameters for the CSA Analyzer. This information can be used to readily identify common storage usage trends.

Note: The CSA Analyzer must be running for data to be available in this workspace. For information on configuring the CSA Analyzer, see Configuring IBM Tivoli OMEGAMON XE for z/OS.

You can link to this workspace from the Address Space Common Storage - Active Users workspace.

This workspace cannot record history data.

Related topics: “Address Space Common Storage - Active Users workspace” on page 250, “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

Address Space CPU Usage Class and Period workspace

This workspace provides information about CPU utilization by address space for a selected service class and period. The workspace contains a table view and bar chart that illustrate the CPU utilization for each address space in the selected class and period.

The name of the system, class, and period being reported on is found directly below the table view within the workspace. The table view displays basic identifying information such as job name and ASID, basic Workload Manager information such as service class and service class period, and various CPU statistics. It also provides enclave data, such as the total number of dependent and independent enclaves owned by the address space that are currently active or inactive.

The CPU Usage bar chart presents the CPU percentage, TCB percentage, and SRB percentage for each job listed in the workspace’s table view. (The bar chart does not reflect enclave CPU utilization. For enclave CPU utilization, see the “Enclave Information workspace” on page 260.)

You can link to the Address Space CPU Usage Class and Period workspace from a row of the WLM Service Class Resources table view in the WLM Service Class Resources workspace.

Related topics: “WLM Service Class Resources workspace” on page 275, “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

Address Space CPU Usage Details workspace

This workspace displays job level CPU percentages and related details for the selected address space in a tabular view.

In addition, the workspace contains are two graphical views. The Short Term CPU Usage bar chart plots this address space’s total CPU, TCB, and SRB percentage usage over the last few seconds (2.3 seconds by default). The Job CPU Usage bar chart shows the CPU accumulated since tracking started for this address space. It shows the chargeable CPU, which is TCB + SRB + Additional SRB Service.

You can link to this workspace from the Address Space CPU Utilization Summary table in the Address Space Overview workspace, or from the Address Space CPU Utilization table in the Address Space CPU Utilization workspace.

Related topics: “Address Space Overview workspace” on page 255, “Address Space CPU Utilization workspace” on page 253, “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

Address Space CPU Usage Enclaves workspace

This workspace provides detailed information about the enclaves in the selected address space.

This workspace contains three default views:

- The Address Space CPU Usage Enclaves table view provides the enclaves usage details, since monitoring started (at Tivoli Enterprise Monitoring Server startup or at the start of the job, whichever comes later).
- The Short Term CPU Usage bar chart plots this address space’s total CPU, TCB, and SRB percentage usage over the last few seconds (2.3 seconds, by default).

- The Job CPU Usage bar chart shows the CPU accumulated since monitoring started for this address space (at Tivoli Enterprise Monitoring Server start up or start of job, whichever comes later). It shows the chargeable CPU, which is TCB + SRB + Additional SRB Service.

You can link to this workspace from the Address Space CPU Utilization Summary table in the Address Space Overview workspace, or from the Address Space CPU Utilization table in the Address Space CPU Utilization workspace.

Related topics: “Address Space Overview workspace” on page 255, “Address Space CPU Utilization workspace,” “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

Address Space CPU Utilization workspace

This workspace provides information about CPU utilization by address space for all address spaces in this LPAR. The workspace contains a table view and bar chart that illustrate the CPU utilization for each address space on the selected system.

The name of the system being reported on is found directly below the **Address Space CPU Utilization** table view within the workspace. For each address space, the table view displays basic identifying information such as job name and address space ID, basic Workload Manager information such as service class and service class period, and various CPU statistics. It also provides enclave data, such as the total number of dependent and independent enclaves owned by the address space that are currently active or inactive.

The **CPU Usage** bar chart shows the CPU percentage, TCB percentage, and SRB percentage for each job listed in the workspace’s table view. (The bar chart does not reflect enclave CPU utilization. For enclave CPU utilization, see the “Enclave Information workspace” on page 260.)

From the **Address Space CPU Utilization** table, you can link to the

- “Enclave Information workspace” on page 260 to view just those enclaves owned by a particular address space
- “WLM Service Class Resources workspace” on page 275 to view service class period information for the service class and period in which the selected address space is executing
- “Address Space CPU Usage Details workspace” on page 252
- “Address Space CPU Usage Enclaves workspace” on page 252
- “Inspect Address Space CPU Use workspace” on page 262 to see data drilled down to the CSECT level of data for the most active TCBs, for a selected address space.

  The help for the Inspect Address Space CPU Use workspace contains information on the Inspect function, the default links, and customizing the Inspect parameters used to collect the data in the workspace.

  There are two links to the Inspect Address Space CPU Use workspace:

  - The Inspect Address Space CPU Use link uses the default parameters of 1000 samples at 5 millisecond intervals to collect the data presented in the workspace.
  - The Inspect with 5000 samples at 2ms interval link uses the specified parameters, but it can also be modified to use the sample count and sampling interval you specify to collect the data to be presented in the workspace.
- “OMEGAMON for MVS - Job Details workspace” on page 267 for information about the job in the selected row.
- The IBM Tivoli OMEGAMON XE for CICS on z/OS Region Overview workspace for the selected address space.

  Note: The OMEGAMON XE for CICS monitoring agent must be installed and running on the target system for the link to be functional.

You can link to the Address Space CPU Utilization workspace from the following locations:

- The Address Space Count view of the “Address Space Overview workspace” on page 255
- A row of the Enclave Information table view in the “Enclave Information workspace” on page 260 workspace to see information related to the address space owning the navigated-from enclave.
- A row of the WLM Service Class Resources table view in the “WLM Service Class Resources workspace” on page 275 to see information related to those address spaces executing in the service class and period of the navigated-from row of the **WLM Service Class** table view.

Related topics:

“Organization of system-level predefined workspaces” on page 224

“Attribute groups used by the system-level predefined workspaces” on page 226

Address Space Details for Job workspace

The Address Space Details for Job workspace provides a more comprehensive overview of the resource being used by a specific address space. The workspace comprises the following views:

- CPU Usage Details table view and CPU Usage Details bar graph. These views provide basic CPU usage, including zIIP and zAAP eligible work running on regular processors. From this view you can link to the following workspaces:
  - “Inspect Address Space CPU Use workspace” on page 262
  - “Address Space CPU Usage Enclaves workspace” on page 252
  - “Address Space Bottlenecks and Impact Analysis workspace” on page 248
  - “Dubbed Address Spaces workspace” on page 259
- Common Storage Usage table view. This table shows total common storage percentage used by the address space for each of the main common storage areas (CSA, ECSA, SQA, and ESQA). From this view you can link to the “Address Space Common Storage - Allocation Details workspace” on page 251 workspace.
- Enqueue Conflicts table view. This view shows enqueue conflicts that the job name may be involved with, either as an owner or as a waiter. In addition, it shows the wait time in seconds that the job has been waiting for each enqueue name. When the time is zero, the job is (probably) the owner of the enqueue. A time of greater than zero seconds indicates that the job has been waiting for the enqueue for that number of seconds. This may or may not indicate a problem, since waiting for certain enqueues may be part of the normal operation of the job. From this view you can link to the “Global Enqueue and Reserve workspace” on page 232.
- Address Space Details table view. This table displays basic address space information, including job start time, duration, and service class and period. From this view you can link to the “WLM Service Class Information for Selected Address Space workspace” on page 276.
- Address Space Real Storage table view. This table shows real storage usage by the address space.
- Address Space Virtual Storage table view. This table shows virtual storage usage details for the address space.
- Address Space Resource Details table view. This table shows resource usage details for the address space.

You can link to this workspace from the Address Space CPU Utilization Summary view of the “Address Space Overview workspace” on page 255.

Related topics:

“Organization of system-level predefined workspaces” on page 224

“Attribute groups used by the system-level predefined workspaces” on page 226

Address Space Overview workspace

The Address Space Overview workspace contains bar charts providing information on CPU usage, central storage frame counts, and fixed storage for a selected managed system. It also contains Address Space Count and Address Space CPU Utilization Summary table views. From these tables, you can navigate to other workspaces which provide more detail about a particular job, TSO user, started task, or APPC address space.

From the **Address Space Counts** table, you can link to the following workspaces. When linked to from this table, these workspaces present information for all relevant address spaces.

- “Address Space CPU Utilization workspace” on page 253
- “Address Space Bottlenecks Summary workspace” on page 250
- “Address Space Storage workspace” on page 256
- “Address Space Common Storage - Active Users workspace” on page 250

From the **Address Space CPU Utilization Summarization** table, you can link to the following workspaces. When linked from this table, these workspaces limit their presentation to the job, TSO user, started task, or APPC address space named in the parent row.

- “Address Space Bottlenecks and Impact Analysis workspace” on page 248
- “Address Space Bottlenecks Detail workspace” on page 248
- “Address Space CPU Usage Details workspace” on page 252
- “Address Space CPU Usage Enclaves workspace” on page 252
- “Address Space Details for Job workspace” on page 254
- “Address Space Storage for Job workspace” on page 256
- “Enclaves Owned by Selected Address Space workspace” on page 261
- “Periods Workspace for Service Class” on page 235
- “OMEGAMON for MVS - Job Details workspace” on page 267
- IBM Tivoli OMEGAMON XE for CICS on z/OS Region Overview workspace
- Tivoli OMEGAMON XE for Mainframe Networks Application Connections workspace

Note: For cross-product links to be available, the monitoring agent for the linked-to product must be installed and running on the target system.

Related topics:

- “Organization of system-level predefined workspaces” on page 224
- “Attribute groups used by the system-level predefined workspaces” on page 226

Address Space Owning Selected Enclave workspace

This workspace displays information about CPU usage by the address space which owns the selected enclave.

Bar charts display usage percentages for the current sampling interval (Short Term CPU) and cumulative percentages since collection began (Job CPU Usage) for, TCBs, SRBs, IFAs, and IFAs on CP.

The CPU Utilization for Address Space Owning Enclave: <Enclave_ID> provides detailed statistics on CPU usage for the address space.

You can access this workspace from the Enclave Information workspace.

This workspace can collect historical data.

Related topics: “Enclave Information workspace” on page 260, “Organization of system-level predefined workspaces” on page 224,“Attribute groups used by the system-level predefined workspaces” on page 226

Address Space Storage workspace

This workspace provides information about the storage allocated to an address space. It contains two table views and two bar charts.

The **Address Space Real Storage** table view provides information about the real storage allocated to an address space in terms of various types of frame counts and slot counts, as well as the management status of a given address space. From this view you can link to the “OMEGAMON for MVS - Job Details workspace” on page 267 for more information on the job in the selected row.

A frame is the basic unit of real and expanded storage. A slot is the basic unit of auxiliary storage. Each frame or slot consists of 4096 bytes of contiguous storage. An address space’s pages are located either on frames or in slots. Central Frames are those frames that are online or accessible to the processor or have been varied offline using a z/OS command or by the system itself. Fixed Frames are those pages of real storage that have been reserved in real storage by the address space and cannot be paged out or used by other address spaces. Expanded Frames are the number of frames of expanded storage owned by an address space. Hiperspace Frames are the number of frames of hiperspace storage owned by an address space. Hiperspace allows a program to store and retrieve data directly from expanded storage, avoiding the overhead of DASD I/O.

The **Address Space Virtual Storage** table view provides information about virtual storage use in three categories: Low, Extended, and Large. This new table provides tuning insight and opportunity for system and applications purposes. At a systems level, the large storage users can be identified, while at an applications level, analysis of how much storage is being used in each category as well as how much of the large storage object pace is in use can be performed.

The **Central Storage Frame Counts** bar chart graphs the counts of central frames and fixed frames. Pages are often fixed to prevent paging delays or when certain operations require that real addresses remain constant over time.

The **Fixed Storage** bar chart graphs fixed frame counts in terms of Low, Extended, and Large storage allocations, in megabytes. Low storage is that below 16 megabytes (MB), often referred to as “below the line.” From 16 MB to 2 gigabytes (GB), storage is referred to as Extended or “above the line,” and storage over 2 GB is referred to as Large or “above the bar.”

You can link to this workspace from the Address Space Counts table of the “Address Space Overview workspace” on page 255.

This workspace can record history.

Related topics:

“Address Space Storage for Job workspace”

“Organization of system-level predefined workspaces” on page 224

“Attribute groups used by the system-level predefined workspaces” on page 226

Address Space Storage for Job workspace

This workspace contains the same information as the Address Space Storage workspace for a selected job. It contains two table views and two bar charts:

The Address Space Real Storage table view provides information about the real storage allocated in terms of various types of frame counts and slot counts, as well as the management status of a given address space.

The Address Space Virtual Storage table view provides information about virtual storage use in three categories: Low, Extended, and Large. Low storage is that below 16 megabytes (MB), often referred to as “below the line.” From 16 MB to 2 gigabytes (GB), storage is referred to as Extended or “above the line,” and storage over 2 GB is referred to as Large or “above the bar.”

The Central Storage Frame Counts bar chart graphs the counts of central frames and fixed frames. Pages are often fixed to prevent paging delays or when certain operations require that real addresses remain constant over time.

The Fixed Storage Frame Counts bar chart graphs fixed frame counts in terms of Low, Extended, and Large storage allocations, in megabytes.

You can link to this workspace from the Address Space CPU Utilization table of the “Address Space Overview workspace” on page 255.

Related topics: “Address Space Overview workspace” on page 255, “Address Space Storage workspace” on page 256, “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

Channel Path Activity workspace

Note: The Resource Measurement Facility (RMF) must be started for this workspace to display data.

The Channel Path Activity workspace shows utilization for each channel path. If the utilization of a channel path is excessive, check

- DASD devices on the path, to determine whether there is contention for the path.
- Workloads, to determine whether they are doing excessive I/O to devices on the path
- Past utilization of the channel path

Excessive use of a channel path often indicates contention for poorly balanced resources. In the case of DASD, heavily used datasets should be evenly distributed among devices and channels. If a path is excessively busy only at a particular time, you may want to reschedule jobs that are affected by this contention.

In addition to a table view that provides information about various factors affecting utilization, two bar graphs provide easy-to-read information.

The LPAR Percent chart graphs the LPAR Percent column of the table view. LPAR Percent is the percentage of time the channel was busy working for the logical partition.This column can be blank or non-blank. The column is blank when: a) the system is running in basic mode; b) the channel path is offline; c) the ESCON Multiple Image Facility (EMIF) is not installed; or d) Channel Path Measurement Facility (CPMF) is unavailable or not installed.

Valid values range from 0 to 100% when the system is running in LPAR mode, EMIF is installed, and the channel path is online. If the channel path is shared, this value represents the LPAR’s utilization of the channel path from the start of the RMF interval. If the channel path is not shared, the LPAR utilization value is the same as the complex-wide utilization value.

This value can be greater than the utilization of the channel by the entire complex, since the LPAR value is obtained from the hardware, whereas the Complex value is sampled and may not include channel utilization.

The Complex Percent chart graphs the values in the Complex Percent column of the table view. Complex percent is the percentage of time that a given channel was busy within the complex. Each bar in the chart represents a given channel path.

This workspace can record history.

Related topics: “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

Common Storage workspace

Note: The Common Storage Area Analyzer (CSA Analyzer), supplied with Tivoli OMEGAMON XE products, must be started for this workspace to display data.

This workspace, accessed from the Navigator, displays information about four important areas of common storage: common service area (CSA), extended CSA (ECSA), system queue area (SQA), and extended SQA (ESQA). The information provided includes

- the percentage and amount of CSA and ECSA currently allocated. The percentage is calculated by dividing the amount of storage currently allocated by the total amount of storage available and includes storage in use and unowned storage.
- the amount of SQA and ESQA that is overflowing into CSA and ECSA, respectively. The overflow values are not included in SQA and ESQA total sizes.
- the amount of storage currently in use
- the total size of each area. If the bottom of CSA or the top of ECSA is not on a megabyte boundary, further rounding occurs to move either the bottom of CSA or the top of ECSA to the next megabyte boundary.
- the amount of allocated storage that is not owned by an address space
- the growth in use during the last interval.

The in-use percentage and amount of growth are also illustrated as bar charts. CSA and SQA use is typically stable, growing only when a subsystem, for example, DB2, is started, at IPL time, or when many new users log on. Therefore, you can use growth in common storage to spot potential problems as well as trends in CSA use.

This workspace can record history.

Related topics: “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

Cryptographic Services workspace

The Cryptographic Services workspace provides an overview of monitored cryptographic coprocessors.

The upper half of the workspace contains status information on situations monitoring cryptographic coprocessors and situation event history data. The lower half contains ICSF Subsystem Status table, which displays current configuration and status values for the subsystem.

Note: Tivoli OMEGAMON XE on z/OS cannot monitor cryptographic service calls unless at least one coprocessor is installed and configured.

Related topics: “Cross-system Cryptographic Coprocessor Overview workspace” on page 231, “Service Call Performance workspace” on page 269, “Top Users Performance workspace” on page 271, “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

DASD MVS workspace

Note: The Resource Measurement Facility (RMF) must be started in order for this workspace to display data.

This workspace assists you in monitoring various error conditions for a group of devices. It identifies problem conditions and lets you know how many devices are in the stated condition. The columns in the table view are represented graphically in the bar chart in the upper right corner of the workspace.

This workspace can record history.

Related topics: "Organization of system-level predefined workspaces" on page 224, "Attribute groups used by the system-level predefined workspaces" on page 226

DASD MVS Devices workspace

Note: The Resource Measurement Facility (RMF) must be started for this workspace to display data.

The DASD MVS Devices workspace permits you to monitor activity for individual DASD devices. Individual bar charts graphically display data found in the various columns of the DASD MVS Devices table view. Use this workspace to determine those devices in your enterprise where problems are occurring.

This workspace can record history.

Related topics: "Organization of system-level predefined workspaces" on page 224, "Attribute groups used by the system-level predefined workspaces" on page 226

Dubbed Address Spaces workspace

The Dubbed Address Spaces workspace contains information about all z/OS address spaces that have been marked as being a user of UNIX System Services requests. Such address spaces are referred to as "dubbed".

Dubbing is the process of making an address space known to the z/OS UNIX System Services kernel. Address spaces created by fork() are automatically dubbed when they are created. Other address spaces become dubbed if they invoke a z/OS UNIX service, either directly or indirectly. *Indirectly* means that the address space does not itself invoke the OS/390® UNIX System Services APIs, but uses an application that does. For example, address space MYJOB does not directly use the z/OS UNIX System Services APIs. Howver, it does use TCP/IP running on the current version of z/OS. Since this version of TCP/IP uses the z/OS UNIX System Services APIs, address space MYJOB is dubbed and is included in the Address Spaces workspace.

This workspace contains two views. The CPU Time% bar chart shows the percentage of CPU time used by each dubbed address space. The UNIX Address Spaces table view provides futher information about those address spaces, including the name of the address space; the type of address space, batch, started task, or TSO user; the TSO userid associated with the address space; and the service class (if the system is in goal mode) or performance group (if the system is in compatibility mode) of the address space. For a batch job or started task, the name of the address space is the jobname from the job statement. For a TSO address space, this is the user ID of the logged-on user. The address space name may also have been generated by z/OS UNIX System Services.

You can link to this workspace from the "z/OS UNIX System Services Overview workspace" on page 278 and the "UNIX Processes workspace" on page 274.

You can link from this workspace to the "UNIX Processes workspace" on page 274 to see information about the processes running in a selected address space.

This workspace can report historical data.

Related topics: "Organization of system-level predefined workspaces" on page 224, "Attribute groups used by the system-level predefined workspaces" on page 226

Enclave Information workspace

An enclave is a transaction that spans more than one dispatchable unit of work and more than one address space, and even multiple systems, but is reported on and managed as a single unit of work.

This workspace contains two views. The Enclave Information table view provides information about the owner, characteristics, and CPU utilization of each enclave. The Enclave CPU % by Address Space bar chart summarizes enclave CPU utilization at the address space level. In the chart, there are two bars for each address space: one for dependent enclaves and one for independent enclaves.

Note: The Enclave Information table view and the Enclave CPU % by Address Space bar chart derive information from different attribute groups. The Enclave Information table view provides data derived from the Enclave Table attributes, while the bar chart provides data derived from the Address Space CPU Utilization attributes. Since the data is not gathered in one atomic operation, there may be discrepancies in the information provided by the table view and bar chart. For example, if an enclave has just completed, it is possible that its owner will be represented in the bar chart even though the enclave does not appear in the Enclave Information table view. When you display this workspace as the result of navigating from another workspace, the Enclave Information table view generally displays only a subset of all of the enclaves in the system. However, the bar chart always reflects all address spaces owning enclaves.

This workspace can be reached from the Navigation view, from the Address Space CPU workspaces, or from the WLM Service Class Resources workspace. When accessed from the Navigator, the Enclave Information workspace provides performance data for all existing enclaves.When accessed by the link from the Address Space CPU Utilization workspace, only the enclaves associated with the address space selected in the link are presented. When accessed from the service class view, only the enclaves in that service class and period are shown.

From the Enclave Information workspace, you can navigate to

- the “Enclave Details workspace,” which shows performance and extended classification information for a selected enclave
- the “Address Space Owning Selected Enclave workspace” on page 255, which shows the CPU utilization for the address space owning the selected enclave
- the “WLM Service Class Resources workspace” on page 275, which provides view service class period information for the service class and period in which the enclave is executing

This workspace can record history.

Related topics: “Address Space CPU Utilization workspace” on page 253, “Address Space Owning Selected Enclave workspace” on page 255, “Enclave Details workspace,” “WLM Service Class Resources workspace” on page 275, “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

Enclave Details workspace

This workspace provides detailed performance and extended classification information for a selected enclave.

This workspace contains two views. The Enclave CPU% bar chart shows all enclave-related CPU usage for the address space that owns the selected enclave. The Enclave Details table view displays the performance and Workload Manager classification data for the selected enclave.

You can link to this workspace from the Enclave Table view of the Enclave Information workspace.

This workspace does *not* support historical data collection.

Related topics: “Enclave Information workspace” on page 260, “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

Enclaves in Selected Service Class and Period workspace

This workspace provides the same information as the Enclave Information workspace, but for a selected service class and period. The Enclave CPU % by Address Space bar chart shows CPU usage by independent and dependent enclaves for all address spaces owning enclaves in the service class and period.

You can link to this workspace from the WLM Service Class Resources workspace, or from the Enclave Information table in the Enclave Information workspace.

This workspace can record history.

Related topics: “Enclave Information workspace” on page 260, “WLM Service Class Resources workspace” on page 275, “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

Enclaves Owned by Selected Address Space workspace

This workspace provides enclave information for a selected address space.

The Enclave CPU % by Address Space bar chart shows information for all address spaces owning enclaves. The data in the Enclaves Owned by <Address Space Name> table is limited to the enclaves owned by a selected address space.

You can link to this workspace from the Address Space CPU Utilization Summary table in the Address Space Overview workspace, the Address Space CPU Utilization workspace, or the Enclave Information workspace.

This workspace can record history.

Related topics: “Enclave Information workspace” on page 260, “Enclave Table attributes” on page 161, “Address Space Overview workspace” on page 255, “Address Space CPU Utilization workspace” on page 253, “Address Space CPU Utilization attributes” on page 146, “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

Enqueue and Reserve Detail workspace

The Enqueue and Reserve Detail workspace provides detailed information on all the global enqueue conflicts and reserves in the system. It shows all of the owners and waiters for a particular resource within the system, together with their identifying information and descriptive information about the type of resource requested.

You link to this workspace from the Enqueue, Reserve, and Lock Summary workspace.

Related topics: “Enqueue, Reserve, and Lock Summary workspace,” “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

Enqueue, Reserve, and Lock Summary workspace

Enqueues, reserves, and locks are methods of serializing resources to ensure only one updater at a time.

