5.5

IBM OMEGAMON for Storage on z/OS Tuning Guide

SC27-4380-04

Note

Before using this information and the product it supports, read the information in “Notices” on page 25.

Edition notice

2023-07-27

This edition applies to Version 5.5 of OMEGAMON for Storage on z/OS and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

Chapter 1. Tips for maximizing performance

Reduce the impact of monitoring activity, improve the performance of IBM Tivoli Monitoring, and make the collection of short and long-term historical data more efficient.

To reduce the impact of monitoring activity:

- Collect data less frequently.
- Filter the data that you collect.

To improve the performance of IBM Tivoli Monitoring:

- Reduce steady-state demand, such as historical data collection, on system resources.
- Reduce on-demand or user-generated demand, such as user requests for data, on system resources.

To make the collection of short and long-term historical data more efficient:

- Distinguish between short and long-term historical data collection.
- Determine the types of historical data to collect and reconfigure interval parameters.

Monitoring activities

Use masks and filters to target the data that you collect and readjust the data collection interval to reduce the demand of monitoring activities on system resources.

Masks

The collection, summarizing, and writing of historical data using system resources.

Use masks to reduce the amount of resources required to retrieve data and improve response times for users. For example, you can select and exclude noncritical devices from monitoring. To find out how to use masks to filter data, see *IBM OMEGAMON for Storage: User's Guide*.

Data collection intervals

You must balance the requirement to collect sufficient information about monitored resources and the impact of this activity on system resources. If you can collect sufficient information about a monitored resource when you collect data every 60 minutes, set the data collection frequency to 60 minutes, and not to a lower frequency, such as 45 minutes.

Tip: Use PARMGEN to change data collection frequency parameters.

Steady-state demand on system resources

You can tune IBM Tivoli Monitoring to reduce steady-state demand on system resources.

Steady-state demand is the constant demand, at all times and without interruption, on the availability of system resources. It exists even when users are not logged in to IBM Enterprise Portal Server. Here are some examples of jobs and processes that cause steady-state demand:

- Background collectors
- Historical collection
- Auto-started situations
- Policies

Reduce steady-state demand by:

- Distinguishing between critical and noncritical thresholds.
- Disabling data collection for noncritical resources.

- Tuning thresholds.
- Configuring data collection intervals.

Effective thresholds

You can use IBM Tivoli Monitoring to administer thresholds effectively and reduce demand on system resources.

IBM Tivoli Monitoring centralizes the administration of thresholds. Smaller staffs can manage larger numbers of systems. Availability threats can be monitored more frequently than less critical events, such as performance events. Instead of defining a refresh or sampling interval for a subsystem, you can define the interval for the threshold. This facility, if used with caution, provides you with an opportunity to lower demand on system resources and provide more availability protection. In addition, IBM Tivoli OMEGAMON XE and IBM Tivoli OMEGAMON DE interfaces use a single collector to provide raw data exceptions to multiple users. In enterprises where many people monitor the same subsystems, huge gains can be obtained.

Monitoring technologies are centered around the monitoring agent. When you upgrade to Tivoli OMEGAMON XE, you use a different approach to tuning than traditional IBM Tivoli OMEGAMON II products. The concepts of categorize, disable, and tune still apply. However, when you use Tivoli OMEGAMON XE, you can monitor applications more closely, and have more control over defining what is most important to detect.

Critical and noncritical thresholds

To determine which thresholds require monitoring, you must distinguish between critical and noncritical thresholds.

To avoid degradation in system performance, you must identify critical thresholds, including situations that require frequent sampling. For example, it is not necessary to monitor conditions that do not undergo frequent changes over long periods of time. Not all data is equally important and not all data requires the same amount of resources to gather. Many exceptions that place a heavy burden on system resources have a minimal effect on availability. Ideally, you want to maximize benefits and minimize resource usage.

You can get different answers to the following questions depending on the application that runs on the subsystem:

- Do you monitor production systems and non production systems?
- Do you plan to use the same sampling intervals for production systems and non production systems?

One size does not fit all. Sensible use of monitoring yields higher benefits at lower cost.

Noncritical data collection

You can disable the collection of data items that place a heavy burden on system resources without yielding significant benefits.

The objective of the IBM Tivoli OMEGAMON monitoring solution is to set low defaults to reduce demand on system resources. However, these defaults are set without prior knowledge of the unique environment or business applications of a customer. If, after careful investigation, you identify data collection items that place a heavy demand on system resources without providing much benefit, you disable them. Disabling such data collection items does not necessarily have an adverse affect on monitoring activity.

Threshold tuning

You can tune or use different collection methods for data collection items that place a high demand on system resources and yield high benefits.

In some products or subsystems, more than one method can be used to detect the same situation. Not all methods place the same demand on system resources.

IBM Tivoli OMEGAMON XE, because it uses Boolean logic, provides more methods to detect complex situations. For example, a critical application suffers poor direct access storage device (DASD) response times. The cause of the poor response times is that the volume is being heavily accessed by multiple applications on other z/OS partitions. You adjust the scope and frequency of numerous data collection settings to improve DASD response times.

IBM Tivoli OMEGAMON II uses indicator lights for thresholds and, typically, provides only one interval for each group of thresholds. With IBM Tivoli OMEGAMON XE, you can have different intervals for each threshold, including warnings and critical thresholds. IBM Tivoli OMEGAMON XE provides you with greater control and flexibility to help you reduce demand on system resources.

Situation monitoring and interval settings

You can take advantage of the architecture and efficiencies built into the design of IBM Tivoli OMEGAMON XE to further reduce or avoid excessive demand on system resources.

Identify the situations that you want to monitor and how often you must monitor them. When you become familiar with the design of the software, you realize that numerous different interval settings can result in excessive demands on system resources. You know when a less-frequent interval can reduce, or increase, demand on system resources.

Take into account how the software gathers data for exception thresholds. The software uses situations to represent exceptions. Situations help you make intelligent choices on setting intervals. The software has a navigator view with leaf names such as Application Summary, Cache CU Performance, and Virtual Tape Subsystems. Leaf names link to workspaces that contain numerous columns of data gathered by a single data collector.

If you right-click the **Address Space** leaf, for example, you can view all of the situations associated with that leaf. All of the situations on the leaf use the same data collector. If you use the same interval setting for all of the situations on the leaf, the situations are grouped. Because the situations are grouped, one call to the data collector is sufficient to collect the data. However, if each situation has a different interval setting, the data collector is called or run for each situation on the leaf.

If you change an interval from 2 minutes to 1 minute, you increase the demand on system resources. The software groups thresholds and collects all of the data from the same table or attribute group at the same time. When you change the interval, you schedule a new data collector and double the demand on system resources. If you change a third threshold and set it to 3 minutes, you triple the demand on system resources. If you want to sample less frequently, you must ensure that you change all situations that belong to the same attribute group to a less-frequent interval.

If a leaf has four or less situations, ensure that warning and critical situations have the same interval setting. If you do, the data collector must run only one time to collect the data. With five or more active situations, there is no gain, because the collector is already running two times. All situations with the same interval setting are scheduled together. For example, if you apply four warning situations for one interval and four critical situations for another higher interval, you can reduce demand on system resources. This approach assumes that the situations have two conditions or items to evaluate, as most product-provided situations have. If you code a complex seven-condition or evaluation criteria situation, the situation will probably not be scheduled with more than one additional situation. The actual limitation is on the size of the rule.

Before you measure the impact of your changes, you must restart the IBM Tivoli Enterprise Monitoring Server hub. The process of combining occurs only at start up for situations that are auto-started. In fact, you have a higher demand on system resources before you restart because anything that was previously combined was uncombined when you changed it.

User demands on system resources

You can reduce user-generated demands on system resources.

Each time a user requests data, a demand is made on system resources such as:

- When users use IBM Enterprise Portal Server clients:

  - To investigate alerts by navigating workspaces
  - With auto-update enabled on workspaces that contain graphs or tables
- When users frequently access workspaces with tabular views that show data from monitoring agents that return large volumes of data. Keep in mind that:
  - Auto-refreshed screens always require memory.
  - Large intermittent reports can cause spikes in memory utilization.

You can:

- Reduce IBM Enterprise Portal client and server demand on memory and processor resources.
- Improve response times for workspace:
  - Data retrievals
  - Screen refreshes

A query assigned to a chart or table view requests data from a specific attribute group. When you open or refresh a workspace, the query runs. You can substantially reduce the number of data samples that a query retrieves and the volume of data that a query retrieves:

- By customizing the query to filter out insignificant data
- By reducing the number of rows and columns in a query
- By applying the same query to multiple views in a workspace

The use of the auto-refresh facility can negatively affect performance. However, if you do use auto-refresh, set the interval setting to a long interval on a graphic view instead of a tabular view, to reduce the impact on performance.

Custom queries

You can use custom queries to reduce network traffic and processing cycles on the IBM Tivoli Enterprise Monitoring Agent and IBM Tivoli Enterprise Monitoring Server.

When you use custom queries, you can:

- Reduce memory usage on the IBM Enterprise Portal Server and client
- Limit the number of rows and columns passed from the agent to the IBM Tivoli Enterprise Monitoring Server

For example, you should use the Dataset Attribute Database feature to write a query that targets data sets on a specific control unit. Furthermore, you should refine your query so that it retrieves only monitoring information about space utilization. You can substantially reduce the amount of data that the query returns and reduce demand on system resources:

- By limiting the query to the specific control unit
- By specifying the attribute that you want to monitor

It is more efficient to use custom queries than it is to use View filters to filter data. View filters do not reduce demand on system resources. In fact, IBM Enterprise Portal requires more processor resources when you use View filters. The client applies filters to the current page. If a query returns more than one page, only a subset of the data is viewed on each page. IBM Enterprise Portal provides users with an option to increase the size of pages. When you select this option, more filtered data is displayed on each page. The memory requirements for the client increases, because the data that the query retrieves is stored in larger pages.

Number of columns in queries

You can reduce user demand on system resources by restricting the number of columns that custom queries return.

When you use predefined queries, data is returned for all of the columns that belong to the attribute group. An attribute group can contain 50 columns. You can create custom queries to target the specific columns that contain the information that you require. For example, the OMEGAMON for Storage Cache CU table has 28 columns. You should create a workspace that reports on storage for a large system. In particular, view information about de staging and staging rates. Create a custom query that selects the columns that contain this information and you exclude the other 26 columns from the query. In this way, you reduce the processor and memory requirements for IBM Enterprise Portal Server and client.

Reuse of custom queries in workspace

If you reuse one custom query for each table to generate data for all views in a workspace, you can reduce demand on system resources.