This workspace provides summary information about all global enqueue conflicts and reserves for the system. The workspace contains all enqueue conflicts and reserves in the system, including its contribution to conflicts between systems in those configurations in which enqueue management spans two or more systems. The workspace also provides information on suspend and spin lock.

This workspace contains the following views:

- **Count of Enqueue and Reserve Waiting Address Spaces** bar chart
- **Wait Time of Enqueue and Reserve Waiting Address Spaces** bar chart
- **Enqueue and Reserve Summary** table
- **Suspend Lock Activity** table
- **Spin Locks Held** table
- **Spin Lock Waiters** table

From this workspace you can link to the Enqueue and Reserve Details workspace.

Historical data collection is available for this workspace.

Related topics: “Enqueue and Reserve Detail workspace” on page 261, “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

Inspect Address Space CPU Use workspace

The information in the Inspect Address Space CPU Use workspace allows you to observe where in the executable code a z/OS address space is spending its time. This workspace contains three views:

- Sampling Statistics view

  This view shows the values used by the agent to collect the data. The four columns of this view show the number of samples requested, the interval at which the samples were taken in milliseconds, the number of samples collected and the number of samples used.

  Normally the number of samples collected will be same as the number requested unless the job being inspected ends before the inspect agent has finished collecting data. In this case, the number of samples collected will be the number collected up to the point where Inspect detected that the target job had ended.

  The number of samples used is the number of times that the Inspect agent saw CPU activity in the target address space and gives you some indication as to the statistical accuracy of the resultant Inspect data. The more samples that actually see activity during the time the Inspect agent is running, the more accurate the overall sample will be.The number of samples used value does not represent the number of rows of Inspect data.
- Agent Messages view

  This view displays anyAppendix C, “Inspect messages,” on page 327 returned by the Inspect agent. These messages help to explain the resultant data (or lack thereof) that you see in the other views. For example, if no CPU activity was seen by Inspect in the address space being inspected, the agent would return a message indicating that; the number of samples used column in the Sampling Statistics view would be zero; and no data would be displayed in the Inspect Data view.
- CPU Usage for ASID view

  This view contains data, drilled down to the agent-selected level of granularity within each CSECT for the most active TCBs, sorted in descending order of CPU usage percentage. (If the agent is unable to determine CSECT information for a module, the data is returned for the load module as a whole.) The Inspect agent returns data only for elements for which it saw CPU activity.

  The information in this table helps you to identify where in the code an address space is spending its time. When used in conjunction with link edit and compile or assembly listings, this information can help you to identify looping code or code that may be a candidate for a rework to improve its efficiency.

“Accessing the workspace”
“About the Inspect links”
“Using custom inspect parameters”

Accessing the workspace

You can link to the Inspect Address Space CPU Use workspace using one of two predefined links from the Address Space CPU Utilization table in the Address Space CPU Utilization workspace. When you select the link for a particular address space, the Inspect workspace is populated with data gathered using the parameters specified in the link definition.

The workspace is not populated by data until the Inspect agent completes on the host system. The time it takes to complete is a function of the number of samples and the sampling interval. For example, if 1000 samples are take at a 5-millisecond interval (the default settings), it will take 5 seconds for the data collection process to complete.

About the Inspect links

There are two Inspect links:

- The Inspect Address Space CPU Use link uses the default parameters of 1000 samples at 5 millisecond intervals.
- The Inspect with 5000 samples at 2ms interval link uses the specified parameters, but it can also be modified to use the sample count and sampling interval you specify.

Using custom inspect parameters

Note: The changes you make to the Inspect link cannot be saved unless you save and rename the workspace using Save As. They are deleted when you close the Inspect workspace.

When you are selecting the values for number of samples and the sampling intervals, bear in mind that if the total time taken to execute the agent exceeds the client timeout value in the Tivoli Enterprise Monitoring Server, the Tivoli Enterprise PortalTivoli Enterprise Portal will return no data, even if the agent subsequently completes normally.

To modify the Inspect with custom parameters link:

1. Navigate to the Address Space CPU Utilization workspace by linking from the Address Space Counts table of the Address Space Overview workspace for the target system.
2. Right-click the link icon beside a row in the Address Space CPU Utilization table and select Link Wizard from the pop-up menu.
   The Link Wizard editor appears.
3. In the Selection area, select Inspect with 5000 samples at 2ms interval, then click Next >.
4. In the Link Identity field, type a new name for the link and a description, if desired, then click Next >. Note that changing the name of the link will change its name in the link pop-up menu.
5. In the Properties tree, select the INTERVAL symbol under Query - Inspect Address Space CPU Use. The Expression editor appears instead of the help in the right frame.
6. In the Expression field, type the interval, in milliseconds, at which you want Inspect to collect samples.
7. In the Properties tree, select the SAMPLES symbol under Query - Inspect Address Space CPU Use.
8. In the Expression field, type the number of samples you want Inspect to use in deriving the data.
9. Click Finish to save your changes and close the editor.

The new name is displayed in the pop-up menu when you right-click a link icon in the Address Space CPU Utilization table.

Related topics: Appendix C, “Inspect messages,” on page 327, “Inspect Address Space CPU Use attributes” on page 166, “Address Space CPU Utilization workspace” on page 253, “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

HiperDispatch Details workspace

This workspace provides HiperDispatch-related topology and utilization statistics at the LPAR and individual logical processor level for the LPAR on which an OMEGAMON XE on z/OS agent is running when the LPAR is in HiperDispatch mode. This information can help you determine how your HiperDispatch configuration is performing.

HiperDispatch makes efficient use of CPU-level hardware caching in processors that employ Non-Uniform Memory Access (NUMA) in their designs. NUMA is a means of minimizing memory access latency when retrieving CPU-level cache contents. Efficient use of NUMA architecture requires the dispatching of work, where possible, on the same physical CPU, or another CPU on the same book, to take advantage of cache hits in L1, L1.5 or local L2 caches. Collaboration between the hardware, the LPAR Hypervisor and the Workload Manager Dispatcher is necessary to fully take advantage of what NUMA offers regarding workload performance.

HiperDispatch Management at the z/OS image level can be deployed on one or more of the LPARs in a complex. HiperDispatch is activated for an LPAR by specifying `HiperDispatch=YES` in the SYS1.PARMLIB IEAOPT member. `HiperDispatch=NO`, the default, physical CPUs will be assigned equally across all of the logical CPUs in an LPAR.

For the HiperDispatch feature to be used, the following conditions must be in effect:

- The HiperDispatch parameter in SYS1.PARMLIB(IEAOPT) must be specified YES.
- The hardware must support HiperDispatch.
- The z/OS operating system must support HiperDispatch.

WLM can then determine if the LPAR should be in HiperDispatch mode or not. WLM will place the LPAR in HiperDispatch mode if there is sufficient work in this LPAR to keep 1.5 physical CPUs busy. If demand falls below this threshold WLM will take the LPAR out of HiperDispatch mode.

Two views present information at the LPAR level:

- The LPAR Attributes table displays the status of the HiperDispatch Management (On, Off, or Unavailable), the LPAR Name, the LPAR Cluster Name, the LPAR Group Name.
- The LPAR Weights table displays the current, minimum, and maximum weights by processor type (standard, zAAP, zIIP).

When HiperDispatch Management is on, three additional views display the logical CPU ID, entitlement (High, Medium or Low), percentage entitlement, physical CPU percentage, physical overhead percentage managing the CPU within the LPAR, and a status (Online, Offline, Parked, Park Pending, or Reserved) for each CPU, by logical type.

You can link to this workspace from a row in the System CPU Utilization workspace.

Related topics:

“System CPU Utilization workspace” on page 269

“Organization of system-level predefined workspaces” on page 224

“Attribute groups used by the system-level predefined workspaces” on page 226

LPAR Clusters workspace

An LPAR cluster is a collection of z/OS images (logical partitions, or LPARs) that are carved from the resources of a single Central Processing Complex (CPC) and are part of a single sysplex. Because the LPAR's resources are all derived from a common pool of resources in the CPC, the resources can be redistributed as needed among the LPARs. This permits the operating system's Workload Manager to move resources from LPARs with less important work to LPARs with more important work when that important work is missing its goals.

The LPAR Clusters workspace assists you in answering the following questions:

1. What is the reason that this work cannot obtain sufficient processor resource?
2. Do I have less important work that is meeting or exceeding its service goals? Is the reason that LPARs where this workload is running are obtaining more than their fair share of the processor resource?

This workspace permits you to examine and monitor performance information for the LPARs or clusters associated with a given workload. The workspace contains two charts:

- The **LPAR Logical Utilization** bar chart graphically shows how well each LPAR is able to obtain CPU resources by charting the Effective Weight Percent (that is, actual) and Logical Weight Percent (target) values for a given LPAR. LPARs that are not accomplishing their target weight show an Effective Weight bar shorter than their Logical % Weight bar. The Effective Weight Index attribute in the CPC LPARs Status report is provided as a measure of this comparison.
- The **LPAR Physical Utilization** bar chart graphically shows, for each LPAR, the Physical CPU Percent Busy (actual utilization) graphed against the Current Weight Percent (target utilization). Examine this chart to locate those LPARs having a significant discrepancy between the height of the two bars. These LPARs may warrant further examination. A significant discrepancy is not necessarily an indication that there is a performance problem. The CPU Index attribute in the CPC LPARs Status report is provided as a measure of this comparison.

Tables in this workspace provide information about CPC status, LPAR cluster information, and CPC LPARs status. The first two tables provide CPC and Cluster summary information. The **CPC Status** table also contains information about model capacity IDs and capacity ratings. The default positioning of the capacity columns is at the end of the columns displayed in the **CPC Status** view. If you deploy Capacity Provisioning Manager or On/Off Capacity on Demand, consider reorganizing the sequence of columns so that these columns appear closer to the beginning of the column sequence.

Note: Some LPARs cannot easily be associated with a cluster. These are collected under the umbrella cluster labeled ″N/A″. Some LPAR attributes may not be available or may not apply to one of more of these LPARs. And because ″N/A″ is not a real cluster, it is not included in the **LPARs Clusters** view.

The **CPC LPARs Status** table provides detailed performance information for the LPARs that are configured on the CPC. The **CPC LPARs Status** table has the following predefined thresholds set:

- If LPAR Name not equal to "PHYSICAL" and Effective Weight Index equal to 0.9, you will see a 🟨 warning indicator
- If LPAR Name not equal to "PHYSICAL" and Effective Weight Index less than 0.9, you will see a 🟥 critical indicator

From the **LPAR Cluster** table, you can link to the "LPARs Assigned to a Cluster workspace" on page 266.

From the **CPC Status** table, you can link to the "OMEGAMON for MVS – LPAR PR/SM Processor Statistics workspace" on page 267.

Related topics:

"Organization of system-level predefined workspaces" on page 224

“Attribute groups used by the system-level predefined workspaces” on page 226

LPARs Assigned to a Cluster workspace

This workspace is displayed when you select a specific row in the LPAR Cluster table of the LPAR Clusters workspace. It provides a filtered view of information for the specific cluster you selected.

The workspace contains two charts: the LPAR Logical Utilization (Velocity) chart and the LPAR Physical Utilization chart.

The LPAR Logical Utilization bar chart graphically shows how well each LPAR is able to obtain CPU resources by charting the Effective Weight Percent (that is, actual) and Logical Weight Percent (target) values for a given LPAR. LPARs that are not accomplishing their target weight show an Effective %Weight bar that is shorter than their Logical %Weight bar. The Effective Weight Index attribute in the CPC LPARs Status report is provided as a measure of this comparison.

The LPAR Physical Utilization bar chart graphically shows, for each LPAR, the Physical CPU Percent Busy (actual utilization) graphed against the Current Weight Percent (target utilization). Examine this chart to locate those LPARs having a significant discrepancy between the height of the two bars. These LPARs may warrant further examination. A significant discrepancy is not necessarily an indication that there is a performance problem. The CPU Index attribute in the CPC LPARs Status report is provided as a measure of this comparison.

Tables in this workspace provide information about CPC status, LPAR Cluster status, and statistics for LPARs assigned to the selected cluster. The first two tables provide CPC and Cluster summary information. The third table provides detailed performance information for the LPARS that are configured in the cluster. The Effective Weight Index column has two predefined thresholds.

- If Effective Weight Index equal to 0.9, you will see a 🟨 warning indicator.
- If Effective Weight Index less than 0.8 you will see a 🟥 critical indicator.

You can link to this workspace from the LPAR Clusters table of the LPAR Clusters workspace.

This workspace can record history.

Related topics: “LPAR Clusters workspace” on page 265, “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

OMEGAMON for MVS - CSA Analyzer workspace

This workspace consists of a single terminal emulator view. The OMEGAMON for MVS screen presented in the view is initialized with the common storage area and job name of the row in the source workspace from which the link was selected. This screen displays summary data on common storage utilization collected by the OMEGAMON for MVS CSAA Immediate command. This command can be used to display other information, such as utilization trends at the job or system level, top common storage users for each of the four areas (CSA, SQA, ECSA, ESQA). Use PF1 to learn more about the CSAA command.

You can link to this workspace from all four table views in the “Address Space Common Storage - Orphaned Elements workspace” on page 251.

Related topics:

“Organization of system-level predefined workspaces” on page 224

“Attribute groups used by the system-level predefined workspaces” on page 226

OMEGAMON for MVS - Job Details workspace

This workspace consists of a single terminal emulator view. The OMEGAMON for MVS screen presented in the view is initialized with the job name from the selected row in the source workspace. This screen provides access to detailed information about that job, including:

- A graphic display of task resource contention
- A graphic display of workload contention, by performance group or individual jobs
- A summary of task resource utilization, swap status, proc and step names
- CPU, I/O, and storage utilization and System Resources Manager (SRM) information
- Enqueues owned by this task
- An Inspect sample of CPU use, by TCB, module, and CSECT
- Common service area, system queue area, extended common service area, and extended system queue area allocation by the task as a whole
- Detailed breakdown of virtual storage usage by task

You can link to this workspace from the following workspaces:

- Address Space CPU Utilization Summary view of the "Address Space Overview workspace" on page 255
- Contention (%) by Resource view of the "Address Space Bottlenecks Summary workspace" on page 250
- Address Space CPU Utilization view of the "Address Space CPU Utilization workspace" on page 253
- Address Space Real Storage view of the "Address Space Storage workspace" on page 256

Related topics:

"Organization of system-level predefined workspaces" on page 224

"Attribute groups used by the system-level predefined workspaces" on page 226

OMEGAMON for MVS – License Manager MSU and WLM Capping workspace

This workspace consists of a single terminal emulator view. The OMEGAMON for MVS screen in the view displays data collected by the XRMSU minor command of the SYS major command. This information includes License Manager MSUs (millions of service units per hour) consumed by the LPAR. The amount of service, capped and uncapped, is displayed at the summary level and with a detailed breakdown of the 4-Hour Rolling Average MSU value used by Workload Manager to implement soft-capping, if a Defined LPAR Capacity limit is in effect. You can use the PF1 key to obtain help for this screen.

You can link to this workspace from the System CPU Utilization table view of the "System CPU Utilization workspace" on page 269.

Related topics:

"Organization of system-level predefined workspaces" on page 224

"Attribute groups used by the system-level predefined workspaces" on page 226

OMEGAMON for MVS – LPAR PR/SM Processor Statistics workspace

This workspace consists of a single terminal emulator view. The OMEGAMON for MVS screen displayed in the view shows LPAR information for all logical partitions and for LPAR management overhead, collected by the LPAR command. If supported by the hardware, a number of statistics relevant to On/Off Capacity on Demand (OOCoD), Capacity Provisioning Management (CPM), and Capacity Backup (BU) may be displayed. Use PF1 to get help with the information displayed.

You can link to this workspace from the CPC Status table view of the “LPAR Clusters workspace” on page 265.

Related topics:

“Organization of system-level predefined workspaces” on page 224
“Attribute groups used by the system-level predefined workspaces” on page 226

Operator Alerts workspace

This workspace permits you to monitor various conditions in your enterprise that can raise operator alerts. Some of these conditions are informational, while others require that the operator take some action to avoid or resolve a problem. The bar charts in the upper right hand corner of the workspace graphically illustrate the information provided in columns of the Operator Alerts table view.

This workspace can record history.

Related topics: “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

Page Dataset Activity workspace

This workspace provides information about availability and response time for a specific page dataset. Page datasets are auxiliary storage datasets that back up all frames of virtual storage. They must be large enough to contain all common and private virtual storage. Page datasets are used when demand for real storage is greater than the space available and data has been paged out. The process of bringing in data is called a page-in and is coordinated by the auxiliary storage manager (ASM). If swap datasets are not defined, page datasets also contain the swapped out part of an address space.

Because the process of paging is very slow when compared to referencing data from real or expanded storage, it is important that page dataset devices be isolated from contention with other kinds of work. This is especially true if there is contention for real and expanded storage, and the page fault rate is high.

The Percent Full and Response Time bar charts in this workspace provide visual representations of the availability of space in the various types of page datasets and the response times for those datasets.

This workspace can record history.

Related topics: “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

Real Storage workspace

This workspace provides information about the use of real storage on your z/OS system in terms of various types of frame counts. The workspace contains four views.

The first bar chart in the workspace illustrates the system level unreferenced interval count. The unreferenced interval count is inversely related to contention for real storage. The lower the unreferenced interval count, the more quickly frames are being referenced. In some cases, a low count is not necessarily indicative of a paging problem. To determine whether real storage is being overused, use this factor in conjunction with the page fault rate displayed in the “System Paging Activity workspace” on page 270.

The second bar chart illustrates available frames. This chart gives you a sense of how much uncommitted storage is available for use.

The Real Storage table view provides information about the factors that may be affecting performance as determined by the number of frames allocated real storage by storage type.

The Real Storage Summary table view provides summary data for all types of real storage.

This workspace can record history.

You can access this workspace directly from the Navigator.

Related topics: “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

Service Call Performance workspace

The Service Call Performance workspace shows bar charts of top ten service calls by Arrival Rate, Service Time, Pending, and Bytes Processed. The lower right table shows details of all 78 service calls monitored.

You can link to this workspace from the following workspaces:

- The Service Call Performance by System table of the Cross-System Cryptographic Coprocessor Overview workspace
- The Top User Performance workspace
- The Cryptographic Services workspace

From this workspace you can link to the following workspaces:

- The Top User Performance workspace
- The Cryptographic Services workspace
- The Cross-System Cryptographic Coprocessor Overview workspace

This workspace can collect history data.

Note: Since the attributes that report data in this workspace report rolling averages of 10 minutes of collected data, historical sampling should occur at 5-minute intervals (half the length of the data averaging interval).

Related topics:

“Cross-system Cryptographic Coprocessor Overview workspace” on page 231

“Cryptographic Services workspace” on page 258

“Top Users Performance workspace” on page 271

“Organization of system-level predefined workspaces” on page 224

“Attribute groups used by the system-level predefined workspaces” on page 226

System CPU Utilization workspace

System CPU utilization is the percentage of time that all processors available to a system were busy dispatching work. This workspace provides information about CPU usage for the selected system. In an LPAR environment, this group also provides partition management statistics.

The System CPU Utilization workspace contains the following views:

- The **System CPU Utilization** table view shows the number of physical processors reported on, number of processors online, the average percentage of time that all processors collectively available in this z/OS system were busy dispatching work, and other information specific to CPU workload and partition workload.

  The System CPU Utilization table also shows the long-term average CPU service used by this LPAR in millions of service units (MSUs) per hour. Data for this attribute is available only if you use defined capacity as a basis for pricing.

If CPU utilization information for an address space requires analysis, see “Address Space CPU Utilization workspace” on page 253.

If the HiperDispatch Management mode is on, you can see details about utilization by linking to the “HiperDispatch Details workspace” on page 264.

If overall system CPU utilization is too high, you may want to review

– how competition for the CPU has affected a workload’s performance

– CPU use throughout the day

In a complex with more than one CPU, the utilization values can be computed based on the total number of processors, or normalized to a maximum of 100%. From this view, you can link to the “OMEGAMON for MVS – License Manager MSU and WLM Capping workspace” on page 267.

- The **Workload CPU Usage** bar chart provides information on CPU workload, showing the percentage of time spent doing SRB work and TCB work, as well as z/OS overhead. The chart permits you to understand what part of your workload is having the greatest impact on CPU resources.
- The **Partition CPU** bar chart illustrates the percentage of time on average that a logical processor was dispatched in the logical partition during the interval, and the percentage of time on average that a physical processor was dispatched in the logical partition during the interval. For a given partition, when the number of logical processors assigned to a partition is the same as the number of physical processors, the partition’s physical and logical CPU utilization should be the same.

This workspace can record history.

Related topics:

“Organization of system-level predefined workspaces” on page 224

“Attribute groups used by the system-level predefined workspaces” on page 226

System Paging Activity workspace

The System Paging Activity workspace provides information about factors that affect system paging rates. Because the process of paging is very slow when compared to referencing data from real or expanded storage, it is important that page dataset devices be isolated from contention with other kinds of work. This is especially true if there is contention for real and expanded storage, and the page fault rate is high.

Bar charts on this workspace graphically illustrate the factors enumerated in the table view. In the Page Rates bar chart, a high page fault rate (the number of page-ins per second) may indicate contention for storage. A high system page rate may indicate high paging activity over one or more of the storage areas (common, system, or private) resulting in storage contention.

A high number of expanded storage pages moved may indicate that expanded storage is being depleted. This may cause increased z/OS overhead.

The unreferenced interval count is inversely related to contention for real storage. The lower the unreferenced interval count, the more quickly frames are being referenced. In some cases, a low count is not necessarily indicative of a paging problem. To determine whether real storage is being overused, use this factor in conjunction with the page fault rate.

The auxiliary storage manager (ASM) is the z/OS component responsible for managing page dataset I/O. A large ASM queue can negatively impact system performance.

This workspace can record history.

Related topics: “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

Tape Drives workspace

The Tape Drives workspace provides information about various factors that affect the performance and availability of the tape drives on your system.

Two bar charts, Tape Errors and I/O Counts, visually illustrate factors affecting performance of specific devices charted by device address. If the number of tape errors for a device is high, the device may need to be cleaned or adjusted, or the tape may be bad.

The Tape Drive Overview thermometer chart illustrates the total number of devices in each of the conditions shown: DDR Swaps in Progress, Dropped Ready, Not Responding. The thermometer chart also shows the total number of Tape Mounts Pending on your system.

A Dropped Ready condition may indicate power supply or other hardware problems, or that the switch may have been turned off at the device.

Dynamic Device Reconfiguration (DDR) permits tapes to be remounted on another drive without ABENDing a running job. Unless acted upon promptly, performing a DDR swap can create severe system I/O performance degradation.

This workspace can record history.

Related topics: “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

Top Users Performance workspace

The Top Users Performance workspace shows bar charts for top ten users by Arrival Rate, Service Time, Pending, and Bytes Processed. The lower right table shows details of up to 32 top users on the current system.

You can link to this workspace from the Top Users by System view of the Cross-System Cryptographic Coprocessor workspace, or from the Service Call Performance or the Cryptographic Services workspace.

This workspace can collect history data.

Note: Since the attributes that report data in this workspace report rolling averages of 10 minutes of collected data, historical sampling should occur at 5 minute intervals (half the length of the data averaging interval).

Related topics:

“Cross-system Cryptographic Coprocessor Overview workspace” on page 231
“Service Call Performance workspace” on page 269
“Cryptographic Services workspace” on page 258
“Organization of system-level predefined workspaces” on page 224
“Attribute groups used by the system-level predefined workspaces” on page 226

UNIX BPXPRMxx Values workspace

As part of the customization of UNIX System Services, you can create BPXPRM*xx* members in SYS1.PARMLIB (where *xx* is a suffix used to distinguish between various customized PARMLIB members) to establish defaults for UNIX System Services.