You want views to use the same custom query even if, for example, a graph view uses one attribute and a table view uses another attribute. If you create a unique query for each view, you increase:

- The workload of the monitoring agent to collect the data
- The demand on system resources

What works best is to use one query for each table in a workspace shared by all views within that workspace. The custom queries that you create must retrieve all the columns and rows that are required to populate all the views with data in your workspace. If you have many tables and charts that use different queries, the entire set of results for each query is stored on the IBM Enterprise Portal Server. IBM Enterprise Portal client stores:

- 100 rows (default page size) from each query that is currently being viewed
- Previous page of any view

Historical data collection

OMEGAMON for Storage provides both short-term and long-term historical data collection. It is important to distinguish between the two when collecting data.

See:

- Short-term historical data
- Long-term historical data

After you activate historical data collection, an icon is displayed in qualifying views in IBM Enterprise Portal workspaces. To produce meaningful reports, you must wait until sufficient historical data is collected and stored. You can click the icon to extend any existing IBM Enterprise Portal view (also called a report) to include historical data. IBM Enterprise Portal reports automatically pull data from short-term and long-term history based upon the time period that you specify for the report.

You can configure different collection intervals for historical and real-time data. To avoid excessive demand on processor resources and to decrease storage consumption, historical data collection is typically performed less frequently than real-time data collection. You can configure a short-term historical data collection interval of 5, 15, 30, or 60 minutes.

Writing data to long-term history can be configured for 24 hours or 1 hour, or you can turn it off. If you configure long-term history, select a warehousing interval of 1 hour to avoid transferring 24 hours worth of historical data at one time. Shorter intervals reduce the duration of processor usage associated with writing data to the warehouse by distributing the writing of data across 24 time periods.

Short-term historical data

Short-term historical data is stored in the persistent data store on z/OS systems or in files on distributed systems. Short-term historical data comprises data that is stored for 24 hours or less. However, the amount and age of short-term historical data that OMEGAMON for Storage keeps, depends on such factors as:

- The number of resources that are being monitored
- The amount of disk space that is configured for use by the persistent data store

To find out how to configure the collection of short-term historical data for the persistent data store, see the appendix of the *IBMOMEGAMON for Storage: Planning and Configuration Guide*.

Short-term historical data is used to investigate and determine the nature of problems that arise.

Tip: Use PARMGEN to configure data sets that persistent data store uses.

Long-term historical data

Long-term historical data is stored in the IBM Tivoli Data Warehouse. The long-term history database can retain data collected by OMEGAMON for Storage monitoring agents for days, weeks, months, or years. See the following documents for detailed information about the collection of historical data:

- *IBM Tivoli Monitoring: Installation and Setup Guide*
- *IBM Tivoli Monitoring: Administrator's Guide*

Use long-term historical data to analyze trends and determine workload balance. Tivoli Data Warehouse is the primary tool for collecting long-term historical data that IBM Tivoli Monitoring gathers. The warehouse architecture requires a relational database. The warehouse installer provides an IBM DB2® relational database. However, the warehouse can use, instead of DB2, an Oracle, orMicrosoft SQL relational databases. Long-term history also requires installation of the Warehouse Proxy Agent software (provided). See *IBM Tivoli Monitoring: Installation and Setup Guide* for information about setting up the collection of long-term historical data. You must activate and configure short-term historical data collection if you want to collect long-term historical data.

Sample collection

Set the following intervals for data collection.

Table 1. Data collection sample interval settings

Type of interval	**Duration in minutes**
Real-time data collection	5
Short-term historical data collection	15
Long-term warehousing	60

The real-time data collection is based on the intervals that you set.

Table 2. Real-time data sample collection schedule

Time	**Real-time data collection**
1:57	First row of data, 798 bytes, collected.
2:02	Second row of data collected (replaces first row of data, which is discarded).
2:07	Third row of data collected (replaces second row of data, which is discarded).
2:12	Fourth row of data collected (replaces third row of data, which is discarded).
2:17	Fifth row of data collected (replaces fourth row of data, which is discarded).

Table 2. Real-time data sample collection schedule (continued)	
Time	**Real-time data collection**
2:22	Sixth row of data collected (replaces fifth row of data, which is discarded).
2:27	Seventh row of data collected (replaces sixth row of data, which is discarded).
2:32	Eighth row of data collected (replaces seventh row of data, which is discarded).
2:37	Ninth row of data collected (replaces eight row of data, which is discarded).
2:42	Tenth row of data collected (replaces ninth row of data, which is discarded).
2:47	Eleventh row of data collected (replaces tenth row of data, which is discarded).
2:52	Twelfth row of data collected (replaces eleventh row of data, which is discarded).
2:57	Start of next cycle of real-time data collection.

The short-term historical data collection is based on the intervals that you set.

Table 3. Short-term historical data sample collection schedule	
Time	**Short-term historical data collection**
1:58	Start collection.
2:00	Use the most recent real-time collection to store the first row of data (816 bytes).
2:15	Use the most recent real-time collection to store the second row of data (816 bytes).
2:30	Use the most recent real-time collection to store the third row of data (816 bytes).
2:45	Use the most recent real-time collection to store the fourth row of data (816 bytes).
3:00	Use the most recent real-time collection to store the fifth row of data (816 bytes).

Every hour, the Warehouse Proxy Agent transfers all rows of data to the database as **long-term historical data** that is maintained by IBM Tivoli Data Warehouse. You can use the following formula to calculate the volume of data, in bytes, that is collected: 4 x 816 x *n*. In the formula, *n* stands for the total number of data sets in all data set groups.

Types of historical data

You must decide which types of historical data to store in short-term and long-term history and you must also decide how long to store the data. Because data collection consumes processor cycles and disk space, there are, inevitably, trade-offs. Writing data to short-term history is cost-effective and typically much less costly than writing to long-term history.

If you choose to configure short-term historical data collection, an additional amount of storage is used on z/OS. If you also choose to configure long-term history (in the Tivoli Data Warehouse) an additional amount of storage is used on a distributed server (for example, on Windows) for the historical data that your company needs to retain for days, weeks, months, or years.

Short-term history is written to disk. This operation is typically performed at the monitoring agent and it consumes processor cycles. Additional processor cycles are used when the Warehouse Proxy Agent receives data from short-term history and transfers it to the data warehouse. If you collect a large amount of data in short-term history, the extraction process significantly increases the use of processor resources by the monitoring agent.

For many companies, the following configurations and settings offer the best utilization of processor and storage resources:

- Short-term history location: persistent data store defined at the z/OS agent.
- Long-term history location: If you are monitoring fewer than 5,000 resources, you might want to use the same computer that hosts the IBM Enterprise Portal Server.
- Short-term history collection interval: 15 minutes or less frequently.
- Warehousing interval: hourly.

Depending on your requirements, you can configure historical data collection for only a subset of attribute groups. This method is effective in limiting storage consumption, in particular, when you choose not to perform historical data collection for the following high-volume attribute groups:

- I/O at the data set level
- Attribute groups that have many bytes per row (many attributes), such as data set group details

Avoid the collection of data that you do not require for historical reports.

You can use this information as a basis for choosing which attribute groups to enable for historical collection. You can select individual attribute groups for historical collection, including specifying different historical collection intervals and warehouse intervals.

By default, historical reports retrieve up to 24 hours of data from short-term history. If your persistent data store is not allocated with sufficient space, you will not have 24 hours of short-term data to retrieve. Allocate your persistent data store to hold a full 24 hours of data or change the default of 24 hours. See the *Configuring IBM OMEGAMON for Storage and OMEGAMON II for SMS* manual for information about how to change the default of 24 hours.

Because historical data accumulates, you must also determine how long you want to keep the data.

- Short-term history data in the persistent data store automatically wraps, and does not need to be maintained. You can also run a KPDXTRA job to write short-term history to flat files, for backup, or for analysis in a statistical or graphing package. See the *IBM OMEGAMON for Storage: Planning and Configuration Guide* for details on the persistent data store and the KPDXTRA job.
- Long-term historical data is stored in the Tivoli Data Warehouse in a DB2, UDB, Oracle, or Microsoft SQL database. To perform aggregation and pruning functions on the data, install the Summarization and Pruning Agent for the warehouse, as described in the following documents:
  - *IBM Tivoli Monitoring: Installation and Setup Guide*.
  - *IBM Tivoli Monitoring: Administrator's Guide*.

Chapter 2. Modifying parameters for data collection

You can modify the default parameters for the collection of data sets to reduce demand on system resources and improve performance.

In particular, you can reconfigure data set members or data set parameters:

- That regulate data collection frequency such as interval parameters
- That exclude volumes from data set collection
- That record and write data to the persistent data store
- That record statistics about direct access storage devices
- That regulate data set monitoring

You use PARMGEN and the reference information in the *IBM OMEGAMON for Storage Parameter Reference* to modify data set members and parameters.

Attention: If you update data set parameters manually, the changes that you make are overwritten the next time you use PARMGEN to update parameters.

Reduce demand on system resources and improve performance

These are some areas in which you can make changes to reduce demand on system resources and improve performance:

Collection controls

The collection controls parameters control the frequency of data collections for cached volumes and control units. They also control whether cache collection is enabled or disabled.

Application collection intervals

The application collection interval parameters control how often OMEGAMON for Storage rebuilds the list of volumes and data sets of the applications that it monitors.

Response time collection intervals

The response time collection interval parameters control how often response time statistics for direct access storage device (DASD) volumes are collected.

Space collection frequency

The space collection frequency parameters control how often OMEGAMON for Storage collects space and fragmentation statistics for direct access storage device (DASD) volumes.

Write parameters for the persistent data store

The historical collection parameters determine whether OMEGAMON for Storage writes information about the direct access storage devices (DASD) and data sets that it monitors toIBM Tivoli Enterprise Monitoring Server's persistent data store (PDS) at the end of each Resource Measurement Facility™ (RMF™) interval.

To improve performance, you can disable the writing of information to IBM Tivoli Enterprise Monitoring Server's persistent data store.

Direct access storage device statistics collection

The collection controls parameters control how cache, response time, and space and fragmentation statistics about direct access storage devices are recorded on the System Management Facilities (SMF).

Space and fragmentation statistics collection

The DASD space exclusion list parameters control the exclusion of volumes from data collection. To optimize the collection of cache statistics, designate a single LPAR to collect statistics. See "Optimizing the collection of cache controller and DASD space statistics" on page 14.

Exclude volumes from the collection of statistics

You can exclude volumes from the collection of space and fragmentation statistics.

Response time and cache statistics collection

The collection controls parameters select a monitoring method for response time and cache statistics collection.

Collection of information for the persistent data store

You can control the recording of information to the persistent data store of the IBM Tivoli Enterprise Monitoring Server. If historical information is recorded to the persistent data store, you can view it online.

Data stores, which can contain one or more data sets, can use either of these processes:

- Switching
- Wrapping

Switching
Switching applies to multiple data sets and it occurs when a data set becomes full. When the data set becomes full, recording stops. It resumes recording with the next empty data set. When recording begins on the last empty data set, the data set containing the oldest data is emptied. This ensures that space is available for switching so that recording continues when the last data set fills.