The UNIX BPXPRMxx Values workspace reports three types of data about these members and their contents:

- A list of active BPXPRM*xx* member suffixes from SYS1.PARMLIB.

- Contents of the BPXPRM*xx* members, with comments eliminated, override rules applied to duplicates, and each individual statement on a separate row of the table view
- Active individual values of many of the system parameters set by BPXPRM*xx* members or overridden by z/OS console commands

Use this workspace to see at a glance what values are in use, and to compare active values with what is specified in SYS1.PARMLIB. Because the table view that is displayed resolves the somewhat complicated override rules of the BPXPRM*xx* members, it shows more clearly the values that are being used and is more easily understood than looking at the BPXPRM*xx* members directly.

You can link to this workspace from the “z/OS UNIX System Services Overview workspace” on page 278.

This workspace can report historical data.

Related topics: “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

UNIX Files workspace

The UNIX Files workspace lists file attributes, path, and time information for the datasets in your monitored systems. Use this report to obtain detailed information about a specific directory or file.

The table views contained in this workspace displays directories and files. If you link to this workspace from the UNIX Kernal workspace, it displays the files and directories in the root directory. To view the files in a subdirectory, right-click in the appropriate row in the Directories in / table and select *Link to > Subdirectory* from the pop-up menu.

You can link to this workspace from the “z/OS UNIX System Services Overview workspace” on page 278 and the “UNIX Mounted File Systems workspace” on page 273. If you link to the workspace from the UNIX Mounted File Systems workspace, you see only the directories and files for the file system using the selected mount point.

From this workspace you can link to the “UNIX Mounted File Systems workspace” on page 273 to see file information for the mount point.

UNIX HFS ENQ Contention workspace

The UNIX Hierarchical File System ENQ Contention workspace shows the enqueue contentions, if any, for the given file system. Use this workspace to investigate Hierarchical File System (HFS) enqueue problems. This workspace permits you to see who is holding a resource and who is waiting for the resource. It also shows the amount of time contention for the resource has been in effect.

You can link to the HFS Enqueue Contention workspace from the table view of the “UNIX Mounted File Systems workspace” on page 273 when an enqueue exists.

Historical data collection is not available for this workspace.

Related topics:“Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

UNIX Kernel workspace

The UNIX Kernel workspace displays information about the kernel address space.

This workspace contains four views:

- The UNIX Identity table view provides the system name, version, and release number of the selected system.

- The Process Utilization bar chart show graphically the number of processes currently utilized by the system and the maximum number used during the reporting interval.
- The UNIX Kernal table provides performance and usage statistics.
- The Take Action view allows you issue commands directly to the UNIX system.

You can link to the the UNIX Kernel workspace from the "z/OS UNIX System Services Overview workspace" on page 278.

This workspace *cannot* display historical data.

Related topics: "Organization of system-level predefined workspaces" on page 224, "Attribute groups used by the system-level predefined workspaces" on page 226

UNIX Logged-on Users workspace

The UNIX Logged-on Users workspace supplies terminal identification and time data for each user logged on to your monitored systems. You can also use the information in this workspace to report process information for individual users. To see information about the process owned by a user, right-click the row of data for the user in the UNIX System Services Logged-on Users table view. If the user logged in with root authority, this view will contain all processes descended from users having root authority.

Use this workspace to obtain information about individual users and to check for problems caused by user behavior. For example, a high Idle Time value for a remote user may indicate a potential security problem.

You can link to this workspace from the "z/OS UNIX System Services Overview workspace" on page 278.

From this workspace to a "UNIX Processes workspace" on page 274 for the single process which represents a logged on user, or to one which shows all processes related to a logged-on user's UID. (A process can have more than one UID associated with it. For this navigation, a process is considered related to the logged-on user if the logged-on user's UID is the same as the process's ″real UID.″

This workspace can display historical data.

Related topics: "UNIX Processes workspace" on page 274, "Organization of system-level predefined workspaces" on page 224, "Attribute groups used by the system-level predefined workspaces" on page 226

UNIX Mounted File Systems workspace

The UNIX Mounted File System workspace contains basic status and activity information for the UNIX System Services mounted file systems. It contains two views.

The UNIX Mounted File Systems table displays information that indicates the path to a given mounted file system and the data set and volume on which it physically resides. Other information includes the size of the system and how much of the system is currently in use. The Mounted File Systems with ENQ Contention table shows those file systems experiencing enqueue contention, if any, for SYSDSN and SYSZDSN resoucres.

Use this workspace to determine which file systems are actually mounted and what state they are in. For example, it is possible to see those file systems having very little free space or having lots of wasted space. You can also see when enqueue contentions occur and can navigate from this workspace to the HFS Enqueue Contention workspace to get more detailed information.

You can link to this workspace from the "z/OS UNIX System Services Overview workspace" on page 278, the "UNIX Files workspace" on page 272, and the "UNIX Processes workspace" on page 274. If you link to the workspace from the UNIX Files workspace, you see information for the mounted file system for the selected directory or file.If you link from the UNIX Processes workspace, you see information for the mount point for the selected process.

From a row in the UNIX Mounted File Systems table, you can link to a “UNIX Files workspace” on page 272 for the mountpoint, a “UNIX Processes workspace” for processes using the file system, and one of two “UNIX HFS ENQ Contention workspace” on page 272s, one for SYSDSN enqueues and one for SYSZDSN enqueues. Links to the contention workspaces are visible only when an enqueue exists.

This workspace can collect historical data.

UNIX Processes workspace

The UNIX Processes workspace provides detailed information about each process your systems are currently executing.

The UNIX Processes workspace contains two views. The UNIX Run Time % bar chart graphically illustrates the percentage of CPU execution for UNIX work. The UNIX Processes table displays identifying and performance information for each process.

Use this workspace to check the status of processes and identify the ones causing system problems. For example, you may find that the execution time of a process is abnormally high. A problem program that is looping might be causing this problem. You can use this workspace to identify the process and its owner.

You can link to the UNIX Processes workspace from the “z/OS UNIX System Services Overview workspace” on page 278, the “Dubbed Address Spaces workspace” on page 259 (for information about the processes associated with a selected address space), the “UNIX Mounted File Systems workspace” on page 273 (for processes using the file system), and the “UNIX Logged-on Users workspace” on page 273 (for user ID and process IDs).

From this workspace you can link to information about children of a selected process, the threads of the process, the address space in which the process is executing, and the mounted file systems being used by the process.

This workspace can record historical data.

Related topics:

“Dubbed Address Spaces workspace” on page 259,
“UNIX Mounted File Systems workspace” on page 273,
“UNIX Threads workspace,”
“Organization of system-level predefined workspaces” on page 224,
“Attribute groups used by the system-level predefined workspaces” on page 226.

UNIX Threads workspace

The UNIX Threads workspace shows all threads for a given process. Use this report to determine which thread or threads use the most CPU time in a busy process.

You can navigate to this workspace only from the Process table view by linking from the row of the particular process for which you want to see thread data.

This workspace can collect historical data.

Related topics: “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

User Response Time workspace

Note: The End to End Response Time collector, supplied with IBM Tivoli OMEGAMON XE products, must be started for this workspace to display data.

The User Response Time workspace provides an overview of the activities of system users logged in through VTAM using SNA, arranged according to user ID. The User Activity Overview table view identifies, for each user ID, the number of transactions a given user has executed and the total response time for the transactions. The total response time is broken down into the average response time attributable to traffic on the network and average response time for host processing.

Use this workspace to identify bottlenecks in your system that may be due to processor or network problems or due to poorly designed transactions. For example, if a user is running a low number of transactions but is experiencing very high host or total response time, it would be advisable to examine the transactions to understand the reasons for the bottleneck.

Two bar charts, User Response Times and Transaction Counts, illustrate the contents of the table view's columns.

This workspace can record history.

Related topics: "Organization of system-level predefined workspaces" on page 224, "Attribute groups used by the system-level predefined workspaces" on page 226

WLM Service Class Resources workspace

Use the WLM Service Class Resources table view to determine how well your defined service classes are performing against the goals you have set for their performance. Use the Performance Index to see how well services are meeting their goals. If the z/OS V1.9 Workload Manager blocked workload capability is enabled, this table will also show the percentage of time service class workloads ran at a promoted dispatching priority.

In addition to the table view, the workspace contains three bar charts that illustrate the information provided in various columns of the WLM Service Class Resources table view:

- The CPU Percentage bar chart permits you to determine those jobs that are consuming the largest percentage of processing time and to compare the processing time consumed by given jobs.
- The I/Os per Second bar chart illustrates the contents of the I/O Rate column in the workspace's table view. This bar chart permits you to determine which jobs are performing the largest amount of I/O and to determine those jobs that are responsible for excessive I/Os.
- The Average Storage Use in Bytes chart illustrates the contents of the Average Storage column in the table view and permits you to see at a glance the amount of storage consumed by a given job.

You can link to this workspace from a row in the Enclave Information table view in the Enclave Information workspace. In this case, information is displayed for the service class period in which the navigated-from enclave is executing.

From a row in the WLM Service Class Resources table view, you can navigate to

- a workspace to view summary information about the bottlenecks that are occurring for the selected service class period
- a workspace that contains information about those enclaves executing in the selected service class and period
- a workspace to view the CPU utilization of those address spaces executing in the selected service class and period

This workspace can record history.

Related topics: “Enclaves in Selected Service Class and Period workspace” on page 261, “Address Space Bottlenecks in Service Class Period workspace” on page 249, “Address Space CPU Usage Class and Period workspace” on page 252, “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

WLM Service Class Information for Selected Address Space workspace

This workspace provides service class period information for a selected address space. This workspace contains four views:

- The CPU Percentage bar chart enables you to determine the percentage of processing time consumed by a given service class.
- The I/Os per Second bar chart illustrates the information in the I/O Rate column in the workspace’s table view. This bar chart enables you to determine the amount of I/O consumed by a given service class.
- The Average Storage Use in Bytes chart illustrates the contents of the Average Storage column in the table view and permits you to see at a glance the amount of storage consumed by a given service class.
- The Service Class and Period Information for ASID: <address_space_ID> table view presents information that enables you to determine how well the service class periods are performing against the goals set for their performance.

You can link to this workspace from:

- the Address Space CPU Utilization Summary view in the Address Space Overview workspace. (Select Service Class Period Information from the link pop-up.)
- the Address Space CPU Utilization table view in the Address Space CPU Utilization workspace.

This workspace can record history.

Related topics: “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

Workloads for Classes workspace

The Workloads for Classes workspace contains individual report classes that you can open to obtain information.

Incoming work can be assigned a report class as well as a service class. Even though workload manager uses the service class to manage the workloads, report classes can be used for additional reporting across workloads or service classes. Workload Management provides data for reporting at a service class level. You can extend this reporting capability to report on similar types of work units, even if the work units are members of different service classes, by assigning these work units to a report class.

Related topics: “Organization of system-level predefined workspaces” on page 224, “Attribute groups used by the system-level predefined workspaces” on page 226

zFS Overview workspace

As its name indicates, this workspace provides an overview of the system-level activity on the z/OS zSeries File System (zFS). zFS is a file system that can be used in addition to the “legacy” UNIX System Services Hierarchical File System (HFS).

This workspace contains six views:

- **Average Operation Wait Time** bar chart
- **zFS Kernel Summary** table

- **zFS Storage Summary** table
- **zFS Directory Cache** table
- **zFS Metadata Cache** table
- **zFS Kernel Detail** table

Together, these views give you a good overall picture of the health of the zFS workload and indicates what zFS operations are taking the longest.

You can link to this workspace from the “z/OS UNIX System Services Overview workspace” on page 278. From the zFS Kernel Summary view you can link to the “zFS User Cache workspace.”

Related topics:

“Organization of system-level predefined workspaces” on page 224

“Attribute groups used by the system-level predefined workspaces” on page 226

“zFS User Cache workspace”

“z/OS UNIX System Services Overview workspace” on page 278

zFS User Cache workspace

This workspace contains the following six views which contain data that can be used to adjust configuration parameters for the user cache.

- **zFS User Cache Statistics** bar chart
- **External Requests** table
- **File System Reads** table
- **File System Writes** table
- **Page Management** table
- **Dataspaces** table

You can link to this workspace from the zFS Overview workspace.

Related topics:

“Organization of system-level predefined workspaces” on page 224

“Attribute groups used by the system-level predefined workspaces” on page 226

“zFS Overview workspace” on page 276

z/OS System Overview workspace

The z/OS System Overview workspace is the default workspace for each z/OS managed system item in the Navigator. This workspace contains six views that summarize key performance aspects of the LPAR:

- Workload CPU Usage bar chart view. This chart shows overall CPU utilization for the LPAR, including both standard CPU and special processors. The view provides a link to the “System CPU Utilization workspace” on page 269.
- Common Storage bar chart view. This chart shows utilization for common service area (CSA), extended common service area (ECSA), System queue area (SQA), and extended system queue area (ESQA). The view provides a link to the “Common Storage workspace” on page 258.
- System Event Console view. This view shows raised situations originating from this managed system.
- Address Space CPU Utilization table view. This table shows CPU usage by each address space in this LPAR, including information about usage of each of the specialty engines, enclave activity, and job level CPU usage, as well as short term usage. The view provides a link to the “Address Space CPU Utilization workspace” on page 253.

- Active Users of Common Storage table view. This table shows how common storage is being used by each address space in the LPAR. The view links to the “Address Space Common Storage - Active Users workspace” on page 250.
- Enqueue and Reserve Summary table view. This table summarizes Major and Minor enqueue names that have contention or are the target of reserve requests. The view provides a link to the “Enqueue, Reserve, and Lock Summary workspace” on page 261.

Related topics:

“Organization of system-level predefined workspaces” on page 224

“Attribute groups used by the system-level predefined workspaces” on page 226

z/OS UNIX System Services Overview workspace

As its name indicates, this workspace provides an overview of the information provided about z/OS UNIX System Services running on a z/OS image. The workspace contains six views, each showing selected attributes for a different attribute group:

- The Dubbed Address Spaces table view shows the address space ID, name and the percentage of central processor time (CPU) used for z/OS address spaces marked as being users of UNIX System Services requests.
- The UNIX Processes table view shows the command name, job name, run time percentage, process status, and execution state for all currently running processes.
- The UNIX Kernel table view shows selected statistics for the kernel address space and zFS.
- The Process Utilization bar chart shows the number of processes currently utilized by the system and the max number used during the reporting interval.
- The UNIX Logged-On Users table view lists provides the login name, user name, and login time for all users logged on to the monitored system.

From the UNIX System Services Overview workspace, you can link to the following workspaces:

“Address Space CPU Usage Details workspace” on page 252 (filtered on the OMVS address space name)

“Address Space CPU Usage Details workspace” on page 252 (filtered on the zFS address space)

“Dubbed Address Spaces workspace” on page 259 (the default link)

“UNIX BPXPRMxx Values workspace” on page 271

“UNIX Files workspace” on page 272

“UNIX Kernel workspace” on page 272

“UNIX Logged-on Users workspace” on page 273

“UNIX Mounted File Systems workspace” on page 273

“UNIX Processes workspace” on page 274

“zFS Overview workspace” on page 276

Related topics:

“Organization of system-level predefined workspaces” on page 224

“Attribute groups used by the system-level predefined workspaces” on page 226

Chapter 12. Situations

OMEGAMON XE on z/OS provides predefined situations which check for system conditions that are typically considered to be problematic or noteworthy and trigger ⊗ Critical or ⚠ Warning event indicators in the Navigator.

Some situations are shipped with very high or very low values, which essentially disable them. You should examine these situations and customize them with values that are meaningful for your installation.

You can examine the conditions or values being monitored using the Situation editor and, if necessary, change them to ones better suited to your environment. You can also use the predefined situations as models for creating your own situations using OMEGAMON XE for z/OS attributes.

Note: If you choose to make changes to a predefined situation, you should change a copy and preserve the original situation in the form in which it was shipped.

Related topics: “Activating predefined situations”

Activating predefined situations

Two of the predefined situations shipped with OMEGAMON XE on z/OS are auto-started and autodistributed: KM5_No_Sysplex_DASD_Filter_Warn and KM5_Weak_Plex DASD_Filter_Warn. This means that these situations are set to run at when the Tivoli Enterprise Monitoring Server starts up and distributed to the MVS_SYSPLEX managed system list. The Tivoli Enterprise Monitoring Server automatically populates this managed system list with each online sysplex proxy node and these situations begin monitoring when the Tivoli Enterprise Monitoring Server starts.

All the other predefined situations must be activated.

To activate a situation, you distribute (assign) it to one or more managed systems or managed system lists and start it. If you want the situation to continue running across Tivoli Enterprise Monitoring Server restarts, check the Run at startup option.

For instructions for activating situations, see “Activating situations” on page 22 or the Tivoli Enterprise Portal online Help.

Sysplex situations

The sysplex situations are described below in alphabetical order. You can access the description of a specific situation by selecting its name in the Contents tab.

KM5_Model_Sysplex_DASD_Filter

KM5_Model_Sysplex_DASD_Filter uses the DASD Device Collection Filtering attributes Average Response time and I/O Rate, but does not provide any preset values. This situation is intended for use as the basis for creating your own DASD filter. The formula is:

```
*VALUE DASD_Device_Collection_Filtering.Average_Response_Time *GT 0 .0
*AND *VALUE DASD_Device_Collection_Filtering.I/O_Rate *GT 0.0
```

No DASD device data is collected unless a DASD filter situation is activated. To create a DASD filter, use the Situation editor to create a copy of this situation and specify values appropriate to your environment. You can further refine the filter criteria using other attributes in the DASD Device Collection Filtering attribute group. Set the situation to autostart, and distribute it to the *MVS_SYSPLEX managed system list.

KM5_No_Sysplex_DASD_Filter_Warn

KM5_No_Sysplex_DASD_Filter_Warn is raised when no data is being reported. This indicates either that no DASD filter situation is being used or that the filter being used is too strong and all devices are eliminated by the criteria. The formula is:

```
*IF *COUNT Sysplex_DASD_Group.Volume_Serial_Number *EQ 1
*AND *VALUE Sy splex_DASD_Group.Volume_Serial_Number *EQ $none$
```

This situation is set to run at startup and distributed to the *MVS_SYSPLEX managed system list.

If there is no sysplex DASD filter in place and you want to start monitoring DASD information, add a sysplex DASD situation. You can use the product-provided situation KM5_Model_Sysplex_DASD_Filter as the basis for your filtering situation. To use the model:

- Adjust the threshold for Average Response Time and I/O Rate to appropriate values for your environment.
- Use the **Save as** option to save the situation with a new name, then distribute it to the *MVS_SYSPLEX managed system list.
- On the **Formula** tab, select the desired sampling interval.
- Set the situation to run at startup.

If there is a sysplex DASD filter in use, weaken the filtering criteria to allow some devices to be monitored.

KM5_Weak_Plex_DASD_Filter_Warn

KM5_Weak_Plex_DASD_Filter_Warn monitors for a high count of DASD devices, which indicates either that the DASD filtering criteria are not strong enough or that the threshold needs to be raised to reduce the number of reported devices. The provided threshold is 1000, but you change the threshold to avoid false positive results. The formula is:

```
*COUNT Sysplex_DASD_Group.Volume_Serial_Number *GT 1000
```

This situation is auto-started and distributed to the *MVS_SYSPLEX managed system list.

MVS_CFStruct_Status_Crit

MVS_ CFStruct_ Status_ Crit monitors the status of coupling facility structures and raises a Critical alert when a structure has failed.

The formula is:

```
VALUE CF_Structures.Structure_Status EQ Failed
```

If this situation is raised, an immediate rebuild is necessary. If another coupling facility is available as a backup, it should be made the primary facility.

MVS_CFStruct_Status_Warn

MVS_ CFStruct_ Status_ Warn monitors the status of coupling facility structures and raises a Warning alert when a structure is in “Rebuild” status.

The formula is:

```
VALUE CF_Structures.Structure_Status GT 0X14
```

If this situation is raised, a coupling facility structure is about to fail. A rebuild may soon be necessary. If another coupling facility is available as a backup, it should be made the primary facility.

MVS_CFStructStat_FalseLock_Crit

MVS_ CFStructStat_ FalseLock_ Crit monitors the statistics related to coupling facility structures and raises a Critical alert when the false lock table contention count reaches a high level.

The formula is:

```
VALUE CF_Structures.False_Lock_Table_Entry_Contention_Count  GT 100
```

False contention should be less than half of total lock contention or less than 0.5% of total lock requests. If this condition occurs frequently, the size of the lock structure should be increased.

MVS_CFStructStat_FalseLock_Warn

MVS_CFStructStat_FalseLock_Warn monitors the statistics related to coupling facility structures and raises a Warning alert when false contention for lock structures is approaching a high level.

The formula is:

```
VALUE CF_Structures.False_Lock_Table_Entry_Contention_Count  GT 0 AND
VALUE CF_Structures.False_Lock_Table_Entry_Contention_Count LE 100
```

False contention should be less than half of total lock contention or less than 0.5% of total lock requests. If this condition occurs frequently, the size of the lock structure should be increased.

MVS_CFStructToMVS_Requests_Crit

MVS_ CFStructToMVS_Requests_Crit monitors the number of requests that the coupling facility makes to a given z/OS system and raises a Critical alert when the coupling facility performance has deteriorated severely.

The formula is:

```
VALUE CF_Structure_to_MVS_System.Average_Queued_Request_Time  GT 100 OR
VALUE CF_Structure_to_MVS_System.Request_Rates *GT 1000
```

If this situation is raised, check the status of the coupling links and the performance data for the coupling facility machine or LPAR. A link may be down or a processor maybe offline or malfunctioning. If the coupling facility is in an LPAR with shared processors, the overall CPC may be fully loaded, thus causing the LPAR weights to come into effect. If more LPs are assigned to the CF than its share of total weight can support, try varying an LP offline.

MVS_ CFStructToMVS_Requests_Warn

MVS_ CFStructToMVS_ Requests_Warn monitors the number of requests that the coupling facility makes to a given z/OS system and raises a Warning alert when the coupling facility performance has deteriorated to the extent that requests may begin to time out or otherwise fail.

The formula is:

```
(VALUE CF_Structure_to_MVS_System.Average_Queued_Request_Time  GT 20.0 AND
VALUE CF_Structure_to_MVS_System.Average_Queued_Request_Time LE 100.0)  OR
(VALUE CF_Structure_to_MVS_System.Request_Rates GT 500.0 AND
VALUE CF_Structure_to_MVS_System.Request_Rates LE 1000.0) OR
VALUE CF_Structure_to_MVS_System.Requests_Converted GT 50
```

If this situation is raised, you may have exhausted the current capacity. Check the status of the coupling links and the performance data for the coupling facility machine or LPAR.

MVS_CFStructUsers_Connect_Crit

MVS_CFStructUsers_Connect_Crit monitors the connectivity status of coupling facility structure users and raises a Critical alert when a connection has failed and a rebuild is not working.

The formula is:

```
VALUE CF_Clients.Connection_Status EQ Failing
```

If another coupling facility is available as a backup, it should be made the primary facility. Check the links and the status of the coupling facility hardware.

MVS_CFStructUsers_Connect_Warn

MVS_CFStructUsers_Connect_Warn monitors the connectivity status of coupling facility structure users and raises a Warning alert when the connection status indicates a reconnection, or when there are any problem connections ("Failed" or "Failed Persistent").

The formula is:

```
VALUE CF_Clients.Connection_Status EQ ConnectedRebuild  OR
VALUE CF_Clients.Connection_Problem_Flag GT 0
```

If this situation is raised, it indicates that coupling facility connections are failing and a rebuild is in effect. If another coupling facility is available as a backup, it should be made the primary facility. Check the links and the status of the coupling facility hardware.

MVS_CFSystems_Performance_Crit

MVS_CFSystems_Performance_Crit monitors the performance of coupling facility systems and raises a Critical alert when CPU utilization exceeds 95%, an indicator shows coupling facility system status as "Failed", or the I/Os per second exceed 500.

The formula is:

```
VALUE CF_Systems.CPU_Percent GT 95 OR
VALUE CF_Systems.Status EQ Failed OR
VALUE CF_Systems.I/Os_Per_Second GT 5000
```

If this situation is raised, it may indicate that coupling facility is in an overload condition. Check the performance data for the coupling facility and, if it is unacceptable, determine if some large structures can be replicated in another facility to reduce the data rate. The long term solution may require an increase in coupling facility capacity or a technology upgrade to faster coupling links.

MVS_CFSystems_Performance_Warn

MVS_CFSystems_Performance_Warn monitors the performance of coupling facility systems and raises a Warning alert when:

- CPU utilization is between 80% and 95%
- a path is not operational
- the system status is ”Reconcile“
- I/O rates are between 200 and 500 per second or structures out-of-policy count is greater than 0.

The formula is:

```
(((VALUE CF_Systems.CPU_Percent GT 80 AND
VALUE CF_Systems.CPU_Percent LE 95) OR
VALUE CF_Systems.Status EQ Reconcile) OR
(VALUE CF_Systems.I/Os_Per_Second GT 2000 AND
VALUE CF_Systems.I/Os_Per_Second LE 5000)) OR
VALUE CF_Systems.Structure_Count_Out_Policy GT 0
```

If this situation is raised, it may indicate that coupling facility is in an overload condition. Check the performance data for the coupling facility and, if it is unacceptable, determine if some large structures can be replicated in another facility to reduce the data rate. The long term solution may require an increase in coupling facility capacity or a technology upgrade to faster coupling links.

MVS_CFSystemPaths_Busy_Crit

MVS_CFSystemPaths_Busy_Crit monitors the number of times a system path was busy when a request was made and raises a Critical alert when a path busy condition occurs greater than 95% of the time.

The formula is:

```
VALUE CF_Path.Contention_Percent GT 95
```

Severe performance degradation will continue until action is taken to reduce the traffic or increase the data bandwidth to the coupling facility. Traffic may be reduced by adding coupling facilities; bandwidth may be increased by using more links or faster links. See the IBM Redbook OS/390 Parallel Sysplex® Configuration Volume 2: Cookbook , publication number SG24-5638, for details and further recommendations.

MVS_CFSystemPaths_Busy_Warn

MVS_CFSystemPaths_Busy_Warn monitors the number of times a system path was busy when a request was made and raises a warning when a path busy condition occurs between 60% and 95% of the time, or a path is not operational.

The formula is:

```
(VALUE CF_Path.Contention_Percent GT 60 AND
VALUE CF_Path.Contention_Percent LE 95) OR
VALUE CF_Path.Status EQ NotOperational
```

This situation is raised when coupling facility path contention is high, or a link is down. Performance degradation may continue until action is taken to reduce the traffic or increase the data bandwidth to the coupling facility. Traffic may be reduced by adding coupling facilities; bandwidth may be increased by using more links or faster links. See the IBM Redbook OS/390 Parallel Sysplex Configuration Volume 2: Cookbook , publication number SG24-5638, for details and further recommendations.

MVS_GRS_RespTime_Crit

MVS_ GRS_ RespTime_ Crit monitors the response time in the global resource serialization (GRS) complex and raises a Critical alert when response time exceeds the specified threshold, 25 milliseconds by default.

The formula is:

```
VALUE GRS_Ring.Response_Time *GT 25
```

If no hardware problem is indicated, this situation indicates that there is probably a bottleneck in one of the connected systems, causing the GRS token to become stalled. If the sysplex is a parallel sysplex, this kind of problem can be eliminated by changing to a GRS Star configuration, using the coupling facility to hold and exchange GRS data across all systems in the sysplex.

MVS_GRS_RespTime_Warn

MVS_GRS_RespTime_Warn monitors the response time in the global resource serialization complex (GRS) and raises a Warning alert when response time is high or one of the systems in the ring is inactive.

The formula is:

```
(VALUE GRS_Ring.Response_Time GT 15 AND
VALUE GRS_Ring.Response_Time LE 25) OR
VALUE GRS_Ring.Status EQ Inactive
```

If this situation is raised, use operator commands to display GRS status and attempt to restart the GRS ring. If no hardware problem is indicated, there is probably a bottleneck in one of the connected systems, causing the GRS token to become stalled. If the sysplex is a parallel sysplex, this kind of problem can be eliminated by changing to a GRS Star configuration, using the coupling facility to hold and exchange GRS data across all systems in the sysplex.

MVS_XCFGroupMembers_Status_Crit

MVS_XCFGroupMembers_Status_Crit monitors the status of the members of cross-coupling facility groups and raises a Critical alert when XCF connection to one or more systems has been lost.

The formula is:

```
VALUE XCF_Member.Status EQ Missing
```

If this situation is raised, check that the systems are up, then use operator commands to try to restart the connections.

MVS_XCFGroupMembers_Status_Warn

MVS_XCFGroupMembers_Status_Warn monitors the status of the members of cross-coupling facility groups and raises a Warning alert when an XCF connection to one or more systems has failed.

If this situation is raised, use operator commands to try to restart the connections.

The formula is:

```
VALUE XCF_Member.Status EQ Failed
```

MVS_XCFSystemPaths_Crit

MVS_XCFSystemPaths_Crit monitors the number of times a cross system coupling facility path was busy when a request was made and raises a Critical alert when that count exceeds 95% of the defined limit.

The formula is:

```
VALUE XCF_Path.Retry_Percent GT 95
```

If this situation is raised, an XCF path is experiencing a very high number of retries and failure may occur soon. If another path is available, it should be brought online.

MVS_XCFSystemPaths_Warn

MVS_XCFSystemPaths_Warn monitors the number of times a cross system coupling facility path was busy when a request was made and raises a warning when the retry percent is between 80% and 95%, or the path status is "Failed", "Rebuilding", or "Quiescing".

The formula is:

```
(VALUE XCF_Path.Retry_Percent GT 80 AND
VALUE XCF_Path.Retry_Percent LE 95) OR
VALUE XCF_Path.Status GT 0X07
```

If this situation is raised, an XCF path is experiencing a high number of retries and failure may occur soon. If another path is available, it should be brought online.

OS_CMD_CF_Systems_ Perform_Crit

OS_CMD_CF_Systems_Perform_Crit monitors the performance of a coupling facility system and issues a Critical alert when a coupling facility is experiencing a very high activity rate or has failed. The situation will send a message to designated TSO user IDs when this situation is true.

The formula is:

```
VALUE CF_Systems.CPU_Percent GT 95 OR
VALUE CF_Systems.Status EQ Failed OR
VALUE CF_Systems.I/Os_Per_Second GT 5000
```

If this situation is raised, use the details view to determine which problem is occurring. If the status is "failed", switch to the secondary coupling facility and restart and rebuild the failed one. If request rate or CPU usage is high, bring the situation to the attention of a system programmer. Reconfiguration or load balancing may be necessary.

OS_CMD_DASD_Device_ContIdx_Warn

OS_CMD_DASD_Device_ContIdx_Warn monitors the contention index of a DASD device and issues a Warning alert when the contention index reaches the specified threshold (0.5 by default). It will send a message to designated TSO user IDs indicating this situation is true.

The formula is:

```
VALUE Sysplex_DASD.Average_Device_Contention_Index GT  .500
```

If this situation is raised, give this information to a system programmer or DASD specialist. If the device is on a cached storage subsystem with RAID, this may not be much of a problem. If that is the case, the threshold should be adjusted to an appropriate value for the subsystem.

OS_CMD_WLM_Performance_Idx_Crit

OS_CMD_WLM_Performance_Idx_Crit monitors the workload manager performance index of a service class period and issues a Critical alert when the performance index exceeds the specified threshold (1.50, by default). It will send a message to designated TSO user IDs indicating this situation is true.

The formula is:

```
VALUE Sysplex_WLM_Service_Class_Period.Performance_Index  GT 1
```

This situation is raised when a service class is failing to meet its goal. If this situation persists, the person responsible for the Workload Manager service definition should determine if the missed goal needs to be changed. Bottleneck analysis of the service class should provide useful information.

Sysplex_DASD_Dev_ContIndx_Warn

Sysplex_DASD_Dev_ContIndx_Warn monitors contention for a DASD device and issues a Warning alert when the index is greater than the specified threshold (.5 by default).

The formula is:

```
VALUE Sysplex_DASD.Average_Device_Contention_Index GT  .500
```

If this situation is raised, notify a system programmer or DASD specialist. If the device is on a cached storage subsystem with RAID, this may not be much of a problem. If that is the case, the threshold should be adjusted to an appropriate value for the subsystem.

Sysplex_DASDSys_VaryStatus_Warn

Sysplex_ DASDSys_ VaryStatus_ Warn monitors DASD devices and issues a Warning alert if a device vary status is ”Boxed“ or ”Not Ready.“

The formula is:

```
VALUE Sysplex_DASD_Device.Vary_Status *EQ 0X04 OR
VALUE Sysplex_DASD_Device.Vary_Status *EQ 0X05
```

This situation is raised when a DASD volume is in a ”Boxed“ or ”Not Ready“ status and is not available. If needed data sets are on the device and no backup devices are available, a system or subsystem failure may occur. If you cannot use the VARY command to return the device to online status, there may be a hardware malfunction and a Severity 1 or 2 hardware incident should be opened with the hardware vendor.

Sysplex_GlobalEnq_Wait_Crit

Sysplex_GlobalEnq_Wait_Crit monitors the wait time on a global enqueue and issues a Critical alert when either the wait time or maximum wait time exceeds 60 seconds.

The formula is:

```
VALUE Global_Enqueues.Maximum_Wait_Time GT 60 OR
VALUE Global_Enqueues.Wait_Time GT 60
```

This situation is raised when enqueue delay is high and has remained so for over a minute. Check the details to determine who is holding the ENQ. If it is a batch job that can be canceled and requeued, the deadlock can be broken by doing so. If it is a started task or on-line user, a system programmer should be called.

Sysplex_GlobalEnq_Wait_Warn

Sysplex_GlobalEnq_Wait_Warn monitors the wait time on a global enqueue and issues a Warning alert when either the wait time or maximum wait time falls between 30 and 60 seconds.

The formula is:

```
(VALUE Global_Enqueues.Maximum_Wait_Time GT 30 AND
VALUE Global_Enqueues.Maximum_Wait_Time LE 60) OR
(VALUE Global_Enqueues.Wait_Time GT 30 AND
VALUE Global_Enqueues.Wait_Time LE 60)
```

This situation is raised when ENQueue delay is high. Check the details to determine who is holding the ENQ. If it is a batch job that can be canceled and requeued, the deadlock can be broken by doing so. If it is a started task or on-line user, a system programmer should be called.

Sysplex_Workloads_PerfIdx_Crit

Sysplex_Workloads_PerfIdx_Crit monitors the performance index related to workload on the sysplex and issues a Critical alert when the performance index of any service class period exceeds the specified threshold (1.50 by default).

The formula is:

```
VALUE Sysplex_WLM_Service_Class_Period.Performance_Index  GT 1.50
```

This situation is raised when a service class is failing to meet its goal. If this situation persists, the person responsible for the Workload Manager service definition should determine if the missed goal needs to be changed. Bottleneck analysis of the service class should provide useful information.

Sysplex_Workloads_PerfIdx_Warn

Sysplex_Workloads_PerfIdx_Warn monitors the performance index related to workload on the sysplex and issues a Warning alert when the performance index of any service class period falls between 1. 20 and 1. 50.

The formula is:

```
VALUE Sysplex_WLM_Service_Class_Period.Performance_Index  GT 1.20 AND
VALUE Sysplex_WLM_Service_Class_Period.Performance_Index LE 1.50
```

This situation is raised when a service class is failing to meet its goal. If this situation persists, the person responsible for the Workload Manager service definition should determine if the missed goal needs to be changed. Bottleneck analysis of the service class should provide useful information.

Sysplex_XCFGroups_Warn

Sysplex_XCFGroups_Warn monitors cross-system coupling facility (XCF) groups and issues a Warning alert when the problem count (member status of "Failed" or "Missing") is greater than 0.

The formula is:

```
VALUE XCF_Group.Problem_Count GT 0
```

This situation is raised when there is a problem in XCF. A system programmer should be contacted to research and correct the problem.

Sysplex_XCFSystems_Status_Crit

Sysplex_XCFSystems_Status_Crit monitors the status of a cross-system coupling facility (XCF) and issues a Critical alert when cross-system communication within the sysplex is disabled.

The formula is:

```
VALUE XCF_System.Status EQ Missing
```

A system programmer should be contacted to determine why XCF is down and to restore it as soon as possible.

Sysplex_XCFSystems_Status_Warn

Sysplex_XCFSystems_Status_Warn monitors the status of a cross-system coupling facility (XCF) and issues a Warning alert when a system's status shows it is being shut down.

The formula is:

```
VALUE XCF_System.Status EQ BeingRemoved
```

When this situation is raised, a sysplex communication failure can be expected imminently. System programmer help is required to determine why the shutdown is occurring and to restart the facility as soon as possible.

System situations

The system-level predefined situations are described below in alphabetical order. You can access the description of a specific situation by selecting its name in the Contents tab under situations in the OMEGAMON XE on z/OS section of the help.

Renamed situations

The names of several pre-existing situations were changed in version 4.1.

Several of the OMEGAMON XE for UNIX System Services predefined situations were renamed when the product was merged with OMEGAMON XE on z/OS. The old and new names are shown in the following table. If you are migrating from OMEGAMON XE for UNIX System Services, existing situations are migrated when support for OMEGAMON XE for z/OS is installed into the hub monitoring server. You may continue to use those situations, in which case you should not start the OMEGAMON XE on z/OS versions. Alternatively, you can delete the original situations and use the OMEGAMON XE on z/OS versions.

OMEGAMON XE for UNIX System Services	OMEGAMON XE on z/OS
Check_Missing_Mount_Point	Check_Missing_UNIX_Mount_Point
Excess_Kernel_CPU_Time	Excess_UNIX_Kernel_CPU_Time
Excess_Process_UNIX_Run_Time	Excess_Process_UNIX_Run_Time
Excess_UNIX_System_Time	Excess_UNIX_System_Time
Excess_UNIX_User_Time	Excess_UNIX_User_Time
ENQ_Contention_Critical	UNIX_ENQ_Contention_Critical
ENQ_Contention_Warning	UNIX_ENQ_Contention_Warning
File_System_Free_Space_Critical	UNIX_File_System_FreeSpace_Crit
File_System_FreeSpace_Warning	UNIX_File_System_Free_Space_Warn
Logged_On_User_Idle	UNIX_Logged_On_User_Idle
Missing_inetd_Process	Missing_UNIX_inetd_Process
Quiecsed_File_System	Quiecsed_UNIX_File_System
Shortage_of_Processes_Critical	Shortage_of_UNIX_Processes_Crit
Shortage_of_Processes_Warning	Shortage_of_UNIX_Processes_Warn
Unwanted_inetd_Process	Unwanted_UNIX_inetd_Process

Four situations that monitor real storage have been renamed for version 4.1 to reflect that fact that z/OS no longer distinguishes between central and expanded storage. If you are running with a mixed environment during a staged upgrade, the older situations will continue to work with V3.1 agents, but you must use the new versions with V4.1 agents. You cannot distribute the new situations to V3.1 agents. When you have completed your upgrade, you should replace and delete the older versions. (They will run, but they will never become true.)

Version 3.1 and before	Version 4.1

OS390_CentralAvailFrames_Crit	OS390_Available_Frames_Crit
OS390_CentralAvailFrames_Warn	OS390_Available_Frames_Warn
OS390_CentralOnlineFrames_Crit	OS390_Frames_Online_Crit
OS390_CentralOnlineFrames_Warn	OS390_Frames_Online_Warn

OS390_Unref_Interval_Cnt_Crit and OS390_Unref_Interval_Cnt_Warn situations are still shipped, but they have changed slightly to allow for the new table data format.

Note: Situations which monitor expanded storage are no longer shipped with V4.1.

Because of architectural changes, the names of the situations listed in the following table have been shortened. You should replace the old situations with the renamed versions and delete the old ones. If you continue to use the older situations, when the situations evaluate to true, you will see correct data in the Initial Situation values column of the Situation Event Console, but not in the Current Situation column.

Old name	New named
OS390_System_PageFault_Rate_Crit	OS390_System_PageFaultRate_Crit
OS390_System_PageFault_Rate_Warn	OS390_System_PageFaultRate_Warn
OS390_Channel_LPAR_Busy_Pct_Crit	OS390_Channel_LPAR_BusyPct_Crit
OS390_Channel_LPAR_Busy_Pct_Warn	OS390_Channel_LPAR_BusyPct_Warn
OS390_Cache_FastWrite_HitPt_Crit	OS390_Cache_FastWriteHitPt_Crit
OS390_Cache_FastWrite_HitPt_Warn	OS390_Cache_FastWriteHitPt_Warn
OS390_Common_PageDS_PctFull_Crit	OS390_Common_PageDSPctFull_Crit
OS390_Common_PageDS_PctFull_Warn	OS390_Common_PageDSPctFull_Warn
OS390_Tape_Permanent_Errors_Crit	OS390_Tape_Permanent_Error_Crit
OS390_Tape_Permanent_Errors_Warn	OS390_Tape_Permanent_Error_Warn

Check_Missing_UNIX_Mount_Point

Check_Missing_UNIX_Mount_Point raises a Critical alert when a specified mount point is missing. The formula is:

```
IF *MISSING USS_Mounted_File_Systems.Mount_Point *EQ ( / )
```

If this alert is raised, a system programmer should be notified.

Crypto_CKDS_80PCT_Full

Crypto_CKDS_80PCT_Full monitors the Cryptographic Key Dataset (CKDS) and issues a Critical alert when it reaches 80% or more of its maximum capacity.

The formula is:

```
VALUE ICSF.Status EQ Active AND VALUE ICSF.CKDS_80Full EQ Yes
```

The CKDS is a VSAM linear dataset used to store keys encryption and authorization keys. If the dataset is at 80% or more of its maximum capacity, a new dataset should be created using a new master key and all keys contained in the dataset should be re-enciphered into the new dataset. The name of the current CKDS is shown in the CKDSname attribute. Refer to the ICSF Administration Guide for further details.

Crypto_CKDS_Access_Disabled

Crypto_CKDS_Access_Disabled monitors the Cyptographic Key Dataset (CKDS) and raises a Warning alert if access has been disabled.

The formula is:

```
VALUE ICSF.CKDSAccess EQ Disabled
```

The CKDS is a VSAM linear dataset used to store keys encryption and authorization keys. Access is normally disabled when a new master key or CKDS is being initialized. This interruption is temporary and access is enabled after key management operations are completed.

Crypto_Internal_Error

Crypto_Internal_Error monitors for internal errors and issues a Critical alert if one is detected.

The formula is:

```
VALUE ICSF.MonStatus NE Enabled
```

Contact IBM Support with the event attributes to report the error and for assistance in correcting the problem. MonStatus = Overrun indicates that an internal queue overflow has been detected. SCEDisabled > 0 indicates one or more service call exits have ABENDed and is no longer collecting performance data.

Crypto_Invalid_Master_Key

Crypto_Invalid_Master_Key monitors for the existence of a valid master key and raises a Critical alert if none is detected.

The formula is:

```
VALUE ICSF.Status EQ Active AND VALUE ICSF.CCMKeyOK EQ  No
```

A valid master key must be loaded into at least one of the cryptographic coprocessors. Use the ICSF ISPF dialog, TKE, or the system element to load the master key into each cryptographic coprocessor. A different master key may be loaded into coprocessors shared by PRSM Logical Partitions. Each LPAR is associated with a separate Domain Index to isolate cryptographic keys. For PCIcoprocessors, the master key must be the same value as the symmetric-keys master key (SYM-MK).

Crypto_Invalid_PKA_Master_Keys

Crypto_Invalid_PKA_Master_Keys monitors for the existence of a valid Key Management Master Key (KMMK) and a valid Signature Master Key (SMK) and raises a Critical alert if either is invalid or missing.

The formula is:

```
VALUE ICSF.PKAMKeys EQ Invalid
```

If this situation is raised, ensure that the KMMK and SMK are loaded into each coprocessor. For PCI coprocessors, the SMK key must be the same value used for the asymmetric-keys master key (ASYM-MK). Use the KMMK and SM attributes to validate the values of the verification hash patterns for these keys.

Crypto_No_Coprocessors

Crypto_No_Coprocessors monitors for cryptographic coprocessors and raises a Critical alert if none is online.

The formula is:

```
VALUE ICSF.Status EQ Active AND VALUE ICSF.1_CC EQ No
```

At least one cryptographic coprocessor must be online for cryptographic services to become available. Verify that at least one coprocessor has been configured for the z/OS system. Use the System Element console to configure the coprocessors for use by systems.

Crypto_No_PCI_Coprocessors

Crypto_No_PCI_Coprocessors monitors for PCI cryptographic coprocessors and raises a Warning alert if none is detected.

The formula is:

```
VALUE ICSF.1_PCI EQ No
```

Several Public Key Algorithm (PKA) service calls will not function without a PCI coprocessor available. Since PCI coprocessors are optimized for operations, PKA services will run slower on CMOS coprocessors.

Crypto_PCI_Unavailable

Crypto_PCI_Unavailable monitors for PCI coprocessors and raises a Critical alert if one is detected but is not online or active.

The formula is:

```
VALUE ICSF.1_PCI EQ Yes AND VALUE ICSF.PCIStatus NE Active
```

PCI coprocessors are optimized for Public Key Algorithm (PKA) operations and will run slower on CMOS coprocessors. Also, several PKA services will not run without a PCI coprocessor available.

Crypto_PKA_Services_Disabled

Crypto_PKA_Services_Disabled monitors PKA services calls and raises a Warning alert if the service calls are disabled.

The formula is:

```
VALUE ICSF.Status EQ Active AND VALUE ICSF.PKACall EQ  Disabled
```

Disable the services only to update PKA Key Management Master Key (KMMK) or Signature Master Key (SMK), or to manage the Public Key Dataset (PKDS). Enable PKA services calls only after PKA management operations are completed.

Crypto_PKDS_Read_Disabled

Crypto_PKDS_Read_Disabled monitors the status of the Public Key Dataset (PKDS) and issues a Warning alert if read operations have been disabled.

The formula is:

```
VALUE ICSF.PKDSRead EQ Disabled
```

The PKDS is a VSAM dataset used to store Public Key Algorithm (PKA) keys used for encryption and authentication. Read operations may be temporarily disabled for management operations on the PKDS. Read access to the PKDS is restored following completion of management operations. The PKDSname attribute displays the name of the current PKDS.

Crypto_PKDS_Write_Disabled

Crypto_PKDS_Write_Disabled monitors the status of the Public Key Dataset (PKDS) and issues a Warning alert if write operations have been disabled.

The formula is:

```
VALUE ICSF.PKDSWrite EQ Disabled
```

The PKDS is a VSAM dataset used to store Public Key Algorithm keys used for encryption and authentication. Write access to the dataset may be temporarily disabled to allow key management operations to occur. Write access should be enabled following completion of PKDS key management operations. The PKDSname attribute displays the name of the current PKDS.

Crypto_Service_Unavailable

Crypto_Service_Unavailable monitors the status of cryptographic services and raises a Critical alert if they are unavailable.

The formula is:

```
VALUE ICSF.CryptoSvcs EQ Inactive
```

If this situation is raised, verify that the ICSF subsystem is running on this system. If the ICSF subsystem is active, ensure that cryptographic coprocessors are online and available to this system. Also verify that a valid master key has been loaded in each coprocessor configured for this system.