Wrapping
Wrapping applies to single data sets and occurs when the data set becomes full. New records replace the oldest records in the data set. The oldest records and their associated indexes must be located and deleted before the new records are inserted. Because of this, wrapping can have a negative impact on system performance.

Data set monitoring: reducing demand on system resources

You can use application monitoring, device selection, and monitoring option keywords, to reduce the demand on system resources by devices that monitor data sets.

You use reference information to reduce the demand on system resources that devices that monitor data sets generate when OMEGAMON II for SMS collects data about address spaces. See:

- "Application monitoring" on page 10.
- "Device selection keywords" on page 11.
- "Device monitoring option keywords" on page 12.

Application monitoring

Perhaps the most effective and efficient way of collecting data set-level I/O statistics is to forego setting collection options for DASD volumes or groups and instead use the Application Monitoring feature. This allows you to monitor all I/O for a job or started task down to the data set level. If the data sets move to different volumes, the monitoring moves with them. In this way, you can avoid incurring the CPU overhead of monitoring for a larger group of DASD than is actually required. For more information about this feature, see *IBM OMEGAMON for Storage User's Guide*.

Device selection keywords

Use device selection keywords to specify a device or range of devices, and to initialize the monitoring status of a device. A keyword value can specify a single device or a list of devices separated by commas or blanks. Each device that you add to a list can also include an optional initial monitoring status such as ON or OFF. If you include an initial monitoring status, you must enclose the device and its monitoring status in parentheses:

(XXXX,ON)

(XXXX,OFF)

By default, the initial monitoring status is set to ON. Keyword values can include the optional sub keyword called SAMPCT which you use to specify how often device I/O is collected. If you include the SAMPCT sub keyword, the value you specify applies to all the devices that you add to the keyword list. You can specify a keyword more than one time if the list of devices does not fit on one line. If no keyword options are specified, device monitoring is not performed.

DEVICES keyword

Use the DEVICES keyword to list one or more hexadecimal device numbers:

DEVICES(030F)

DEVICES(030A (0318,OFF))

DEVRANGE keyword

Use the DEVRANGE keyword to specify a range of hexadecimal device numbers:

DEVRANGE((0210:021A,OFF),0310:031A)

DEVRANGE(021B:021F,0304:0308)

VOLSERS keyword

Use the VOLSERS keyword to list on one or more device VOLSERs. You can use a question mark (?) as a wildcard character. For each VOLSER that you add, you can use only one wildcard character. You must use the wildcard character to specify the last character of the VOLSER number.

VOLSERS(VOL111,VOL2?)

VOLSERS((VOL22?,OFF),VOL333)

ALL keyword

Use the ALL keyword to set the initial monitoring status of all devices, without listing them, to on. You can use the SAMPCT optional subkeyword with the ALL keyword.

ALL

ALL(SAMPCT=1)

SAMPCT subkeyword

Use the SAMPCT optional subkeyword to specify how often device I/O data is collected. You can use the SAMPCT optional subkeyword with all, including the ALL keyword, device keywords. The value of SAMPCT can range from 1 to 99. The default is 1. When SAMPCT is 1, data is collected for every device I/O. When SAMPCT is 5, data is collected every fifth I/O cycle. The value of SAMPCT is applied to each device keyword separately.

ALL(SAMPCT=2)

VOLSERS(VOL111,VOL2?,SAMPCT=2)

VOLSERS(VOL3?,SAMPCT=5)

DEVICES(030A (0318,OFF),SAMPCT=4)

MSR subkeyword

Use the MSR optional subkeyword to specify that device I/O is to be collected on an exception basis. Data is not collected for the device until the response time for the device crosses the threshold a specified number of times. The keyword MSRTARG is used to control the number of times an exception must occur before monitoring takes place. You can use the MSR optional subkeyword with all device keywords including the ALL keyword.

Tip: To find out how to set the threshold for exceptions, see the MSRTARG keyword section in "Device monitoring option keywords" on page 12.

ALL(MSR=35)

VOLSERS(VOL111,VOL2?,MSR=25)

Device monitoring option keywords

Use keywords for device monitoring to configure exception thresholds and turn data collection for the SMS storage class on or off.

MSRTARG keyword

Use the MSRTARG keyword to specify how often a response time exception must occur. Exception level monitoring begins when an exception occurs in a discrete sample of 100 consecutive I/O transactions for a device. Valid range is 1 to 99 and the default value, if not specified, is 51.

Tip: To find out how you configure devices to collect I/O data when exceptions occur, see the MSR subkeyword section in "Device selection keywords" on page 11.

MSRTARG(51)

SCNAME keyword

Use the SCNAME keyword to turn the collection of SMS storage class name (ON) or (OFF).

SCNAME(ON)

Chapter 3. Tuning product features

Tune OMEGAMON for Storage product features to improve performance and reduce demand on system resources.

You can reduce the impact of product-specific features on system resources:

- By following definition guidelines and maintaining data set masks
- By following historical data collection guidelines
- By following wildcard usage guidelines

You can also improve performance and reduce the demand on system resources:

- By using a single LPAR to collect cache controller and DASD space information
- By sharing information collected by an LPAR with all LPARs reporting to the same IBM Tivoli Enterprise Monitoring Server

Definition guidelines for data set masks

Follow the guidelines for defining data set masks and delete or deactivate unused data set masks.