Excess_Process_UNIX_Run_Time

Excess_Process_UNIX_Run_Time detects when a process exceeds 50% UNIX run time. The formula is:

```
IF *VALUE USS_Processes.UNIX_Run_Time% *GT 50
```

CPU utilization attributable to UNIX work for the indicate process exceeds the threshold defined as excessive. This condition might not be a matter of immediate concern. It may indicate that the process is in a loop requesting UNIX System Services.

Excess_UNIX_Kernel_CPU_Time

Excess_UNIX_Kernel_CPU_Time detects when the UNIX kernel is using more than 50% CPU. The formula is:

```
IF *VALUE USS_Kernel.CPU% *GT 50
```

UNIX System Services kernel CPU utilization exceeds the threshold defined as excessive. This condition might not be a matter of immediate concern. It may indicate that a process or address space is in a loop requesting kernel services. Look for processes or adress spaces with abnormally high CPU utilization.

Excess_UNIX_System_Time

Excess_UNIX_System_Time detects when a dubbed address space exceeds 50% UNIX run time. The formula is:

```
IF *VALUE USS_Address_Spaces.UNIX_System_Time% *GT 50
```

CPU utilization attributable to execution of z/OS UNIX System Services kernel code exceeds the threshold currently defined as excessive. This condition might not be a matter of immediate concern. It may indicate that the address space is in a loop requesting UNIX System Services.

Excess_UNIX_User_Time

Excess_UNIX_User_Time detects when a dubbed address space exceeds 50% UNIX user time. The formula is:

```
IF *VALUE USS_Address_Spaces.UNIX_User_Time% *GT 50
```

Address space CPU utilization attributable to UNIX work exceeds the threshold currently defined as excessive. This condition might not be a matter of immediate concern. It may indicate that the address space is in a loop.

INET_Max_Sockets__Critical

INET_Max_Sockets_Critical detects when the percentage of internet sockets in use has reached 95%. The formula is:

```
*VALUE USS_Kernel.ISock_Curr_Pct *GE 95.0
```

This situation indicates when internet socket usage is near maximum. If the maximum is reached, UNIX System Services are not accessible by Internet connections. Help from a system programmer is needed immediately. Consider increasing the value of NETWORK DOMAINNAME(AF_INET)- MA XSOCKETS(). Use the SETOMVS RESET command to dynamically change the MAXSOCKETS value. To make a permanent change, edit the BPXPRM*xx* member in SYS1.PARMLIB.

INET_Max_Sockets_Warning

INET_Max_Sockets_Warning detects when the percentage of internet sockets in use is between 80 and 95%. The formula is:

```
*IF *VALUE USS_Kernel.ISock_Curr_Pct *GE 80.0 *AND
*VALUE USS_Kernel.ISock_Curr_Pct *LT 95.0
```

This situation indicates when internet socket usage is approaching the maximum. If the maximum is reached, UNIX System Services are not accessible by Internet connections. Consider increasing the value of NETWORK DOMAINNAME(AF_INET)- MA XSOCKETS(). Use the SETOMVS RESET command to dynamically change the MAXSOCKETS value. To make a permanent change, edit the BPXPRM*xx* member in SYS1.PARMLIB.

INET6_Max_Sockets_Critical

INET6_Max_Sockets_Critical detects when the percentage of Internet v6 sockets in use has reached 95%. The formula is:

```
*IF *VALUE USS_Kernel.I6Sock_Curr_Pct *GE 95.0
```

This situation indicates when Internet V6 socket usage is near the maximum. If the maximum is reached, UNIX System Services will not be accessible by Internet connections. Help from a system programmer is needed immediately. Consider increasing the value of NETWORK DOMAINNAME(AF_INET6)- MA XSOCKETS(). Use the SETOMVS RESET command to dynamically change the MAXSOCKETS value or to make a permanent change, edit the BPXPRM*xx* member in SYS1.PARMLIB.

INET6_Max_Sockets_Warning

INET6_Max_Sockets_Warning detects when the percentage of Internet V6 sockets in use is between 80 and 95%. The formula is:

```
*IF *VALUE USS_Kernel.I6Sock_Curr_Pct *GE 80.0 *AND
*VALUE USS_Kernel.I6Sock_Curr_Pct *LT 95.0
```

This situation indicates when Internet V6 socket usage is approaching the maximum. If the maximum is reached, UNIX System Services will not be accessible by Internet connections. Consider increasing the value of NETWORK DOMAINNAME(AF_INET6)- MA XSOCKETS(). Use the SETOMVS RESET command to dynamically change the MAXSOCKETS value or to make a permanent change, edit the BPXPRM*xx* member in SYS1.PARMLIB.

KM5_CPU_Loop_Warn

This situation detects when the value of all CPU, zIIP, zIIP on CP, zAAP, and zAAP on CP using and waiting counts divided by total sample count is greater than 95.0%.

The formula is:

```
*IF *VALUE Address_Space_Bottlenecks.CPU_Loop_Index > 95.0
```

This situation indicates that address space CPU usage is high. A high value indicates either that the address space is in a CPU loop or that it is in a very CPU-intensive phase of processing. Examine each address space. If the application is known to be a heavy CPU user, then continue monitoring it. It may not actually be in a loop. If the address space is using unexpectedly large amounts of CPU, it is a candidate for cancellation. Note that very CPU intensive jobs may read high without being in a loop, so the CPU Loop Index value is a guide, not a guarantee.

Missing_UNIX_inetd_Process

Missing_UNIX_inetd_Process detects a missing inetd process. The formula is:

```
IF *MISSING USS_Processes.Command_Name *EQ ( inetd )
```

The inetd daemon is not active. The inetd daemone provides UNIX networking services. Start inetd daemon.

OS390_Allocated_CSA_Crit

OS390_Allocated_CSA_Crit monitors to determine whether the percentage of the Common Storage Area allocated is equal to or greater than 95% and issues a Critical alert if the condition is true.

The formula is:

```
IF VALUE Common_Storage.Area EQ CSA AND
VALUE Common_Storage.Allocation_Percent GE 95
```

A system crash can occur because of exhausted CSA. Use this situation to identify the address spaces using high amounts of CSA and stop or cancel nonessential address spaces with high usage.

OS390_Allocated_CSA_Warn

OS390_ Allocated_CSA_Warn monitors to determine whether the percentage of the Common Storage Area (CSA) allocated is between 90% and 94.9% inclusive and issues a Warning if the condition is true.

The formula is:

```
IF VALUE Common_Storage.Area EQ CSA AND
VALUE Common_Storage.Allocation_Percent GE 90 AND
VALUE Common_Storage.Allocation_Percent LT 95
```

A system crash can occur due to exhausted CSA. Identify the address spaces using high amounts of CSA and stop or cancel nonessential address spaces with high usage.

OS390_AvgCPU_Pct_Crit

OS390_AvgCPU_Pct_Crit monitors to determine the average percentage of time that all processors available in this z/OS system were busy dispatching work and issues a Critical alert if the average percent value is equal to or greater than 100.

The formula is:

```
IF VALUE System_CPU_Utilization.Average_CPU_Percent GE  100
```

This condition might not be a matter of immediate concern. If it arises suddenly on a uniprocessor, it may indicate that a unit of work is in a loop. If it is a chronic condition, it may be that the system is kept busy with low priority work. However, if service classes are missing their goals, a capacity increase may be needed.

OS390_AvgCPU_Pct_Warn

OS390_AvgCPU_Pct_Warn monitors to determine the average percent of time that all processors available in this system were busy dispatching work, and issues a Warning if the average percent value is between 95 and 99% inclusive.

The formula is:

```
IF VALUE System_CPU_Utilization.Average_CPU_Percent GE  95 AND
VALUE System_CPU_Utilization.Average_CPU_Percent LT 100
```

This condition might not be a matter of immediate concern. If it arises suddenly on a uniprocessor, it may indicate that a unit of work is in a loop. If it is a chronic condition, it may be that the system is kept busy with low priority work. However, if service classes are missing their goals, a capacity increase may be needed.

OS390_Cache_FastWriteHitPt_Crit

OS390_Cache_FastWrite_HitPt_Crit monitors the percentage of successful I/O requests to write data to the cache and issues a Critical alert if the percentage is between 0 and 50%. If there is no service class that is missing its goal, this situation’s thresholds may need to be adjusted.

The formula is:

```
IF VALUE DASD_MVS_DEVICES.Fast_Write_Hit_Percent GT 0  AND
VALUE DASD_MVS_Devices.Fast_Write_Hit_Percent LE 50
```

If service class periods are missing goals due to delay from the indicated device, datasets may need to be moved so that the Fast Write Cache capacity is better matched to the workload.

OS390_Cache_FastWriteHitPt_Warn

OS390_Cache_FastWrite_HitPt_Warn monitors the percentage of successful I/O requests to write data to the cache and issues a Warning alert if the percentage is between 50% and 70% inclusive.

The formula is:

```
IF VALUE DASD_MVS_DEVICES.Fast_Write_Hit_Percent LE 70  AND
VALUE DASD_MVS_DEVICES.Fast_Write_Hit_Percent GT 50
```

If there is no service class that is missing its goal, this situation’s thresholds may need to be adjusted. If service class periods are missing goals due to delay from the indicated device, datasets may need to be moved so that the Fast Write Cache capacity is better matched to the workload.

OS390_Cache_Read_HitPct_Crit

OS390_Cache_Read_HitPct_Crit monitors the percent of successful I/O requests to read data from the cache and issues a Critical alert if the percentage is greater than 0 and less than or equal to 50%.

The formula is:

```
IF VALUE DASD_MVS_DEVICES.Cache_Read_Hit_Percent GT 0  AND
VALUE DASD_MVS_DEVICES.Cache_Read_Hit_Percent LE 50
```

If this situation raises, determine whether dataset placement should be adjusted (I/O tuning). If goals are being missed, tuning may be required. If no goals are being missed, the threshold may need to be adjusted.

OS390_Cache_Read_HitPct_Warn

OS390_Cache_Read_HitPct_Warn monitors the percent of successful I/O requests to read data from the cache and issues a Warning if the percentage is between 51% and 70% inclusive.

The formula is:

```
IF VALUE DASD_MVS_Devices.Cache_Read_Hit_Percent LE 70  AND
VALUE DASD_MVS_Devices.Cache_Read_Hit_Percent GT 50
```

If this situation raises, determine whether dataset placement should be adjusted (I/O tuning). If goals are being missed, tuning may be required. If no goals are being missed, the threshold may need to be adjusted.

OS390_Cache_Write_HitPct_Crit

OS390_Cache_Write_HitPct_Crit monitors the percent of successful I/O requests to write temporary data to the cache and issues a Critical alert if the percentage is greater than 0 and less than or equal to 50%.

The formula is:

```
IF VALUE DASD_MVS_Devices.Cache_Write_Hit_Percent GT 0  AND
DASD_MVS_Devices.Cache_Write_Hit_Percent LE 50
```

If this situation is raised, determine whether any address spaces that have this device number allocated are in service classes that are missing their goals. If goals are being missed, dataset placement (I/O tuning) may be required.

OS390_Cache_Write_HitPct_Warn

OS390_Cache_Write_HitPct_Warn monitors the percent of successful I/O requests to write temporary data to the cache and issues a Warning alert if the percentage is between 51% and 70% inclusive.

The formula is:

```
IF VALUE DASD_MVS_Devices.Cache_Write_Hit_Percent LE 70  AND
VALUE DASD_MVS_Devices.Cache_Write_Hit_Percent GT 50
```

If this situation is raised, determine whether any address spaces that have this device number allocated are in service classes that are missing their goals. If goals are being missed, dataset placement (I/O tuning) may be required.

OS390_Available_Frames_Crit

OS390_Available_Frames_Crit monitors to determine when available frames of real storage are less than the specified threshold and issues a Critical alert when the condition is true. This problem should correct itself in a short time by means of page stealing. However, if the problem occurs more often than once a day, there may be a performance problem in the paging subsystem or an address space is using an excessive number of pages.

The formula is:

```
IF VALUE Real_Storage.Storage_Type EQ Summary AND
VALUE Real_Storage.Available_Frames LE 0
```

The available frame queue is a list of frames that are available to the system. If the number of frames is too low, the system resources manager (SRM) automatically replenishes the queue by stealing frames that have not been recently referenced. Controlling the available frame queue is the method SRM uses to manage central storage use. A low available frame queue is a problem only if it causes contention for central storage.

OS390_Available_Frames_Warn

OS390_Available_Frames_Warn monitors to determine when available frames of real storage are less than the specified threshold and issues a Warning alert when the condition is true. This problem should correct itself in a short time by means of page stealing. However, if the problem occurs more often than once a day, there may be a performance problem in the paging subsystem or an address space is using an excessive number of pages.

The formula is:

```
IF VALUE Real_Storage.Storage_Type EQ Summary AND
VALUE Real_Storage.Available_Frames LT 1 AND
VALUE Real_Storage.Available_Frames GT 0
```

The available frame queue is a list of frames that are available to the system. If the number of frames is too low, the system resources manager (SRM) automatically replenishes the queue by stealing frames that have not been recently referenced. Controlling the available frame queue is the method SRM uses to manage central storage use. A low available frame queue is a problem only if it causes a high page fault rate.

OS390_Frames_Online_Crit

OS390_CentralOnlineFrames_Crit monitors the central storage online frame count and issues a Critical alert when the condition is true. This situation indicates that central (real) storage available to this system is less than the threshold. If this alert results from a deliberate reconfiguration action, you should reset this situation’s threshold using the Situation editor. Otherwise, check for a possible hardware problem.

The formula is:

```
IF VALUE Real_Storage.Storage_Type EQ Summary AND
VALUE Real_Storage.Online_Frames LE 0
```

Set the critical threshold for central storage that should be online. The value must be less than the warning threshold. To be alerted of any loss of storage, you might want to set the threshold to the amount of central storage that should always be online.

OS390_Frames_Online_Warn

OS390_CentralOnlineFrames_Warn monitors the central storage online frame count and issues a Warning alert when the condition is true. This situation indicates that central (real) storage available to this system is less than the threshold. If this alert results from a deliberate reconfiguration action, you should reset this situation's threshold using the Situation editor. Otherwise, check for a possible hardware problem.

The formula is:

```
IF VALUE Real_Storage.Storage_Type EQ Summary AND
VALUE Real_Storage.Online_Frames LT 1 AND
VALUE Real_Storage.Online_Frames GT 0
```

Set the warning threshold (in frames) for online central storage. The value must be greater than the critical threshold. You might want to leave this situation off, and use the critical situation to alert you to any storage loss.

OS390_ChannelComplexBusy_Crit

OS390_ChannelComplexBusy_Crit monitors channel path activity and has determined that one or more channel paths is busier on all systems than the current Critical threshold. A Critical alert is issued. Check to determine the particular channels that are unusually active and if the threshold provided in this situation is low, adjust it using the Situation editor. Note that acceptable busy levels for tape channels, ESCON channels, and FICON® channels are typically much higher than for parallel DASD channels.

The formula is:

```
IF VALUE Channel_Paths.Complex_Percent GE 100
```

OS390_ChannelComplexBusy_Warn

OS390_ChannelComplexBusy_Warn monitors channel path activity and has determined that one or more channel paths is busier on all systems than the current Warning threshold. A Warning is issued. Check to determine the particular channels that are unusually active and if the threshold provided in this situation is low, adjust it using the Situation editor. Note that acceptable busy levels for tape channels, ESCON channels, and FICON channels are typically much higher than for parallel DASD channels.

The formula is:

```
IF VALUE Channel_Paths.Complex_Percent GE 100
```

OS390_Channel_LPAR_BusyPct_Crit

OS390_Channel_LPAR_Busy_Pct_Crit monitors the activity of the channel paths and issues a Critical alert if one or more channel paths is busier than the threshold. Identify the particular channel or channels that are unusually active. Note that the acceptable busy levels for tape channels, ESCON channels, and FICON channels are much higher than typical levels for parallel DASD channels. If the threshold for this situation is too low, adjust it using the Situation editor.

The formula is:

```
IF VALUE Channel.Paths.LPAR_Percent GE 100
```

OS390_Channel_LPAR_BusyPct_Warn

OS390_Channel_LPAR_Busy_Pct_Warn monitors the activity of the channel paths and issues a Warning alert if one or more channel paths is busier than the threshold. Identify the particular channel or channels that are unusually active. Note tht the acceptable busy levels for tape channels, ESCON channels, and FICON channels are much higher than typical levels for parallel DASD channels. If the threshold for this situation is too low adjust, it using the Situation editor.

The formula is:

```
IF VALUE Channel.Paths.LPAR_Percent GE 100
```

OS390_Channel_Path_Offline_Crit

OS390_Channel_Path_Offline_Crit monitors to determine whether a channel path is offline and issues a critical alert when this condition is true. This may be a normal condition if the indicated channel path is dynamically managed.

The formula is:

```
IF VALUE Channel_Paths.Online EQ N
```

If this situation is raised, check the configuration matrix for the current image and determine wther this channel path should be online. If so, attempt to VARY it online.

OS390_Channel_Path_Offline_Warn

OS390_Channel_Path_Offline_Warn monitors to determine whether a channel is offline and issues a warning when this condition is true. This may be a normal condition if the indicated channel path is dynamically managed. Check the configuration matrix for the current image and determine whether this channel path should be online. If so, attempt to VARY it online.

The formula is:

```
IF VALUE Channel_Paths.Online EQ N
```

If this situation is raised, check the configuration matrix for the current image and determine wther this channel path should be online. If so, attempt to VARY it online.

OS390_Common_PageDSPctFull_Crit

OS390_Common_PageDS_PctFull_Crit monitors to determine whether the percentage of slots in use on the common page dataset is greater than or equal to 80% and issues a Critical alert if the condition is true. If the common page data set becomes full, a system crash is imminent.

The formula is:

```
IF VALUE Page_Dataset_Activity.Dataset_Type EQ Common  AND
VALUE Page_Dataset_Activity.Percent_Full GE 80
```

If this situation is raised, determine which address spaces are using the largest number of common slots and terminate those that can be shut down at this time. If this situation occurs more frequently than once a month, a larger common page data set should be created and activated at the next IPL.

OS390_Common_PageDSPctFull_Warn

OS390_Common_PageDS_PctFull_Warn monitors to determine whether the percentage of slots in use on the common page dataset is greater than or equal to 60% and less than 80% and issues a Warning if the condition is true. If the common page data set becomes full, a system crash is imminent.

The formula is:

```
IF VALUE Page_Dataset_Activity.Dataset_Type EQ Common  AND
VALUE Page_Dataset_Activity.Percent Full GE 60 AND
VALUE Page_Dataset_Activity.Percent Full LT 80
```

If this situation is raised, determine which address spaces are using the largest number of common slots and terminate those that can be shut down at this time. If this situation occurs more frequently than once a month, a larger common page data set should be created and activated at the next IPL.

OS390_CSA_Growth_Crit

OS390_CSA_Growth_Crit monitors to determine whether the growth in use of the Common Storage Area is greater than or equal to 50 and issues a Critical alert if the condition is true. The formula is:

```
IF VALUE Common_Storage.Area EQ CSA AND
VALUE Common_Storage.Growth GE 50
```

If this situation is raised, identify the address spaces using high amounts of CSA and showing rapid growth in its use. Stop or cancel nonessential address spaces to avert a crash.

OS390_CSA_Growth_Warn

OS390_CSA_Growth_Warn monitors to determine whether the growth in use of the Common Storage Area is between 35 and 49 inclusive and issues a Warning alert if the condition is true. The formula is:

```
IF VALUE Common_Storage.Area EQ CSA AND
VALUE Common_Storage.Growth GE 35 AND
VALUE Common_Storage.Growth LT 50
```

If this situation is raised, identify the address spaces using high amounts of CSA and showing rapid growth in its use. Stop or cancel nonessential address spaces to avert a crash.

OS390_DASD_Busy_Percent_Crit

OS390_DASD_Busy_Percent_Crit monitors DASD device utilization and issues a Critical alert when the percentage of time a device is busy is greater than or equal to 100. The formula is:

```
IF VALUE DASD_MVS_Devices.Percent_Busy GE 100
```

This condition may represent a current or pending performance problem if any service class period is missing its goal because of I/O delay for this device. This threshold is set to 100% by default, which can be set to a lower value only when troubleshooting a chronic DASD performance problem.

OS390_DASD_Busy_Percent_Warn

OS390_DASD_Busy_Percent_Warn monitors DASD device utilization and issues a Warning alert when the percentage of time a device is busy is greater than or equal to 100. The formula is:

```
IF VALUE DASD_MVS_Devices.Percent_Busy GE 100
```

This condition may represent a current or pending performance problem if any service class period is missing its goal because of I/O delay for this device. This threshold is set to 100% by default and should be set to a lower value only when pursuing a chronic DASD performance problem.

OS390_DASD_Dropped_Ready_Crit

OS390_DASD_Dropped_Ready_Crit monitors the count of devices in this condition and issues a Critical alert if the number is greater than or equal to 5. The formula is:

```
IF VALUE DASD_MVS.Dropped_Ready GE 5
```

Should this rare situation occur, a hardware service person should be notified.

OS390_DASD_Dropped_Ready_Warn

OS390_DASD_Dropped_Ready_Warn monitors the count of devices in this condition and issues a Warning alert if the number is greater than 0 but less than 5. The formula is:

```
IF VALUE DASD_MVS.Dropped_Ready *GT 0 AND
VALUE DASD_MVS.Dropped_ Ready LT 5")
```

Should this rare situation occur, a hardware service person should be notified.

OS390_DASD_NoDynamicReconn_Critical

OS390_DASD_NoDynamicReconn_Critical monitors the count of devices in this condition and issues a Critical alert if the number is greater than or equal to 5.

The formula is:

```
IF VALUE DASD_MVS.No_Dynamic Path_Reconnect GE 5
```

This problem should be referred to appropriate personnel to determine whether the devices should be offloaded.

OS390_DASD_NoDynamicReconn_Warn

OS390_DASD_NoDynamicReconn_Warn monitors the count of devices in this condition and issues a Warning if the number is greater than 0 and less than 5.

The formula is:

```
IF VALUE DASD_MVS.No_Dynamic_Path_Reconnect GT 0 AND
VALUE DASD_MVS.No_Dynamic_Path_Reconnect LT 5
```

This problem should be referred to appropriate personnel to determine whether the devices should be offloaded.

OS390_DASD_Not_Responding_Crit

OS390_DASD_Not_Responding_Crit monitors the count of devices in this condition and issues a Critical alert if the number is greater than or equal to 5.

The formula is:

```
IF VALUE DASD_MVS.Not_Responding GE 5
```

Should this rare situation occur, a hardware service person should be notified.

OS390_DASD_Not_Responding_Warn

OS390_DASD_Not_Responding_Warn monitors the count of devices in this condition and issues a Warning if the number is greater than 0 but less than 5. Should this rare situation occur, a hardware service person should be notified.

The formula is:

```
IF VALUE DASD_MVS.Not_Responding GT 0 AND
VALUE DASD_MVS.Not_Responding LT 5
```

OS390_DASD_Response_Time_Crit

OS390_DASD_Response_Time_Crit monitors the response time for a DASD device and issues a Critical alert when the threshold value is exceeded. This situation is distributed as disabled by default and should be activated only when attempting to solve a problem where excessive DASD response time is likely to be the cause.

The formula is:

```
IF VALUE DASD_MVS_Devices.Response GE 1000000000
```

OS390_DASD_Response_Time_Warn

OS390_DASD_Response_Time_Warn monitors the response time for a DASD device and issues a Warning when the threshold value is exceeded. This situation is distributed as disabled by default and should be activated only when attempting to solve a problem where excessive DASD response time is likely to be the cause.