When you define data set masks:

- Keep amount of data small.
- Keep amount of time to collect data as short as possible.
- Avoid duplication of data in a group:
  - By ensuring that the defined masks are unique
  - By ensuring that they do not overlap

When you specify a data set mask, you can clear the Space Data option:

- When space-related data, such as tracks allocated or used or both, are not required
- When the last reference date is not required

Tip: Delete the data set masks and groups that you no longer use. If you do not want to permanently delete a data set mask or group, you can deactivate it by editing it and clearing the Active option.

Guidelines for collecting historical data

Follow the guidelines for collecting historical data and determining data collection rates.

If you use historical recording for data set groups:

- Collect only the data that you need.
- Choose an appropriate data collection rate.

You can set historical data collection rates to:

- Every minute
- Every 5 minutes
- Every 15 minutes
- Every 30 minutes
- Every 60 minutes
- Every 1,440 minutes (every day)

If you choose a high frequency rate for collecting historical data:

- Historical files fill up or wrap faster.

- More processor and I/O resources of the IBM Tivoli Enterprise Monitoring Server are used.

For more information about setting collection rates, see *IBM OMEGAMON for Storage: Planning and Configuration Guide*.

Guidelines for using wildcards

To reduce demand on system resources, follow the guidelines about using wildcards.

Try not to use wildcard characters, such as %, or +, or ++ in the first qualifier of a data set mask. When the highest-level qualifier contains a wildcard, the software cannot use the standard mechanism to efficiently locate all matching data sets. Instead, the software must search and examine all of the catalogs to find matching data sets. This resource-intensive activity occurs every time the data is refreshed. When you retrieve a list, make sure that you delete the mask from the group.

If in a full catalog search, the total number of data sets on the system is large, a request may not begin before the next data collection cycle begins. In such cases, the collection cycle is adjusted dynamically so that it takes at least as long as the time it took to complete the request.

Optimizing the collection of cache controller and DASD space statistics

To optimize the collection of cache statistics, designate a single LPAR to collect statistics.

In shared DASD environments where multiple LPARs report to the same Tivoli Enterprise Monitoring System hub:

- You can designate a single LPAR to collect cache space statistics.
- You can deactivate all other LPARs from collecting cache space statistics.

You must use PARMGEN to designate the LPAR that collects cache controller statistics. You must also disable non designated LPARs from collecting cache controller statistics.

The statistics that are collected by the designated LPAR are made available to the LPARs that have been deactivated. So, you can view the statistics that are collected by one LPAR on all LPARs.

Prerequisites

To make cache controller and DASD space statistics available to LPARs that do not collect statistics:

- You myust install OMEGAMON for Storage, component 5608A 1000

To make DASD space statistics available to LPARs that do not collect statistics:

- You must exclude all volumes from collecting DASD space statistics

If you specify a volume to collect DASD space statistics, you disable the sharing of DASD space statistics with LPARs that do not collect statistics.

Restriction: The sharing of cache controller and DASD space statistics is not available to customers who only configure Tivoli OMEGAMON II for SMS (component 5608A0500).

Related tasks:

You modify the cache collection interval parameter to control how often or whether OMEGAMON for Storage collects cache statistics for cached volumes and control units.

Chapter 4. Modifying parameters for data collection

You can modify the default parameters for the collection of data sets to reduce demand on system resources and improve performance.

In particular, you can reconfigure data set members or data set parameters:

- That regulate data collection frequency such as interval parameters
- That exclude volumes from data set collection
- That record and write data to the persistent data store
- That record statistics about direct access storage devices
- That regulate data set monitoring

You use PARMGEN and the reference information in the *IBM OMEGAMON for Storage Parameter Reference* to modify data set members and parameters.

Attention: If you update data set parameters manually, the changes that you make are overwritten the next time you use PARMGEN to update parameters.

Reduce demand on system resources and improve performance

These are some areas in which you can make changes to reduce demand on system resources and improve performance:

Collection controls

The collection controls parameters control the frequency of data collections for cached volumes and control units. They also control whether cache collection is enabled or disabled.

Application collection intervals

The application collection interval parameters control how often OMEGAMON for Storage rebuilds the list of volumes and data sets of the applications that it monitors.

Response time collection intervals

The response time collection interval parameters control how often response time statistics for direct access storage device (DASD) volumes are collected.

Space collection frequency

The space collection frequency parameters control how often OMEGAMON for Storage collects space and fragmentation statistics for direct access storage device (DASD) volumes.

Write parameters for the persistent data store

The historical collection parameters determine whether OMEGAMON for Storage writes information about the direct access storage devices (DASD) and data sets that it monitors toIBM Tivoli Enterprise Monitoring Server's persistent data store (PDS) at the end of each Resource Measurement Facility™ (RMF™) interval.

To improve performance, you can disable the writing of information to IBM Tivoli Enterprise Monitoring Server's persistent data store.

Direct access storage device statistics collection

The collection controls parameters control how cache, response time, and space and fragmentation statistics about direct access storage devices are recorded on the System Management Facilities (SMF).

Space and fragmentation statistics collection

The DASD space exclusion list parameters control the exclusion of volumes from data collection. To optimize the collection of cache statistics, designate a single LPAR to collect statistics. See "Optimizing the collection of cache controller and DASD space statistics" on page 14.

Exclude volumes from the collection of statistics

You can exclude volumes from the collection of space and fragmentation statistics.

Response time and cache statistics collection

The collection controls parameters select a monitoring method for response time and cache statistics collection.

Collection of information for the persistent data store

You can control the recording of information to the persistent data store of the IBM Tivoli Enterprise Monitoring Server. If historical information is recorded to the persistent data store, you can view it online.

Data stores, which can contain one or more data sets, can use either of these processes:

- Switching
- Wrapping

Switching
Switching applies to multiple data sets and it occurs when a data set becomes full. When the data set becomes full, recording stops. It resumes recording with the next empty data set. When recording begins on the last empty data set, the data set containing the oldest data is emptied. This ensures that space is available for switching so that recording continues when the last data set fills.

Wrapping
Wrapping applies to single data sets and occurs when the data set becomes full. New records replace the oldest records in the data set. The oldest records and their associated indexes must be located and deleted before the new records are inserted. Because of this, wrapping can have a negative impact on system performance.

Data set monitoring: reducing demand on system resources

You can use application monitoring, device selection, and monitoring option keywords, to reduce the demand on system resources by devices that monitor data sets.

You use reference information to reduce the demand on system resources that devices that monitor data sets generate when OMEGAMON II for SMS collects data about address spaces. See:

- "Application monitoring" on page 10.
- "Device selection keywords" on page 11.
- "Device monitoring option keywords" on page 12.

Application monitoring

Perhaps the most effective and efficient way of collecting data set-level I/O statistics is to forego setting collection options for DASD volumes or groups and instead use the Application Monitoring feature. This allows you to monitor all I/O for a job or started task down to the data set level. If the data sets move to different volumes, the monitoring moves with them. In this way, you can avoid incurring the CPU overhead of monitoring for a larger group of DASD than is actually required. For more information about this feature, see *IBM OMEGAMON for Storage User's Guide*.

Device selection keywords

Use device selection keywords to specify a device or range of devices, and to initialize the monitoring status of a device. A keyword value can specify a single device or a list of devices separated by commas or blanks. Each device that you add to a list can also include an optional initial monitoring status such as ON or OFF. If you include an initial monitoring status, you must enclose the device and its monitoring status in parentheses:

(XXXX,ON)

(XXXX,OFF)

By default, the initial monitoring status is set to ON. Keyword values can include the optional sub keyword called SAMPCT which you use to specify how often device I/O is collected. If you include the SAMPCT sub keyword, the value you specify applies to all the devices that you add to the keyword list. You can specify a keyword more than one time if the list of devices does not fit on one line. If no keyword options are specified, device monitoring is not performed.

DEVICES keyword

Use the DEVICES keyword to list one or more hexadecimal device numbers:

DEVICES(030F)

DEVICES(030A (0318,OFF))

DEVRANGE keyword

Use the DEVRANGE keyword to specify a range of hexadecimal device numbers:

DEVRANGE((0210:021A,OFF),0310:031A)

DEVRANGE(021B:021F,0304:0308)

VOLSERS keyword

Use the VOLSERS keyword to list on one or more device VOLSERs. You can use a question mark (?) as a wildcard character. For each VOLSER that you add, you can use only one wildcard character. You must use the wildcard character to specify the last character of the VOLSER number.

VOLSERS(VOL111,VOL2?)

VOLSERS((VOL22?,OFF),VOL333)

ALL keyword

Use the ALL keyword to set the initial monitoring status of all devices, without listing them, to on. You can use the SAMPCT optional subkeyword with the ALL keyword.

ALL

ALL(SAMPCT=1)

SAMPCT subkeyword

Use the SAMPCT optional subkeyword to specify how often device I/O data is collected. You can use the SAMPCT optional subkeyword with all, including the ALL keyword, device keywords. The value of SAMPCT can range from 1 to 99. The default is 1. When SAMPCT is 1, data is collected for every device I/O. When SAMPCT is 5, data is collected every fifth I/O cycle. The value of SAMPCT is applied to each device keyword separately.

ALL(SAMPCT=2)

VOLSERS(VOL111,VOL2?,SAMPCT=2)

VOLSERS(VOL3?,SAMPCT=5)

DEVICES(030A (0318,OFF),SAMPCT=4)

MSR subkeyword

Use the MSR optional subkeyword to specify that device I/O is to be collected on an exception basis. Data is not collected for the device until the response time for the device crosses the threshold a specified number of times. The keyword MSRTARG is used to control the number of times an exception must occur before monitoring takes place. You can use the MSR optional subkeyword with all device keywords including the ALL keyword.

Tip: To find out how to set the threshold for exceptions, see the MSRTARG keyword section in "Device monitoring option keywords" on page 12.

ALL(MSR=35)

VOLSERS(VOL111,VOL2?,MSR=25)

Device monitoring option keywords

Use keywords for device monitoring to configure exception thresholds and turn data collection for the SMS storage class on or off.

MSRTARG keyword

Use the MSRTARG keyword to specify how often a response time exception must occur. Exception level monitoring begins when an exception occurs in a discrete sample of 100 consecutive I/O transactions for a device. Valid range is 1 to 99 and the default value, if not specified, is 51.

Tip: To find out how you configure devices to collect I/O data when exceptions occur, see the MSR subkeyword section in "Device selection keywords" on page 11.

MSRTARG(51)

SCNAME keyword

Use the SCNAME keyword to turn the collection of SMS storage class name (ON) or (OFF).

SCNAME(ON)

Chapter 5. Tuning product features

Tune OMEGAMON for Storage product features to improve performance and reduce demand on system resources.

You can reduce the impact of product-specific features on system resources:

- By following definition guidelines and maintaining data set masks
- By following historical data collection guidelines
- By following wildcard usage guidelines

You can also improve performance and reduce the demand on system resources:

- By using a single LPAR to collect cache controller and DASD space information
- By sharing information collected by an LPAR with all LPARs reporting to the same IBM Tivoli Enterprise Monitoring Server

Definition guidelines for data set masks

Follow the guidelines for defining data set masks and delete or deactivate unused data set masks.