The formula is:

```
If VALUE DASD_MVS_Devices.Response GE 1000000000
```

OS390_ECSA_Allocation_Pct_Crit

OS390_ECSA_Allocation_Pct_Crit monitors to determine whether the percentage of the Extended Common Storage Area allocated is greater than or equal to 95% and issues a Critical alert if the condition is true. Check the current size of ECSA (the second CSA subparameter in IEASYSxx). The value may need to be adjusted before the next IPL. Attempt to determine who is using excessive ECSA or causing it to grow rapidly.

The formula is:

```
IF VALUE Common_Storage.Area EQ ECSA AND
VALUE Common_Storage.Allocation_Percent GE 95
```

OS390_ECSA_Allocation_Pct_Warn

OS390_ECSA_Allocation_Pct_Warn monitors to determine whether the percentage of the Extended Common Storage Area allocated is between 90% and 94.9% inclusive and issues a Warning if the condition is true. Check the current size of ECSA (the second CSA subparameter in IEASYSxx). The value may need to be adjusted before the next IPL. Attempt to determine who is using excessive ECSA or causing it to grow rapidly.

The formula is:

```
IF VALUE Common_Storage.Area EQ ECSA AND
VALUE Common_Storage.Allocation_Percent GE 90 AND
VALUE Common_Storage.Allocation_Percent LT 95
```

OS390_ExpandedToCentralStor_Crit

OS390_ExpandedToCentralStor_Crit monitors the page movement rate from expanded storage to central storage and issues a Critical alert when the threshold value is exceeded. This situation is disabled by default and should be activated only when attempting to solve a problem where excessive page movement is likely to be a cause.

The formula is:

```
IF VALUE Real_Storage.Storage_Type EQ CentralStorage AND
VALUE Real_Storage.Pages_Read_From_Expanded GE 1000000000
```

OS390_GlobalEnqueueReserve_Crit

OS390_GlobalEnqueueReserve_Crit monitors to determine whether the maximum wait time or the current wait time of any enqueue is greater than 60 seconds and issues a Critical alert if the condition is true. Check to determine who is holding the ENQ. If it is a batch job that can be cancelled and requeued, you can break the deadlock by so doing. If it is a started task or online user, report the problem to the appropriate personnel.

The formula is:

```
IF VALUE Enqueues.Maximum_Wait_Time GT 60 OR
VALUE Enqueues.Wait_Time GT 60
```

OS390_GlobalEnqueueReserve_Warn

OS390_GlobalEnqueueReserve_Warn monitors to determine whether the maximum wait time or the current wait time of any enqueue is between 31 and 60 seconds inclusive and issues a Warning if the condition is true. Check to determine who is holding the ENQ. If it is a batch job that can be cancelled and requeued, you can break the deadlock by so doing. If it is a started task or online user, report the problem to the appropriate personnel.

The formula is:

```
IF (VALUE Enqueues.Maximum_Wait_Time GT 30 AND
VALUE Enqueues.Maximum_Wait_Time LE 60) OR
(VALUE Enqueues.Wait_Time GT 30 and VALUE Enqueues.Wait_Time LE 60)
```

OS390_GRS_Broken_Crit

OS390_GRS_Broken_Crit monitors to determine whether the Global Resource Serialization (GRS) complex is broken and issues a Critical alert if it is. If the GRS complex is broken, it may be necessary to attempt to restart GRS from the console. You can display the status of the channel-to-channel adaptors on each system by entering the command D GRS.

The formula is:

```
IF VALUE Operator_Alerts.GRS_Status EQ Broken
```

OS390_GRS_Broken_Warn

OS390_GRS_Broken_Warn monitors to determine whether the Global Resource Serialization (GRS) complex is broken and issues a Warning if it is. If the GRS complex is broken, it may be necessary to attempt to restart GRS from the console. You can display the status of the channel-to-channel adaptors on each system by entering the command **D GRS**.

The formula is:

```
IF VALUE Operator_Alerts.GRS_Status EQ Broken
```

OS390_GTF_Active_Crit

OS390_GTF_Active_Crit monitors to determine whether the Generalized Trace Facility is active and issues a Critical alert if the condition is true. While the Generalized Trace Facility is a useful diagnostic tool, it can cause performance degradation. Ensure that GTF is active for the minimum time required to obtain the needed data.

The formula is:

```
IF VALUE Operator_Alerts.GTF_Active EQ True
```

OS390_GTF_Active_Warn

OS390_GTF_Active_Warn monitors to determine whether the Generalized Trace Facility is active and issues a Warning if the condition is true. While the Generalized Trace Facility is a useful diagnostic tool, it can cause performance degradation. Ensure that GTF is active for the minimum time required to obtain the needed data.

The formula is:

```
IF VALUE Operator_Alerts.GTF_Active EQ True
```

OS390_HSM_RecallWait_Crit

OS390_HSM_RecallWait_Crit monitors to determine whether the wait time in seconds of the longest single HSM recall that is waiting is greater than or equal to 1200 seconds and issues a Critical alert if the condition is true. Make sure that there is no outstanding tape mount for an HSM tape. In some cases, a wait can occur when a Migration Level 1 volume is tied up by a RESERVE or other conflicting activity such as a volume backup.

The formula is:

```
IF VALUE Operator_Alerts.HSM_Recall_Wait_Time GE 1200
```

OS390_HSM_RecallWait_Warn

OS390_HSM_RecallWait_Warn monitors to determine whether the wait time in seconds of the longest single HSM recall that is waiting is between 600 and 1199 seconds inclusive and issues a Warning if the condition is true. Make sure that there is no outstanding tape mount for an HSM tape. In some cases, a wait can occur when a Migration Level 1 volume is tied up by a RESERVE or other conflicting activity such as a volume backup.

The formula is:

```
IF VALUE Operator_Alerts.HSM_Recall_Wait_Time GE 600 and
VALUE Operator_Alerts.HSM_Recall_Wait_Time LT 1200
```

OS390_Indexed_VTOC_Lost_Crit

OS390_Indexed_VTOC_Lost_Crit monitors the count of devices in this condition and issues a Critical alert if the count is greater than or equal to 5. Refer this problem to an appropriate storage management specialist.

The formula is:

```
IF VALUE DASD_MVS.Indexed_VTOC_Lost GE 5
```

OS390_Indexed_VTOC_Lost_Warn

OS390_Indexed_VTOC_Lost_Warn monitors the count of devices in this condition and issues a Warning if the number is greater than 0 but less than 5. Refer this problem to an appropriate storage management specialist.

The formula is:

```
IF VALUE DASD_MVS.Indexed_VTOC_Lost GT 0 AND
VALUE DASD_MVS.Indexed_VTOC_Lost LT 5
```

OS390_Local_PageDS_Errors_Crit

OS390_Local_PageDS_Errors_Crit monitors to determine whether the number of errors in a local page dataset is greater than or equal to 5 and issues a Critical alert if the condition is true. Identify the failing dataset or datasets. Check for a spare page dataset slot, and if there is no spare, increase the PAGETOT parameter in IEASYSxx. There should be at least one spare slot per two page datasets. Remove the failing dataset from the PAGE parameter in IEASYSxx. If there is a spare slot, PAGEADD a dataset and use PAGEDEL REPLACE to move the pages to a good dataset.

The formula is:

```
IF VALUE Page_Dataset_Activity.Errors GE 5
```

OS390_Local_PageDS_Errors_Warn

OS390_Local_PageDS_Errors_Warn monitors to determine whether the number of errors in a local page dataset is greater than or equal to 1 and less than 5 and issues a Warning if the condition is true. Identify the failing dataset or datasets. Check for a spare page dataset slot, and if there is no spare, increase the PAGETOT parameter in IEASYSxx. There should be at least one spare slot per two page datasets. Remove the failing dataset from the PAGE parameter in IEASYSxx. If there is a spare slot, PAGEADD a dataset and use PAGEDEL REPLACE to move the pages to a good dataset.

The formula is:

```
IF VALUE Page_Dataset_Activity.Errors GE 1 and
VALUE Page_Dataset_Activity.Errors LT 5
```

OS390_Local_PageDS_PctFull_Crit

OS390_Local_PageDS_PctFull_Crit monitors to determine whether the percentage of slots in use on a local page dataset is greater than or equal to 35% and issues a Critical alert if the condition is true. When usage approaches 30%, paging efficiency begins to decline, and blocked paging disappears at about 35% occupancy. If this situation occurs, prepare to PAGEADD another dataset if the critical threshold is passed. If the current PAGTOTL setting in IEASYSxx does not allow another dataset to be added, it should be increased before the next IPL.

The formula is:

```
IF VALUE Page_Dataset_Activity.Dataset_Type EQ Local and
VALUE Page_Dataset_Activity.Percent_Full GE 35
```

OS390_Local_PageDS_PctFull_Warn

OS390_Local_PageDS_PctFull_Warn monitors to determine whether a local page dataset is greater than or equal to 25% full and less than 35% full and issues a Warning if the condition is true. When usage approaches 30%, paging efficiency begins to decline, and blocked paging disappears at about 35% occupancy. If this situation occurs, prepare to PAGEADD another dataset if the critical threshold is passed. If the current PAGTOTL setting in IEASYSxx does not allow another dataset to be added, it should be increased before the next IPL.

The formula is:

```
IF VALUE Page_Dataset_Activity.Dataset_Type EQ Local and
VALUE Page_Dataset_Activity.Percent Full GE 25 and
VALUE Page_Dataset_Activity.Percent Full LT 35
```

OS390_LPAR_OverheadPercent_Crit

OS390_LPAR_OverheadPercent_Crit monitors to determine whether the percentage of time the system spends managing a logical partition is greater than or equal to 20% and issues a Critical alert if the condition is true. Possible causes include saturation of the CPU capacity leading to excessive overhead switching CPUs between LPARs. This can be compounded when there are too many logical processors assigned to an LPAR.

The formula is:

```
IF VALUE System_CPU_Utilization.Partition_Overhead% GE  20
```

OS390_LPAR_OverheadPercent_Warn

OS390_LPAR_OverheadPercent_Warn monitors to determine whether the percentage of time the system spends managing a logical partition is greater than or equal to 10% and less than 20% and issues a Warning if the condition is true. Possible causes include saturation of the CPU capacity leading to excessive overhead switching CPUs between LPARs. This can be compounded when there are too many logical processors assigned to an LPAR.

The formula is:

```
IF VALUE System_CPU_Utilization.Partition_Overhead% GE  10 and
VALUE System_CPU_Utilization.Partition_Overhead% LT 20
```

OS390_LPAR_STATUS_Crit

OS390_LPAR_STATUS_Crit monitors to determine whether LPAR CPU Management Overhead or Velocity Index have exceeded thresholds and if so, issues a Critical alert. These conditions might not be of immediate concern. In the case of LPAR CPU Management Overhead, if the number of configured LPARs is substantial, it may trigger this situation. If the conditions persist, you may consider reducing the number of configured LPARs. In the case of the Velocity Index, you may want to adjust LPAR weights if the LPARs' workloads are not meeting expected service levels.

The formula is:

```
IF (VALUE LPAR_Clusters.LPAR_Name NE _CLTotal AND
VALUE LPAR_Clusters.LPAR_Name NE _CPTotal AND
VALUE LPAR_Clusters.LPAR_Name NE PHYSICAL AND
VALUE LPAR_Clusters.LPAR_Effective_Weight_Index LT 0.9) OR
```

```
(VALUE LPAR_Clusters.LPAR_NAME NE _CLTotal AND
VALUE LPAR_Clusters.LPAR_Name NE _CPTotal AND
VALUE LPAR_Clusters.LPAR_Name NE PHYSICAL AND
VALUE LPAR_Clusters.Host_LPAR_Flag EQ Y AND
VALUE LPAR_Clusters.CPC_CPU_Overhead GT 15.0)
```

OS390_LPAR_STATUS_Warn

OS390_LPAR_STATUS_Warn monitors to determine whether LPAR CPU Management Overhead or Velocity Index have exceeded thresholds and if so, issues a Warning. These conditions might not be of immediate concern. In the case of LPAR CPU Management Overhead, if the number of configured LPARs is substantial, it may trigger this situation. If the conditions persist, you may consider reducing the number of configured LPARs. In the case of the Velocity Index, you may want to adjust LPAR weights if the LPARs' workloads are not meeting expected service levels.

The formula is:

```
IF (VALUE LPAR_Clusters.LPAR_Name NE _CLTotal AND
VALUE LPAR_Clusters.LPAR_Name NE _CPTotal AND
VALUE LPAR_Clusters.LPAR_Name NE PHYSICAL AND
VALUE LPAR_Clusters.LPAR_Effective_Weight_Index LT 1.0) OR
(VALUE LPAR_Clusters.LPAR_NAME NE _CLTotal AND
VALUE LPAR_Clusters.LPAR_Name NE _CPTotal AND
VALUE LPAR_Clusters.LPAR_Name NE PHYSICAL AND
VALUE LPAR_Clusters.Host_LPAR_Flag EQ Y AND
VALUE LPAR_Clusters.CPC_CPU_Overhead GT 10.0)
```

OS390_MAX_ASIDs_in_Use_Crit

OS390_MAX_ASIDs_in_Use_Crit monitors to determine whether the percentage that represents the maximum number of address space vector table slots that are in use or unavailable is greater than or equal to 90% and issues a Critical alert if the condition is true. Check the values of the MAXUSER, RSVNONR, and RSVSTART parameters as well as for any problems that could lead to address space IDs becoming unusable.

The formula is:

```
IF VALUE Operator_Alerts.ASVT_Slot_Utilization GE 90
```

OS390_MAX_ASIDs_in_Use_Warn

OS390_MAX_ASIDs_in_Use_Warn monitors to determine whether the percentage that represents the maximum number of address space vector table slots that are in use or unavailable is between 80 and 89% inclusive, and issues a Warning if the condition is true. Check the values of the MAXUSER, RSVNONR, and RSVSTART parameters as well as for any problems that could lead to address space IDs becoming unusable.

The formula is:

```
IF VALUE Operator_Alerts.ASVT_Slot_Utilization GE 80 and
VALUE Operator_Alerts.ASVT_Slot_Utilization LT 90
```

OS390_Network_ResponseTime_Crit

OS390_Network_ResponseTime_Crit monitors the Network Response Time and when it equals or exceeds 10, issues a Critical alert. Appropriate personnel should be notified if the condition persists.

The formula is:

```
IF VALUE User_Response_Time.Network_Response GE 10
```

OS390_Network_ResponseTime_Warn

OS390_Network_ResponseTime_Warn monitors the Network Response Time and when it equals or exceeds 5 but is less than 10, issues a Warning. Appropriate personnel should be notified if the condition persists.

The formula is:

```
IF VALUE User_Response_Time.Network_Response GE 5 AND
VALUE User_Response_Time.Network_Response LT 10
```

OS390_OLTEP_Active_Crit

OS390_OLTEP_Active_Crit monitors to determine whether OLTEP is active and issues a Critical alert if the situation is true. Determine who is using OLTEP and minimize the time of its use.

The formula is:

```
IF VALUE Operator_Alerts.OLTEP_Active EQ True
```

OS390_OLTEP_Active_Warn

OS390_OLTEP_Active_Warn monitors to determine whether OLTEP is active and issues a Warning alert if it is. Determine who is using OLTEP and minimize the time of its use.

The formula is:

```
IF VALUE Operator_Alerts.OLTEP_Active EQ True
```

OS390_Outstanding_WTORs_Crit

OS390_Outstanding_WTORs_Crit monitors to determine whether the number of outstanding Write to Operator with Reply requests is greater than or equal to 12 and issues a Critical alert if the condition is true. Check the operator console for outstanding replies and address these. If all of the outstanding replies are correct and routine, you may want to adjust this situation's threshold.

The formula is:

```
IF VALUE Operator_Alerts.Outstanding_Operator Replies  GE 12
```

OS390_Outstanding_WTORs_Warn

OS390_Outstanding_WTORs_Warn monitors to determine whether the number of outstanding Write to Operator with Reply requests is between 10 or 11 inclusive and issues a Warning if the condition is true. Check the operator console for outstanding replies and address these. If all of the outstanding replies are correct and routine, you may want to adjust this situation's threshold.

The formula is:

```
IF VALUE Operator_Alerts.Outstanding_Operator_Replies  GE 10 AND
VALUE Operator_Alerts.Outstanding_Operator_Replies LT 12
```

OS390_PageDSNotOperational_Crit

OS390_PageDSNotOperational_Crit monitors the number of page datasets in this condition and issues a Critical alert if the number is greater than or equal to 5. Verify that paging devices are operational. If a device is not operational, attempt to VARY it online. If a page data set was drained by a prior PAGEDEL

DRAIN command, it may now be removed by a PAGEDEL DELETE command. If this alert occurs without warning, an IPL may be imminent. Prepare to shut down and request appropriate assistance.

The formula is:

```
IF VALUE System_Paging_Activity.Datasets_Not_Operational  GE 5
```

OS390_PageDSNotOperational_Warn

OS390_PageDSNotOperational_Warn monitors the number of page datasets in this condition and issues a Warning if the number is from 1 to 4 inclusive. Verify that paging devices are operational. If a device is not operational, attempt to VARY it online. If a page data set was drained by a prior PAGEDEL DRAIN command, it may now be removed by a PAGEDEL DELETE command. If this alert occurs without warning, an IPL may be imminent. Prepare to shut down and request appropriate assistance.

The formula is:

```
IF VALUE System_Paging_Activity.Datasets_Not_Operational  GT 0 AND
VALUE System_Paging_Activity.Datasets_Not_Operational LT 5
```

OS390_Page_Rate_Crit

This situation monitors the current paging rate and raises a Critical alert when the rate exceeds the threshold. The formula is:

```
IF VALUE Page_Dataset_Activity.Page_Rate *LT 0
```

Excessive paging may increase application wait and response time. Because system page rate is dependent on processor type, real storage configuration, and workload you may need to adjust your paging system based on your installation defined service requirements. This situation is disabled by default and should be activated only when necessary to characterize a chronic excessive paging problem.

OS390_Page_Rate_Warn

This situation monitors the current paging rate and raises a Warning alert when the rate exceeds the threshold. The formula is:

```
IF VALUE Page_Dataset_Activity.Page_Rate *LT 0
```

Because system page rate is dependent on processor type, real storage configuration, and workload you may need to adjust your paging system based on your installation defined service requirements. This situation is disabled by default and should be activated only to diagnose a chronic high paging rate.

OS390_Physical_CPUs_Online_Crit

OS390_Physical_CPUs_Online_Crit monitors the number of online CPUs and issues a Critical alert when the number is less than the current threshold. This situation is disabled (set to 0) by default and should be activated only to diagnose chronic configuration problems.

The formula is:

```
IF VALUE System_CPU_Utilization.Physical_CPU_Count LT  0
```

OS390_Physical_CPUs_Online_Warn

OS390_Physical_CPUs_Online_Warn monitors the number of online CPUs and issues a warning when the number is less than the current threshold. This situation is disabled (set to 0) by default and should be activated only to diagnose chronic configuration problems.

The formula is:

```
IF VALUE System_CPU_Utilization.Physical_CPU_Count LT  0
```

OS390_RMF_Not_Active_Crit

OS390_RMF_Not_Active_Crit monitors to determine whether the RMF monitor is inactive and issues a Critical alert if the condition is true. RMF data is essential to performance management and problem analysis. If you cannot restart the RMF, notify appropriate personnel.

The formula is:

```
IF VALUE Operator_Alerts.RMF_Not_Active EQ True
```

OS390_RMF_Not_Active_Warn

OS390_RMF_Not_Active_Warn monitors to determine whether the RMF monitor is inactive and issues a Warning if the condition is true. RMF data is essential to performance management and problem analysis. If you cannot restart the RMF, notify appropriate personnel.

The formula is:

```
IF VALUE Operator_Alerts.RMF_Not_Active EQ True
```

OS390_SMF_Not_Recording_Crit

OS390_SMF_Not_Recording_Crit monitors to determine whether the SMF is recording information and issues a Critical alert if the condition is true. SMF data has numerous uses including resource accounting and capacity management. Check the SMF datasets and restart the collection process as soon as possible. If you cannot restart the SMF datasets, notify appropriate personnel.

The formula is:

```
IF VALUE Operator_Alerts.SMF_Not_Recording EQ True
```

OS390_SMF_Not_Recording_Warn

OS390_SMF_Not_Recording_Warn monitors to determine whether the SMF is recording information and issues a Warning if the condition is true. SMF data has numerous uses including resource accounting and capacity management. Check the SMF datasets and restart the collection process as soon as possible. If you cannot restart the SMF datasets, notify appropriate personnel.

The formula is:

```
IF VALUE Operator_Alerts.SMF_Not_Recording EQ True
```

OS390_SYSLOG_Not_Recording_Crit

OS390_SYSLOG_Not_Recording_Crit monitors to determine whether the System Log is recording information and issues a Critical alert if the condition is true. Determine why logging has stopped. A possibility is that JES spool space is exhausted.

The formula is:

```
IF VALUE Operator_Alerts.SYSLOG_Not_Recording EQ True
```

OS390_SYSLOG_Not_Recording_Warn

OS390_SYSLOG_Not_Recording_Warn monitors to determine whether the System Log is recording information and issues a Warning if the condition is true. Determine why logging has stopped. A possibility is that JES spool space is exhausted.

The formula is:

```
IF VALUE Operator_Alerts.SYSLOG_Not_Recording EQ True
```

OS390_System_Page_Rate_Crit

This situation monitors the system page rate and raises a Critical alert when the threshold is reached. The formula is:

```
IF VALUE System_Paging_Activity.System_Page_Rate *LT 0
```

Excessive paging may increase application wait and response time. Because system page rate is dependent on processor type, real storage configuration, and workload you may need to adjust your paging system based on your installation defined service requirements. This situation is disabled by default and should be activated only to diagnose a chronic high paging rate.

OS390_System_Page_Rate_Warn

This situation monitors the system page rate and raises a Warning alert when the threshold is reached. The formula is:

```
IF VALUE System_Paging_Activity.System_Page_Rate *LT 0
```

Excessive paging may increase application wait and response time. Because system page rate is dependent on processor type, real storage configuration, and workload you may need to adjust your paging system based on your installation defined service requirements. This situation is disabled by default and should be activated only to diagnose a chronic high paging rate.

OS390_System_PageFault_Rate_Crit

OS390_System_PageFault_Rate_Crit monitors the system page fault rate and issues a Critical alert when the threshold is exceeded. This situation is shipped disabled by default.

The formula is:

```
IF VALUE System_Paging_Activity.Page_Fault_Rage GE 1000000000
```

OS390_System_PageFault_Rate_Warn

OS390_System_PageFault_Rate_Warn monitors the system page fault rate and issues a Warning when the threshold is exceeded. This situation is shipped disabled by default.

The formula is:

```
IF VALUE System_Paging_Activity.Page_Fault_Rage GE 1000000000
```

OS390_Tape_Dropped_Ready_Crit

OS390_Tape_Dropped_Ready_Crit monitors the number of tape drives in this condition and issues a Critical alert if the threshold is exceeded. Check the devices and attempt to make them ready. If this is not possible, report the condition to the appropriate personnel.

The formula is:

```
IF VALUE Tape_Drives.Dropped_Ready GE 5
```

OS390_Tape_Dropped_Ready_Warn

OS390_Tape_Dropped_Ready_Warn monitors the number of tape drives in this condition and issues a Warning if the threshold is exceeded. Check the devices and attempt to make them ready. If this is not possible, report the condition to the appropriate personnel.

The formula is:

```
IF VALUE Tape_Drives.Dropped_Ready GT 0 AND
VALUE Tape_Drives.Dropped_Ready LT 5
```

OS390_Tape_Mount_Pend_Time_Crit

This situation issues a Critical alert when any tape unit has been waiting for a tape mount for more than 1200 seconds (20 minutes) or more. The situation has a monitoring interval of 5 minutes. The situation message displays the volume or volumes being requested to allow tape operators to know which tape volumes require mounts.

The formula is:

```
IF VALUE Tape_Drives.Tape_Mount_Pending_Time *GE 1200
```

A tape mount pending time that exceeds the threshold might require contacting operations to ensure that the MOUNT request has been recognized by personnel responsible for mounting tapes on the requested tape unit.

OS390_Tape_Mount_Pend_Time_Warn

This situation issues a Warning alert when any tape unit has been waiting for a tape mount for more than 600 seconds (10 minutes) and less than 1200 seconds (20 minutes). The situation has a monitoring interval of 2 minutes. The warning situation message displays the volume or volumes being requested to allow tape operators to know which tape volumes require mounts.

The formula is:

```
IF VALUE Tape_Drives.Tape_Mount_Pending_Time *GT 600 *AND
*VALUE Tape_Drives.Tape_Mount_Pending_Time *LT 1200
```

A tape mount pending time that exceeds the threshold might require contacting operations to ensure that the MOUNT request has been recognized by personnel responsible for mounting tapes on the requested tape unit.

OS390_Tape_Not_Responding_Crit

OS390_Tape_Not_Responding_Crit monitors the number of tape drives in this condition and issues a Critical alert if the threshold is exceeded. If the condition is persistent and the devices cannot be activated by VARYing them online, report the problem to the appropriate personnel.

The formula is:

```
IF VALUE Tape_Drives.Not_Responding GE 5
```

OS390_Tape_Not_Responding_Warn

OS390_Tape_Not_Responding_Warn monitors the number of tape drives in this condition and issues a Warning if the threshold is exceeded. If the condition is persistent and the devices cannot be activated by VARYing them online, report the problem to the appropriate personnel.

The formula is:

```
IF VALUE Tape_Drives.Not_Responding GT 0 AND
VALUE Tape_Drives.Not_Responding LT 5
```

OS390_Tape_Permanent_Error_Crit

OS390_Tape_Permanent_Errors_Crit monitors the count of permanent errors on a tape drive and issues a Critical alert if the number is greater than or equal to 30.

The formula is:

```
IF VALUE Tape_Drives.Permanent_Errors GE 30
```

OS390_Tape_Permanent_Error_Warn

OS390_Tape_Permanent_Errors_Warn monitors the count of permanent errors on a tape drive and issues a Warning if the number is between 5 and 29 inclusive.

The formula is:

```
IF VALUE Tape_Drives.Permanent_Errors GE 5 AND
VALUE Tape_Drives.Permanent_Errors LT 30
```

OS390_Tape_Temp_Errors_Crit

OS390_Tape_Temp_Errors_Crit monitors the count of temporary errors on a tape drive and issues a Critical alert if the number is greater than or equal to 30. The problem could be caused either by the media or by the device. Monitor to determine whether there is additional degradation and if so, report the problem to appropriate personnel.

The formula is:

```
IF VALUE Tape_Drives.Temporary_Errors GE 30
```

OS390_Tape_Temp_Errors_Warn

OS390_Tape_Temp_Errors_Warn monitors the count of temporary errors on a tape drive and issues a Warning if the number is between 5 and 29 inclusive. The problem could be caused either by the media or by the device. Monitor to determine whether there is additional degradation and if so, report the problem to appropriate personnel.

The formula is:

```
IF VALUE Tape_Drives.Temporary_Errors GE 5 and
VALUE Tape_Drives.Temporary_Errors LT 30
```

OS390_Undispatched_Tasks_Crit

OS390_Undispatched_Tasks_Crit monitors to determine whether the number of tasks or address spaces that have not been dispatched by the SRM due to constraints is greater than or equal to 20 and issues a Critical alert if the condition is true. If the condition persists for more than an hour, a capacity upgrade may be required. Determine whether any important service classes are missing their goals.

The formula is:

```
IF VALUE System_CPU_Utilization.Undispatched_Tasks GE  20
```

OS390_Undispatched_Tasks_Warn

OS390_Undispatched_Tasks_Warn monitors to determine whether the number of tasks or address spaces that have not been dispatched by the SRM due to constraints is greater than or equal to 5 and less than 20 and issues a Warning if the condition is true. If the condition persists for more than an hour, a capacity upgrade may be required. Determine whether any important service classes are missing their goals.

The formula is:

```
IF VALUE System_CPU_Utilization.Undispatched_Tasks GE  05 AND
VALUE System_CPU_Utilization.Undispatched_Tasks LT 20
```

OS390_Unowned_Common_Stor_Crit

OS390_Unowned_Common_Stor_Crit monitors the amount of unowned storage in the Common Services Area and issues a Critical alert if the threshold is exceeded. Ensure that the CSA Analyzer collector is running.

The formula is:

```
IF VALUE Common_Storage.Area EQ CSA AND
VALUE Common_Storage.Unowned GE 1000000000
```

OS390_Unowned_Common_Stor_Warn

OS390_Unowned_Common_Stor_Warn monitors the amount of unowned storage in the Common Services Area and issues a Warning alert if the threshold is exceeded. Ensure that the CSA Analyzer collector is running.

The formula is:

```
IF VALUE Common_Storage.Area EQ CSA AND
VALUE Common_Storage.Unowned GE 1000000000
```

OS390_Unref_Interval_Cnt_Crit

OS390_Unref_Interval_Cnt_Crit monitors to determine whether the storage type is Summary and the amount of time, in seconds, that the oldest frame of pageable storage has gone without being referenced is less than or equal to 10 seconds, and issues a Critical alert if the condition is true.

The formula is:

```
IF VALUE Real_Storage.Storage_Type EQ Summary AND
VALUE Real_Storage.Unreferenced_Interval_Count LE 10
```

If this situation is raised, determine whether any important service classes are failing to meet their goals and if Private Page-in Wait is a significant reason. If so, central storage may be overcommitted, possibly the result of a capacity problem.

Because lower values indicate resource contention, the unreferenced interval count (UIC) is an alert where the critical threshold must be less than the warning threshold.

OS390_Unref_Interval_Cnt_Warn

OS390_Unref_Interval_Cnt_Warn monitors to determine whether the storage type is Summary and the amount of time, in seconds, that the oldest frame of pageable storage has gone without being referenced is greater than 10 and less than 20 seconds, and issues a Warning if the condition is true. Determine whether any important service classes are failing to meet their goals and, if Private Page-in Wait is a significant reason. If so, central storage may be overcommitted, possibly the result of a capacity problem.

The formula is:

```
IF VALUE Real_Storage.Storage_Type EQ Summary AND
VALUE Real_Storage.Unreferenced_Interval_Count LT 20 AND
VALUE Real_Storage.Unreferenced_Interval_Count GT 10
```

Because lower values indicate resource contention, the unreferenced interval count (UIC) is an alert where the warning threshold must be greater than the critical.

OS390_User_Host_Resp_Time_Crit

OS390_User_Host_Resp_Time_Crit monitors to determine whether the host (internal) response time for the indicated TSO user is exceeding the Critical threshold and if so, issues a Critical alert. If the user's service class is meeting its goal, there may be a specific problem in this user's address space. If the service class is missing its goal, the goal may be too demanding and may need to be adjusted.

The formula is:

```
IF VALUE User_Response_Time.Host_Response GE 100000000
```

OS390_User_Host_Resp_Time_Warn

OS390_User_Host_Resp_Time_Warn monitors to determine whether the host (internal) response time for the indicated TSO user is exceeding the Warning threshold and if so, issues a Warning alert. If the user's service class is meeting its goal, there may be a specific problem in this user's address space. If the service class is missing its goal, the goal may be too demanding and may need to be adjusted.

The formula is:

```
IF VALUE User_Response_Time.Host_Response GE 100000000
```

OS390_User_Total_Resp_Time_Crit

OS390_User_Total_Resp_Time_Crit monitors the total response time (host plus network) for a TSO user and issues a Critical alert if the threshold is exceeded. If the service class for this user is meeting its goal, the problem may be a network response time problem. The formula is:

```
IF VALUE User_Response_Time.Total_Response GE 100000000.0
```

OS390_User_Total_Resp_Time_Warn

OS390_User_Total_Resp_Time_Warn monitors the total response time (host plus network) for a TSO user and issues a Warning alert if the threshold is exceeded. If the service class for this user is meeting its goal, the problem may be a network response time problem.

The formula is:

```
IF VALUE User_Response_Time.Total_Response GE 100000000.0
```

OS390_WTO_Buffers_Left_Crit

OS390_WTO_Buffers_Left_Crit monitors to determine whether the remaining WTO buffer pool is becoming dangerously small and issues a Critical alert if the condition is true. Determine whether a console device is down and if so, switch the message stream to another device.

The formula is:

```
IF VALUE Operator_Alerts.WTO_Buffers_Remaining LE 20
```

OS390_WTO_Buffers_Left_Warn

OS390_WTO_Buffers_Left_Warn monitors to determine whether the remaining WTO buffer pool is becoming short of resources and issues a Warning if the condition is true. Determine whether a console device is down and if so, switch the message stream to another device.

The formula is:

```
IF VALUE Operator_Alerts.WTO_Buffers_Remaining GT 20 AND
VALUE Operator_Alerts.WTO_Buffers_Remaining LE 100
```

Quiesced_UNIX_File_System

Quiesced_UNIX_File_System detects a quiesced file system. The formula is:

```
IF *VALUE USS_Mounted_File_Systems.Status *EQ Quiesced
```

The file system indicated is in a Quiesced state. This condition might not be a matter of immediate concern. For example, this could be due to an HSM backup recovery in progress. If this condition persists, a system programmer should be notified.

Shortage_of_UNIX_Processes_Crit

Shortage_of_UNIX_Processes_Crit checks to determine if the current number of processes is very close to 90. The formula is:

```
*VALUE USS_Kernel.Used_Processes *GE 90
```

If the maximum is reached, no more processes can be started. If this condition is due to expected growth, increase the maximum value. Otherwise, a system programmer should be notified.

Shortage_of_UNIX_Processes_Warn

Shortage_of_UNIX_Processes_Warn checks to determine if the current number of processes is between 80 and 90. The formula is:

```
*VALUE USS_Kernel.Used_Processes *GE 80 *AND *VALUE USS_Kernel.Used_Processes *LT 90
```

If the maximum number of processes is reached, no more processes can be started. If this condition is due to expected growth, increase the maximum value. Otherwise, a system programmer should be notified.

UNIX_ENQ_Contention_Critical

UNIX_ENQ_Contention_Critical detects when an HFS enqueue contention has lasted 30 seconds or more. The formula is:

```
*VALUE USS_HFS_ENQ_Contention.Time *GE 30
```

If this condition occurs, check the details to determine who is holding the enqueue. If it is a batch job that can be canceled and requeued, the deadlock can be broken by doing that. If it is a started task, a UNIX process, or an online user, a system programmer should be notified.

UNIX_ENQ_Contention_Warning

ENQ_Contention_Warning detects when an HFS enqueue contention has lasted 10 seconds or more. The formula is:

```
*VALUE USS_HFS_ENQ_Contention.Time *GE 10 *AND
*VALUE USS_HFS_ENQ_Contention.Time *LT 30
```

If this condition occurs, check the details to determine who is holding the enqueue. If it is a batch job that can be canceled and requeued, the deadlock can be broken by doing that. If it is a started task, a UNIX process, or an online user, a system programmer should be notified.

UNIX_File_System_FreeSpace_Crit

UNIX_File_System_Free_Space_Critical detects when any file system has less than 10% free space. The formula is:

```
IF *VALUE USS_Mounted_File_Systems.Percent_Used *GE 90
```

If this condition occurs, the file system space should be extended. This can be accomplished with UNIX System Services commands. If this condition becomes chronic, a system programmer should be notified.

UNIX_File_System_FreeSpace_Warn

UNIX_File_System_Free_Space_Warning detects when any file system has less than 20% free space. The formula is:

```
IF *VALUE USS_Mounted_File_Systems.Percent_Used *GE 80 *AND
*VALUE USS_Mounted_File_Systems.Percent_Used *LT 90
```

If this condition occurs, consider extending the file system space. This can be accomplished with UNIX System Services commands. If this condition becomes chronic, a system programmer should be notified.

UNIX_Logged_On_User_Idle

Logged_On_User_Idle detects a logged-on user with excessive idle time. The formula is:

```
IF *VALUE USS_Logged_on_Users.Idle_Time_Mins *GT 480
```

This condition might not be a matter of immediate concern. Consult installation procedures for appropriate action.

UNIX_Max_Sockets_Critical

UNIX_Max_Sockets_Critical detects when the percentage of UNIX sockets in use has reached 95%. The formula is:

```
IF *VALUE USS_Kernel.USock_Curr_Pct *GE 95.0
```

This situation indicates that usage of UNIX sockets is near the maximum. If the maximum is reached, UNIX System Services functionality will be adversely affected. Help from a system programmer is needed immediately. Consider increasing the value of NETWORK DOMAINNAME(AF_UNIX)- MAXSOCKETS(). Use the SETOMVS RESET command to dynamically change the MAXSOCKETS value, or, to make a permanent change, edit the BPXPRM*xx* member in SYS1.PARMLIB.

UNIX_Max_Sockets_Warning

UNIX_Max_Sockets_Warning indicates that the percentage of UNIX sockets in use is between 80 and 95%. The formula is:

```
*IF *VALUE USS_Kernel.USock_Curr_Pct *GE 80.0 *AND
*VALUE USS_Kernel.US ock_Curr_Pct *LT 95.0
```

This situation indicates that usage of UNIX sockets is approaching the maximum. If the maximum is reached, UNIX System Services functionality will be adversely affected. Consider increasing the value of NETWORK DOMAINNAME(AF_UNIX)- MAXSOCKETS(). Use the SETOMVS RESET command to dynamically change the MAXSOCKETS value, or, to make a permanent change, edit the BPXPRM*xx* member in SYS1.PARMLIB.

Unwanted_inetd_UNIX_Process

Unwanted_inetd_UNIX_Process detects a missing or unwanted inetd process. The formula is:

```
IF *VALUE USS_Processes.Command_Name *EQ inetd
```

The inetd daemon provides UNIX networking services. If the active process is unwanted, consult the installation procedures for appropriate action. If the inetd daemon should be, stop or do not start this situation.

Appendix A. Documentation library

This appendix contains information about the publications in the OMEGAMON XE on z/OS library and about other publications related to OMEGAMON XE on z/OS.

See *IBM Tivoli Monitoring and OMEGAMON XE Products: Documentation Guide*, SC23-8816, for information about accessing and using the publications. You can find the *Documentation Guide* in the IBM Tivoli Monitoring and OMEGAMON XE Information Center at http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/.

To find a list of new and changed publications, click **What's new** on the Welcome page of the IBM Tivoli Monitoring and OMEGAMON XE Information Center. To find publications for the previous version of a product, click **Previous information centers** on the Welcome page for the product.

OMEGAMON XE on z/OS library

The following documents are available for OMEGAMON XE on z/OS:

- *Program Directory* GI11-4113

  Contains information about the material and procedures associated with the installation of IBM Tivoli OMEGAMON XE on z/OS. The Program Directory is intended for the system programmer responsible for program installation and maintenance.
- *Planning and Configuration Guide* SC27-2354

  Provides information that helps plan the deployment and configuration of OMEGAMON XE on z/OS and the required common services component. It also provides detailed instructions for configurating product components. This document is intended for system administrators and others who are responsible for configuring OMEGAMON XE on z/OS.
- *User's Guide* SC27-2353

  Introduces the features, workspaces, attributes, and predefined situations for the OMEGAMON XE on z/OS product and supplements the user assistance provided with this product. This document is written for data center operators and analysts responsible for monitoring and troubleshooting system performance and availability or performing trend analysis for resource planning.
- *Troubleshooting Guide* GC27-2355

  Provides explanations for the messages issued by the OMEGAMON XE on z/OS product, its OMEGAMON II for MVS component, and common agent components. This booksalso provides troubleshooting advice for installation and configuration, security, and usage problems, and instructions for setting up tracing on z/OS.
- *OMEGAMON for MVS User's Guide* SC27-2356

  Describes the features and commands used in OMEGAMON for MVS. Reference information for OMEGAMON major and minor commands is included by functional area, along with a description of the following features: User Profile Facility, Exception Analysis, CSA Analyzer, End-to-End Response Time Feature, Bottleneck Analysis, DEXAN®, Impact Analysis, Workload Profile Facility.
- *OMEGAMON for MVS User's Guide* SC27-2356

  Contains complete descriptions of OMEGAMON for MVS commands, organized alphabetically by command name. Includes a chapter on "Command Groupings" that is an introduction organized by topic (exception analysis, hiperspace, paging, and so on) where you can refresh your memory as to the proper spelling of a command or keyword.
- *EPILOG User's Guide* SC27-2356

  Describes the basic reporting features of EPILOG for MVS. The introduction provides a product overview and a discussion of the EPILOG approach to performance management. The rest of the manual explains how to use the reporter, including the various types of reports and the use of the

DISPLAY command. Topics, such as advanced reporting options, the Workload Profile Facility, exception filtering, exporting historical data, and reporting with SAS graphics are also documented.

- *EPILOG Command Reference* .SC27-2356

  Contains complete descriptions of EPILOG for MVS commands, organized alphabetically by command name.

OMEGAMON XE and Tivoli Management Services on z/OS common library

The books in this library are common to some or all of the OMEGAMON XE products or Tivoli Management Services on z/OS:

- *Quick Start Guide*, GI11-8918

  Provides an overview of the installation and setup process for a monitoring agent on z/OS.
- *Common Planning and Configuration Guide*, SC23-9734

  Covers planning and configuration information common to the OMEGAMON XE V4.2.0 monitoring agents and to the components of Tivoli Management Services on z/OS V6.2.1.
- *Upgrade Guide*, SC23-9745

  Provides an overview and instructions for performing the upgrades from prior versions of OMEGAMON XE monitoring agents and Tivoli Management Services components.
- *End-to-End Response Time Feature Reference*, SC27-2303

  Documents the End to End Response Time feature, a common component used by four OMEGAMON XE monitoring agents on z/OS: CICS, z/OS, IMS, and Mainframe Networks.
- *Reports for Tivoli Common Reporting*, SC27-2304

  Provides information about the Tivoli Common Reporting tool that is specific to products that run under the Tivoli Enterprise Portal and use the Tivoli Data Warehouse database.

OMEGAMON II for MVS V520 library

The OMEGAMON II documentation has not been updated since V520. Any changes relevant to configuration of the version currently incorporated into OMEGAMON XE on z/OS V4.2.0 are documented in the *IBM Tivoli OMEGAMON XE on z/OS: Planning and Configuration Guide*. The existing documents refer to the Candle Support structure rather than to IBM software support. References to Candle® Support processes and procedures are invalid. Direct questions to IBM Software Support. For details about the IBM support structure, see “Support information” on page 331.

The following information sources are available in the OMEGAMON II for MVS V520 library.

- *Configuration and Customization Guide*, GC32-9277

  Describes how to configure and customize the OMEGAMON II for MVS product. It provides background on the product components, addresses maintenance and migration considerations, gives an overview of the configuration and customization process, and documents step-by-step procedures.
- *User's Guide*, GC32-9280

  Contains an overview of OMEGAMON II features, the types of panels displayed, and how to navigate from one panel to another; instructions for adjusting an OMEGAMON II environment; usage scenarios describing how to use OMEGAMON II to monitor realtime and historical performance; instructions for using some of the commands for creating OMEGAMON screen spaces described in the OMEGAMON for MVS Command Language Reference Manual and commands for generating EPILOG resports described in the EPILOG for MVS Command Language Reference Manual.
- *OMEGAMON for MVS Command Language Reference Manual*, GC32-9276

  Provides the syntax and available keywords for OMEGAMON II for MVS commands
- *EPILOG for MVS Command Language Reference Manual*, GC32-9275

Documents the syntax and available keywords for EPILOG for MVS commands.

IBM Tivoli Monitoring library

The following publications provide information about IBM Tivoli Monitoring V6.2 and about the commonly shared components of Tivoli Management Services:

- *Quick Start Guide*, GI11-8058

  Introduces the components of IBM Tivoli Monitoring.
- *Installation and Setup Guide*, GC32-9407

  Provides instructions for installing and configuring IBM Tivoli Monitoring components on Windows, Linux, and UNIX systems.
- *Program Directory for IBM Tivoli Management Services on z/OS*, GI11-4105

  Gives instructions for the SMP/E installation of the Tivoli Management Services components on z/OS.
- *Configuring the Tivoli Enterprise Monitoring Server on z/OS*, SC27-2313

  Gives detailed instructions for using the Configuration Tool to configure Tivoli Enterprise Monitoring Server on z/OS systems. Includes scenarios for using batch mode to replicate monitoring environments across the z/OS enterprise. Also provides instructions for setting up security and for adding application support to a Tivoli Enterprise Monitoring Server on z/OS.
- *Administrator's Guide*, SC32-9408

  Describes the support tasks and functions required for the Tivoli Enterprise Portal Server and clients, including Tivoli Enterprise Portal user administration.
- Tivoli Enterprise Portal online help

  Provides context-sensitive reference information about all features and customization options of the Tivoli Enterprise Portal. Also gives instructions for using and administering the Tivoli Enterprise Portal.
- *User's Guide*, SC32-9409

  Complements the Tivoli Enterprise Portal online help. The guide provides hands-on lessons and detailed instructions for all Tivoli Enterprise Portal features.
- *Command Reference*, SC32-6045

  Provides detailed syntax and parameter information, as well as examples, for the commands you can use in IBM Tivoli Monitoring.
- *Troubleshooting Guide*, GC32-9458

  Provides information to help you troubleshoot problems with the software.
- *Messages*, SC23-7969

  Lists and explains messages generated by all IBM Tivoli Monitoring components and by z/OS-based Tivoli Management Services components (such as Tivoli Enterprise Monitoring Server on z/OS and TMS:Engine).

Other sources of documentation

You can also obtain technical documentation about Tivoli Monitoring and OMEGAMON XE products from the following sources:

- IBM Tivoli Open Process Automation Library (OPAL)

  http://www.ibm.com/software/tivoli/opal

  OPAL is an online catalog that contains integration documentation as well as other downloadable product extensions. This library is updated daily.
- Redbooks

  http://www.redbooks.ibm.com/

  IBM Redbooks®, Redpapers, and Redbooks Technotes provide information about products from platform and solution perspectives.

- Technotes

  You can find Technotes through the IBM Software Support Web site at http://www.ibm.com/software/support/probsub.html, or more directly through your product Web site, which contains a link to Technotes (under **Solve a problem**).

  Technotes provide the latest information about known product limitations and workarounds.

Appendix B. Acronyms used

The following acronyms are used in the Tivoli OMEGAMON XE on z/OS help:

CBU Capacity Backup. Enables alternative System z processors to take over workload from another server in case of an emergency or unusually high demand.

CF Coupling facility. On z/OS, a special logical partition that provides high-speed caching, list processing, and locking functions in a sysplex.

CFRM Coupling Facility Resource Manager. A component of z/OS that provides the services to manage coupling facility resources in a sysplex. This management includes the enforcement of CFRM policies to ensure that the coupling facility and structure requirements are satisfied.

CP Central processor. The part of the computer that contains the sequencing and processing facilities for instruction execution, initial program load, and other machine operations.

CPC Central processor complex. In a z/OS or OS/390 environment, a physical collection of hardware (such as an ES/3090™) that consists of main storage, one or more central processors, timers, and channels.

CPM Capacity provisioning management. The ability to dynamically provision physical processors to, and deprovision them from, a System z10 complex based on an installation policy, with or without operator intervention.

CPU Central processing unit. The part of a computer that includes the circuits that control the interpretion and running of instructions.

CSA Common service area. In z/OS, a part of the common area that contains data areas that can be addressed by all address spaces, but is protected during its use by the key of the requester.