When you define data set masks:

- Keep amount of data small.
- Keep amount of time to collect data as short as possible.
- Avoid duplication of data in a group:
  - By ensuring that the defined masks are unique
  - By ensuring that they do not overlap

When you specify a data set mask, you can clear the Space Data option:

- When space-related data, such as tracks allocated or used or both, are not required
- When the last reference date is not required

Tip: Delete the data set masks and groups that you no longer use. If you do not want to permanently delete a data set mask or group, you can deactivate it by editing it and clearing the Active option.

Guidelines for collecting historical data

Follow the guidelines for collecting historical data and determining data collection rates.

If you use historical recording for data set groups:

- Collect only the data that you need.
- Choose an appropriate data collection rate.

You can set historical data collection rates to:

- Every minute
- Every 5 minutes
- Every 15 minutes
- Every 30 minutes
- Every 60 minutes
- Every 1,440 minutes (every day)

If you choose a high frequency rate for collecting historical data:

- Historical files fill up or wrap faster.

- More processor and I/O resources of the IBM Tivoli Enterprise Monitoring Server are used.

For more information about setting collection rates, see *IBM OMEGAMON for Storage: Planning and Configuration Guide*.

Guidelines for using wildcards

To reduce demand on system resources, follow the guidelines about using wildcards.

Try not to use wildcard characters, such as %, or +, or ++ in the first qualifier of a data set mask. When the highest-level qualifier contains a wildcard, the software cannot use the standard mechanism to efficiently locate all matching data sets. Instead, the software must search and examine all of the catalogs to find matching data sets. This resource-intensive activity occurs every time the data is refreshed. When you retrieve a list, make sure that you delete the mask from the group.

If in a full catalog search, the total number of data sets on the system is large, a request may not begin before the next data collection cycle begins. In such cases, the collection cycle is adjusted dynamically so that it takes at least as long as the time it took to complete the request.

Optimizing the collection of cache controller and DASD space statistics

To optimize the collection of cache statistics, designate a single LPAR to collect statistics.

In shared DASD environments where multiple LPARs report to the same Tivoli Enterprise Monitoring System hub:

- You can designate a single LPAR to collect cache space statistics.
- You can deactivate all other LPARs from collecting cache space statistics.

You must use PARMGEN to designate the LPAR that collects cache controller statistics. You must also disable non designated LPARs from collecting cache controller statistics.

The statistics that are collected by the designated LPAR are made available to the LPARs that have been deactivated. So, you can view the statistics that are collected by one LPAR on all LPARs.

Prerequisites

To make cache controller and DASD space statistics available to LPARs that do not collect statistics:

- You myust install OMEGAMON for Storage, component 5608A 1000

To make DASD space statistics available to LPARs that do not collect statistics:

- You must exclude all volumes from collecting DASD space statistics

If you specify a volume to collect DASD space statistics, you disable the sharing of DASD space statistics with LPARs that do not collect statistics.

Restriction: The sharing of cache controller and DASD space statistics is not available to customers who only configure Tivoli OMEGAMON II for SMS (component 5608A0500).

Related tasks:

You modify the cache collection interval parameter to control how often or whether OMEGAMON for Storage collects cache statistics for cached volumes and control units.

Index

Accessibility

Accessibility features help users with physical disabilities, such as restricted mobility or limited vision, to use software products successfully. OMEGAMON® XE monitoring products support several user interfaces. Product functionality and accessibility features vary according to the interface.

The major accessibility features in this product enable users in the following ways:

- Use assistive technologies, such as screen-reader software and digital speech synthesizer, to hear what is displayed on the screen. Consult the product documentation of the assistive technology for details on using those technologies with this product.
- Operate specific or equivalent features using only the keyboard.
- Magnify what is displayed on the screen.

In addition, the product documentation was modified to include the following features to aid accessibility:

- All documentation is available in both HTML and convertible PDF formats to give the maximum opportunity for users to apply screen-reader software.
- All images in the documentation are provided with alternative text so that users with vision impairments can understand the contents of the images.

Interface information

The Tivoli® Enterprise Portal interface offers the greatest range of functionality, but is not entirely accessible. The OMEGAMON Enhanced 3270 user interface offers more limited functionality, but is entirely accessible. (The enhanced 3270 user interface supports all the accessibility features supported by your emulator. If you are using IBM® Personal Communications, you can find information on its accessibility features at http://publib.boulder.ibm.com/infocenter/pcomhelp/v6r0/index.jsp?topic=/com.ibm.pcomm.doc/books/html/quick_beginnings10.htm. If you are using a third-party emulator, see the documentation for that product for accessibility information.)

The OMEGAMON ("classic") and OMEGAMON II (CUA) 3270 interfaces use an ISPF style interface. Standard and custom PF Key settings, menu options, and command line interface options allow for short cuts to commonly viewed screens. While basic customization options allow for highlights and other eye-catcher techniques to be added to the interface, the customization options are limited.

Related accessibility information

You can view the publications using the Adobe Acrobat Reader.

IBM and accessibility

See the IBM Human Ability and Accessibility Center for more information about the commitment that IBM has to accessibility.

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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, MD-NC119
Armonk, NY 10504-1785
US

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

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 jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites 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 Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

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 information 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.

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.

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

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.

SC27-4380-04