CSECT Control section.

ECB Event control block. A control block used to represent the status of an event.

ECKD Extended Count Key Data. An extension of the count-key-data (CKD) architecture. Count-key data is a data-record format employing self-defining record formats in which each record is represented by up to three fields: a count field identifying the record and specifying its format, an optional key field that can be used to identify the data area contents, and an optional data field that typically contains the user data.

ECSA Extended common service area. A major element of z/OS virtual storage above the 16MB line. This area contains pageable system data areas that are addressable by all active virtual storage address spaces. It duplicates the common system area (CSA) which exists below the 16MB line.

ESCON Enterprise Systems Connection. A peripheral interface for an Enterprise Systems Architecture/390 and zSeries computer. The I/O interface uses ESA/390 logical protocols over a serial interface that configures attached units to a communication fabric.

ESQA Extended system queue area. A major element of z/OS virtual storage above the 16MB line. This storage area contains tables and queues relating to the entire system. It duplicates above the 16MB line the system queue area (SQA) .

GRS ring Global resource serialization ring. Consists of one or more systems connected to each other by communication links. The links are used to pass information about requests for global resources from one system to another in the complex. Requests are made by passing a message or token, called the ring system authority (RSA) message, between systems in a round-robin or ring fashion.

HSA Hardware system area. A logical area of central storage, not addressable by application programs, used to store Licensed Internal Code and control information.

HSM Hierarchical storage management. A function that automatically distributes and manages data on disk, tape, or both by regarding devices of these types and potentially others as levels in a storage hierarchy that range from fast, expensive devices to slower, cheaper, and possibly removable devices. The objectives are to minimize access time to data and maximize available media capacity.

ICF Integrated Coupling Facility. A processor used by multiple systems to coordinate work. ICFs must be assigned to LPARs that then become coupling facilities.

ICSF Integrated Cryptographic Service Facility. Provides the administrative interface and a large set of application interfaces to the cryptographic coprocessor hardware.

IFA Integrated Facility for Applications. A special type of processor.

IRD Intelligent Resource Director. A key feature of the zSeries architecture, which automatically directs resources to priority workloads. IRD gives users the capability of managing resource and workload across z/OS LPARs that are members of a common group called an LPAR cluster. The Intelligent Resource Director is made up of three parts: LPAR CPU Management, Dynamic Channel Path Management (DCM), and Channel Subsystem Priority Queuing (CSSPQ).

JES Job Entry Subsystem. An IBM licensed program that receives jobs into the system and processes all output data that is produced by jobs.

MSU Millions of service units. A measure of the amount of processing a computer can perform in one hour.

OOCoD On/Off Capacity On Demand. The ability of z990 and later servers to temporarily activate and deactivate processors and memory units to handle fluctuating workloads.

RMF Resource Management Facility. An IBM licensed program or optional element of z/OS that measures selected areas of system activity and presents the data collected in the format of printed reports, System Management Facility (SMF) records, or display reports. RMF is used to evaluate system performance and identify reasons for performance problems.

RSA Ring system authority. Requests between systems in a GRS ring system are sent in a round-robin or ring fashion by passing a message or token, called the ring system authority (RSA) message.

SQA System queue area. An area of virtual storage below the 16MB line reserved for system-related control blocks.

SRB Service request block. A unit of work that is scheduled to execute in another address space.

SRM System resources manager. A group of programs that controls the use of system resources, such as programs, devices, and storage areas that are assigned for use in jobs.

TCB Task control block. A z/OS control block used to communicate information about tasks, within an address space, that are connected to a z/OS subsystem.

VIO Virtual input/output. Temporary data sets can be handled by a facility called virtual input/output (VIO). VIO data sets reside in the paging space but appear to the access method to reside on a direct access storage device.

VTOC Volume table of contents. A table on a direct access volume that describes the location, size and other characteristics of each data set on the volume. Each entry is a data set control block (DSCB).

WLM Workload Manager. A z/OS component that prioritizes workloads running on z/OS and matches workloads with available resources.

XCF Cross-system coupling facility. A special logical partition that provides high-speed caching, list processing, and locking functions in a sysplex. XCF provides the z/OS coupling services that allow authorized programs on z/OS systems in a multisystem environment to communicate with (send data to and receive data from) authorized programs on other z/OS systems.

zAAP zSeries Application Assist Processor. A special class of assist processor designed to run Java workloads. For reporting purposes, a zAAP is usually referred to as an integrated facility for applications (IFA).

zIIP zSeries Integrated Information Processor. A special class of assist process used mostly for DB2 workloads.

Appendix C. Inspect messages

The following messages are issued by the Inspect agent and displayed in the Agent Messages view of the Inspect Address Space CPU Use workspace:

KM3IN001I: NO CPU ACTIVITY WAS DETECTED BY INSPECT

The Inspect agent did not see any CPU activity when it looked at the target address space. The target address space may be consuming CPU time but Inspect only considers it if it sees the address space executing during a sample since it needs to assign the CPU time to a component within the address space.

System action: None.

User Action: Refresh the workspace to run Inspect again.

KM3IN002E: ASID PASSED TO INSPECT IS ZERO

The address space ID passed to the Inspect agent by the query was zero. Either the Inspect workspace was invoked from the physical navigation tree or the address space was not assigned to the query by the link.

System action: No data is returned.

User action: Determine why the address space ID value being passed to the query is zero. Possibly the workspace linkage has been changed.

KM3IN003E: ASID PASSED TO INSPECT IS INVALID

The address space ID passed to the Inspect agent by the query was not a valid address space number on the z/OS system. The link to the Inspect workspace may be assigning the incorrect field to the ASID parameter of the query.

System action: No data is returned.

User action: Determine why the ASID value being passed to the query is invalid. Possibly the workspace linkage has been changed

KM3IN004I: ADDRESS SPACE PRIORITY GREATER THAN INSPECT

The priority of the target address space is greater than that of the Tivoli Enterprise Monitoring Server in which the Inspect agent is executing. Inspect may not obtain enough CPU cycles to collect a reasonable sample of data from the target address space.

System action: Processing continues.

User Action: None

KM3IN005I: DATA COLLECTION MAY BE INCOMPLETE

Part two of message KM3IN004I informing the user of the consequences of the priority difference.

System action: Processing continues.

User Action: None

KM3IN006E: SAMPLER TASK ABENDED WITH CODE *code*

The Inspect agent sampler program task ended with the specified code.

System action: No data is returned.

User Action: Determine the cause of the ABEND and possibly obtain a system dump of the ABEND.

KM3IN007E: SELECTED JOB NOT IN TARGET ADDRESS SPACE

The Inspect agent will check that the job name passed is executing within the target ASID. If not, either the wrong job name or ASID was passed by the link to the workspace or the job may have ended. If the job has not ended, check that the link to the Inspect workspace is passing the correct values to the Inspect query.

System action: No data is returned.

User Action: If the job has not ended, check that the link to the Inspect workspace is passing the correct values to the Inspect query.

KM3IN008I: GRANULARITY SET TO `size`

The Inspect agent attempts to limit the amount of detail returned to the client to about 100 rows of data by dynamically calculating the ″grain″ size used split each CSECT of code into ″blocks″ to which CPU activity is attributed.

Because the user cannot know in advance how diverse the execution activity map will be, nor the sizes of the load modules and CSECT that have activity, nor even which modules and CSECTs will have activity, it is not reasonable to allow the user to set this parameter before the Inspect agent is executed, therefore the Inspect agent attempts to calculate a suitable size once all the Inspect data has been collected.

Size is displayed as a hexadecimal value.

System Action: Processing continues.

User Action: None

KM3IN009E: ATTACH OF `program` FAILED WITH RC= `rc`

The Inspect agent attempted to attach the Inspect sampling program program as a separate z/OS task within the Tivoli Enterprise Monitoring Server address space but the ATTACH failed with return code `rc` .

System action: No data is returned.

User Action: Determine the meaning of the return code for the ATTACH function and correct the error.

K**M3IN010I: INSPECT WAS NOT EXECUTED**

Part two of messages KM3IN009E and KM3IN011E. Explains that no data was collected because of the error reported by the previous message.

KM3IN011E UNABLE TO LOAD `program`

The Inspect agent attempted to load the Inspect sampling program but the load operation failed.

System action: No data is returned.

User Action: Determine the cause of the failure to load pogram by scanning the Tivoli Enterprise Monitoring Server and z/OS logs for any appropriate messages.

KM3IN012E OPEN ERROR FOR DATASET `dsn`

The Inspect agent attempts to obtain CSECT information for each load module by accessing and load libraries allocated to the job step TCB of the target address space.

The Inspect agent was unable to open the specified dataset.

System action: Processing continues:

User Action: The person responsible for Tivoli Enterprise Monitoring Server operations on the target host should browse the monitoring server and z/OS logs to determine the cause of the failure.

KM3IN013I: SOME CSECT DATA MAY NOT BE AVAILABLE

Part two of messages KM3IN012E, KM3IN013E, KM3IN015E and KM3IN016E. Indicates that CSECT information may not be available for some load modules due to the previous error.

KM3IN014E: UNABLE TO ALLOCATE DATASET `dsn`

The Inspect agent attempts to obtain CSECT information for each load module by accessing and load libraries allocated to the job step TCB of the target address space.

The Inspect agent was unable to allocate the specified dataset.

System action: Processing continues.

User Action: The person responsible for Tivoli Enterprise Monitoring Server operations on the target host should browse the monitoring server and z/OS logs to determine the cause of the failure.

KM3IN015E: UNABLE TO READ FORMAT 1 DSCB FOR DATASET `dsn`

The Inspect agent attempts to obtain CSECT information for each load module by accessing and load libraries allocated to the job step TCB of the target address space.

The Inspect agent was unable to read the format 1 DSCB for the specified dataset.

System action: Processing continues.

User Action: The person responsible for Tivoli Enterprise Monitoring Server operations on the target host should browse the monitoring server and z/OS logs to determine the cause of the failure.

KM3IN016E UNABLE TO SUPPORT PDSE FOR DATASET `dsn`

The Inspect agent attempts to obtain CSECT information for each load module by accessing and load libraries allocated to the job step TCB of the target address space.

The Inspect agent determined that the specified dataset was a PDSE but the IEWBIND program could not be loaded by the Inspect agent.

System action: Processing continues.

User Action: The person responsible for Tivoli Enterprise Monitoring Server operations on the target host should browse the monitoring server and z/OS logs to determine why IEWBIND could not be located by the Inspect agent.

Support information

If you have a problem with your IBM software, you want to resolve it quickly. IBM provides the following ways for you to obtain the support you need:

Online
Go to the IBM Software Support site at http://www.ibm.com/software/support/probsub.html and follow the instructions.

IBM Support Assistant
The IBM Support Assistant (ISA) is a free local software serviceability workbench that helps you resolve questions and problems with IBM software products. The ISA provides quick access to support-related information and serviceability tools for problem determination. To install the ISA software, go to http://www.ibm.com/software/support/isa.

Troubleshooting Guide
For more information about resolving problems, see the product's Troubleshooting Guide.

Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785 U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION ″AS IS″ WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement might not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758 U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

If you are viewing this information in softcopy form, the photographs and color illustrations might not be displayed.

Trademarks

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

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

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

UNIX is a registered trademark of The Open Group in the United States and other countries.

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

Glossary

A

agent. See *monitoring agent*.

alert. A warning message that appears at a console to indicate an event has occurred that may require intervention.

alert adapter. An agent that monitors and relays alert information from another product to Tivoli Enterprise Portal. Sources of alerts include message logs, system consoles, network management products, and system management products.

attribute. A system or application element being monitored by the agent, such as Disk Name and Disk Read/Writes Per Second. An attribute can also be a field in an ODBC-compliant database.

attribute group. A set of related attributes that can be combined in a view or a situation. Each type of IBM Tivoli OMEGAMON XE agent has a set of attribute groups associated with it.

C

Common user access (CUA). A Systems Application Architecture® (SAA®) specification that gives a series of guidelines describing the way information should be displayed on a screen, and the interaction techniques between users and computers. The OMEGAMON II for MVS interface conforms to CUA guidelines.

component. A separate product or feature of a product provided by IBM

configuring. Making a product operational by completing the configuration of the product using the Configuration Tool and completing the manual steps required outside of the Configuration Tool

coupling facility. A special logical partition that provides high-speed caching, list processing, and locking functions in a parallel sysplex.

cross-system coupling facility (XCF). A component of z/OS which provides services that allow authorized applications to coordinate and manage communication, resource sharing and recovery among other instances of the application in a sysplex. XCF provides Group Services to define the application structure, Monitoring Services to determine when failures have occurred and Signalling Services to provide communication among the members. A sysplex requires a couple dataset to manage the sysplex. The couple dataset includes information related to the systems in the sysplex, XCF groups and members.

CUA. See common user access.

cumulative maintenance. Maintenance through a given date that is customer approved.

customizing. Modifying the defaults for options and settings and other changes that reflect the needs of your site.

E

enclave. A collection of the routines that make up an application. The enclave is the equivalent of any of the following:

- A run unit, in COBOL
- A program, consisting of a main C function and its sub-functions, in C and C++
- A main procedure and all of its subroutines, in PL/I
- A program and its subroutines, in Fortran

Enclaves are a key feature of the program management model.

enqplex. a group of z/OS images in two or more sysplexes under common enqueue management. A resource in one enqplex is distinct from a resource having the same name in another enqplex. Two or more sysplexes having the same enqplex name share qname/rname resources.

enqueue. A shared memory structure that serializes access to database resources. Enqueues are local to one instance if Real Application Clusters is not enabled. When you enable Real Application Clusters, enqueues can be global to a database. *See also* global enqueue. An enqueue is a locking mechanism.

event. A change in the status of a situation you are monitoring.

event workspace. The workspace that opens when you select a situation from the flyover list of situations that have fired.

F

filter. A criterion or set of criteria used to limit the amount of information returned to the user in response to a query.

function. A method of evaluating the information that an attribute supplies. Used in situations, the default function is Value of expression. The functions available depend on the type of attribute used in the condition.

G

goal. A performance target for a workload.

global enqueue. An enqueue that is shared across systems.

global resource serialization (GRS). A base z/OS function that provides applications as well as systems level functions with the ability to serialize a resource between multiple units of work. The resource is abstract (whatever one wants it to be) and can be serialized at various levels (scopes). User resource naming conversions and scope specifications ensure proper intersect between different requesters. Scopes/levels range from multiple systems to single address spaces or job steps. GRS is widely used and is critical to the stability/integrity of a z/OS operating system and all its exploiters.

GRS. *See* global resource serialization.

H

historical reporting. A viewing option for tables that collects data for a specified period of time, such as over the past hour.

hub. The Tivoli Enterprise Monitoring Server that has been designated to act as the focal point to which all Tivoli Enterprise Portal Server connect. A non-hub, or remote, Tivoli Enterprise Monitoring Server passes its collected data to the hub to be made available to clients, thereby creating an Enterprise-wide view.

I

Integrated Facility for Applications (IFA). A special type of processor. *See* zAAP z/Series Application Assist Processor.

Intelligent Resource Director (IRD). A key feature of the zSeries architecture, which automatically directs resources to priority workloads. IRD gives users the capability of managing resource and workload across z/OS LPARs that are members of a common group called an LPAR cluster.

L

LPAR cluster. A set of z/OS LPARs that share a central processing complex (CPC) and are members of a common sysplex. An LPAR cluster comprises the scope of the Intelligent Resource Director management control.

M

managed system. A particular operating system, subsystem, or application in your enterprise that being monitoring with an OMEGAMON or IBM Tivoli OMEGAMON XE agent.

managed system name. (1) From the standpoint of OMEGAMON XE on z/OS, sysplexes and systems are *managed systems*. In the Tivoli Enterprise Portal Navigator, managed systems are identified by *managed system names*. (2) Sysplex managed system names take the form *plexname*:MVS:SYSPLEX, where *plexname* is normally the true name of the sysplex, but could be configured to be an alias for the sysplex. (3) System managed system names take the form *plexname*:*smfid*:MVSSYS, where *plexname* is normally the true name of the sysplex, but could be configured to be an alias for the sysplex.

monitor interval. A specified time, scalable for seconds, minutes, hours, or days, for how often the Tivoli Enterprise Monitoring Server checks to see if a situation has become true. The minimum monitor interval is 30 seconds; the default is 15 minutes.

monitoring agent. Component that monitors systems, subsystems, or applications on the system where they are installed.

N

Navigator. The left pane of the Tivoli Enterprise Portal application window. The Navigator physical view shows your network enterprise as a physical hierarchy of systems grouped by platform. OMEGAMON DE users can also create other views to create logical hierarchies grouped as you specify, such as by department or function.

O

OMEGAMON II. Component that collects and displays data in the OMEGAMON II user interface(s). These include:

- The menu driven CUA interface that is IBM SAA/CUA compliant
- The facility that allows multiple OMEGAMON IIs to execute in the same address space and that communicates with all of them;
- The Common Interface (CI) for some OMEGAMON IIs, the command driven Classic interface z/OS

P

Parallel Sysplex (sysplex). A set of z/OS systems that communicate and cooperate with each other through multisystem hardware components and software services to process customer workloads. A

Parallel Sysplex combines parallel processing with read/write data sharing across multiple systems with full data integrity.

performance index (PI). A mean.

persistent datastore. Component that records and stores historical data.

policy. An automated system process that you set up to perform actions and automates manual tasks. It comprises a series of automated steps, called activities, whose order of execution you control.

predefined situations. A set of ready-made situations for you to use as-is or to modify without having to create your own.

predefined workspaces. A set of workspaces that come with your product for you to use and modify for your environment.

presentation files. Installed with the Server, presentation.dat and presentation.idx store the workspace definitions, link definitions, and terminal emulator scripts.

preventive maintenance. Fixes that can be applied to avoid known problems

product code. The three-letter code used to identify a monitoring product in certain contexts. For example, the product code for OMEGAMON XE on z/OS is KM5.

Properties editor. A dialog for specifying the properties of the individual views that make up a workspace.

Q

query. A query is a request for data from the agent.

R

report class. Work for which reporting information is to be collected separately. A report class can combine work from different service class or a single transaction.

Resource Measurement Facility (RMF). An IBM licensed program or optional element of z/OS that measures selected areas of system activity and presents the data collected in the format of printed reports, System Management Facility (SMF) records, or display reports. RMF is used to evaluate system performance and identify reasons for performance problems.

S

service class. Represents a grouping of work with similar resource and performance requirements. Each service class can have one or more periods. Each service class period has a goal.

service definition. An explicit definition of all the workloads and processing capacity in a sysplex. A service definition includes service policies, workloads, service classes, resource groups, and classification rules.

service policy. A set of performance goals for all z/OS images using z/OS workload management in a sysplex. There can be only one active service policy for a sysplex, and all subsystems in goal mode within that sysplex process towards that policy. However, you can create several service policies, and switch between them to cater for the different needs of different processing periods. The Workload Manager (WLM) will dynamically balance the system resources according to the active service policy.

sample. The data that the product collects for the server instance.

sample interval. The time between data samplings.

severity . The value or relative importance you assign to a particular state. If two or more states occur at the same time, the Navigator level that contains these states shows the indicator for the highest severity level.

situation. A set of conditions that, when met, creates an event. A condition consists of an agent attribute, an operator such as greater than or equal to, and a value. It can be read as, “If - system condition - compared to - value - is true”. An example of a situation is: IF - CPU usage - GT - 90% - TRUE. IF and TRUE are part of every situation. The expression “CPU usage GT 90%” is the situation condition.

SQL. Structured Query Language. SQL is a programming language for getting information from and updating a database. The Tivoli Enterprise Portal Queries editor enables you to write SQL queries to ODBC data sources for retrieval and display in table and chart views.

SRB. Service request block. A unit of work that is scheduled to execute in another address space.

state. An indication associated with an icon, color, and severity level assigned to a situation at any particular point in time. A situation can reflect one of the following states: Critical, Warning, or Informational.

status. The true or false condition of a situation.

sysplex. A set of systems communicating and cooperating with each other, through multisystem hardware components and software services, in order to process workloads. *See also* Parallel Sysplex.

sysplex proxy. A Tivoli Enterprise Monitoring Server that acts as a data consolidation point for sysplex monitoring. Historical data for the sysplex is collected at the proxy, and sysplex situations are evaluated there.

System Management Facility (SMF). A z/OS facility that collects and records a variety of system and job-related information.

System z9 Integrated Information Processor (zIIP). A special class of processor designed to handle eligible data workloads, beginning with several types of DB2 V8.x workloads.

T

Tivoli Data Warehouse. A long-term data store for the performance and analysis data collected by monitoring agents.

The warehoused data is written to Microsoft SQL Server relational database. You can view the data stored in the warehouse in Tivoli Enterprise Portal workspaces, or use third-party analysis and reporting tools on it.

Tivoli Enterprise Portal. The Java-based graphical user interface used to display and work with data provided by the monitoring products.

Tivoli Enterprise Portal Server. A collection of software services for the Tivoli Enterprise Portal that enables retrieval, manipulation and, analysis of data from the monitoring agents running on systems in your enterprise. The Tivoli Enterprise Portal Server connects to the hub Tivoli Enterprise Monitoring Server.

Tivoli Enterprise Monitoring Server. The component of the Tivoli Management Services that:

- Consolidates the data collected by the monitoring agents and distributes the data to the Tivoli Enterprise Portal.
- In some cases, receives commands from the Tivoli Enterprise Portal and distributes them to the appropriate agent or OMEGAMON XE product.
- Stores historical data and prototypes for configuration in the form of application-specific data.

Tivoli Management Services. The infrastructure shared by OMEGAMON XE, IBM Tivoli Monitoring and other products, whose components include Tivoli Enterprise Portal desktop client, Tivoli Enterprise Portal Server, Tivoli Enterprise Monitoring Server, and Tivoli Data Warehouse.

V

Velocio. A type of goal set for a service class period. Velocio is a percentage value indicating how fast work should execute when ready, without being delayed for processor or storage access. Velocity goals reflect the ratio of time when a unit of work was ready to use the CPU to the time work was actually using the CPU or I/O. This type of goal is useful for managing long-running nontransaction servers such as JES or HSM. A high percentage indicates that work should process quickly; a low percentage indicates that a greater amount of delay is acceptable.

velocity. The measure of how fast work should run when ready, without being delayed for processor or storage access. Velocity is expressed as a percentage from 1 to 99. In other words, the velocity goal defines the acceptable amount of delay for work when work is ready to run. Velocity goals should be used for long-running jobs and for address spaces; they are the most appropriate goal for any started tasks that require a goal.

view. A windowpane, or frame, in a workspace. It may contain data from an agent in a chart or table, or it may contain a terminal session or notepad, for example. A view can be split into two separate, autonomous views.

W

warehouse proxy. A process that periodically moves data from the binary history files maintained at the monitoring server or monitoring agent to the warehouse.

workload. Work to be tracked, managed and reported as a unit. A group of service classes.

workspace. A window comprised of one or more views. Every item listed in the Navigator has its own default workspace and may have multiple workspaces.

X

XCF. *See* Cross-system coupling facility

Z

z/OS. An IBM operating system for the IBM zSeries family of enterprise servers that includes and integrates functions previously provided by many IBM software products (including the MVS and OS/390 operating systems). z/OS is an open, secure operating system for the IBM zSeries family of enterprise servers, complies with industry standards, is enabled for network computing and e-business, and supports technology advances in networking server capability, parallel processing, and object-oriented programming.

zAAP. *See* zSeries Application Assist Processor.

zSeries Application Assist Processor (zAAP) . A special class of assist processor designed to run Java workloads. For reporting purposes, a zAAP is usually referred to as an integrated facility for applications (IFA).

zIIP. *see* System z9 Integrated Information Processor.

Index

Printed in USA

SC27-2353-00