IBM® Z OMEGAMON® AI for Storage®

IBM Z OMEGAMON AI for Storage User's Guide

6.1

SC31-5751-00

IBM

Note

Before using this information and the product it supports, read the information in “Notices” on page 414.

Edition notice

This edition applies to version 6.1 of IBM® Z OMEGAMON® AI for Storage® and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

User's guide 13
Workspaces 14
Nodes in the Navigator that are linked to workspaces 15
Accessing workspaces 15
Standard path and method for navigating workspaces 16
Historical navigation 17
Integration with other products 18
Links to workspaces of related products 18
About dynamic linking to cross-product workspaces 20
Associating workspaces with attributes 20
Using IBM Z OMEGAMON AI for Storage information 20
Defining workspace properties 21
Investigating a situation event 21
Filtering information 21
Sorting information 21
Adding a workspace to your favorites 21
Monitoring applications 21
Navigating the workspaces of the Application Summary node 22
Application Summary workspace 22
Summary information for the workspaces in this node 24
Managing application definitions 26
Adding an application definition 26
Editing an application definition 26
Deleting an application definition 27
Monitoring channel paths 27
Navigating the workspaces of the Channel Path node 27
Channel Path workspace 28
Summary information for the workspaces in this node 29
Monitoring cache CU performance 30
Navigating the workspaces of the Cache CU Performance node 30
Cache CU Performance workspace 32
Summary information for the workspaces in this node 33
Monitoring cache CU status 42
Navigating the workspaces of the Cache CU Status node 42
Cache CU Status workspace 44
Summary information for the workspaces in this node 44
Monitoring logical control units 45
Navigating the workspaces of the Logical Control Unit node 45
Logical Control Unit workspace 47
Summary information for the workspaces in this node 48
Monitoring tape groups 49
Navigating the workspaces of the Tape Group node 50
Tape Group workspace 50
Summary information for the workspaces in this node 50
Monitoring Virtual Tape Subsystems (VTS) 51
Navigating the workspaces of the Virtual Tape Subsystems node 51
Virtual Tape Subsystems workspace 52
Summary information for the workspaces in this node 54
Monitoring SMS storage groups performance 58
Navigating the workspaces of the SMS Storage Groups Performance node 58
SMS Storage Groups Performance workspace 61
Summary information for the workspaces in this node 61
Monitoring SMS storage groups space 66
Navigating the workspaces of the SMS Storage Groups Space node 66
SMS Storage Groups Space workspace 67
Summary information for the workspaces in this node 68
Monitoring the performance of user DASD groups 70
Navigating the workspaces of the User DASD Groups Performance node 71
User DASD Groups Performance workspace 73
Summary information for the workspaces in this node 73
Adding a user DASD group 74
Editing a user DASD group 75
Deleting a user DASD group 75
Monitoring the space of user DASD groups 75
Navigating the workspaces of the User DASD Groups Space node 76

User DASD Groups Space workspace ... 76
Issuing Group level Toolkit commands ... 77
Summary information for the workspaces in this node ... 77
Adding a user DASD group ... 78
Editing a user DASD group ... 79
Deleting a user DASD group ... 79
Monitoring user storage groups ... 79
Linked workspaces ... 80
User Groups for Storage Groups workspace ... 80
Adding a user storage group ... 81
Editing a user storage group ... 82
Deleting a user storage group ... 82
Monitoring DFSMShsm™ status ... 82
Navigating the workspaces of the DFSMShsm™ Status node ... 82
DFSMShsm™ Status workspace ... 83
Summary information for the workspaces in this node ... 84
Monitoring Tape Management status ... 86
Navigating the workspaces of the Tape Management Status node ... 86
Monitoring DFSMSrmm™ status ... 86
Monitoring CA 1 status ... 96
Monitoring Record Level Sharing ... 102
Navigating the workspaces of the Record Level Sharing Summary node ... 102
Record Level Sharing Summary workspace ... 103
Summary information for the workspaces in this node ... 109
Monitoring Dataset Attributes System Summary data ... 112
Navigating the workspaces of the Dataset Attributes System Summary node ... 112
Dataset Attributes System Summary workspace ... 114
Summary information for the workspaces in this node ... 115
Monitoring dataset groups summary data ... 118
Navigating the workspaces of the Dataset Group Summary node ... 119
Dataset Group Summary workspace ... 119
Summary information for the workspaces in this node ... 121
Monitoring SMS configuration ... 123
Navigating the workspaces of the SMS Configuration node ... 123
SMS Configuration workspace ... 123
Summary information for the workspaces in this node ... 124
Monitoring Storage Toolkit action requests ... 126
Navigating the workspaces of the Storage Toolkit node ... 126
Storage Toolkit Action Requests workspace ... 126
Summary information for the workspaces in this node ... 127
Monitoring Copy Services ... 127
Navigating the workspaces of Copy Services node ... 128
z/GlobalMirror (XRC) Sessions workspace ... 128
Summary information for the workspaces in this node ... 129
Monitoring ADABAS Database Details ... 130
Navigating the workspaces of the ADABAS Summary node ... 130
ADABAS Summary workspace ... 131
Summary information for the workspaces in this node ... 133
Directly monitoring situations ... 133
Realtime Dataset Metrics ... 134
Realtime Dataset Metrics overview ... 134
Launching Realtime Dataset Metrics ... 134
Selecting target system ... 134
Realtime Dataset Metrics basic features ... 135
View By ... 135
Column Set ... 136
Filters ... 136
Run Query ... 137
Export ... 137
Inline editing ... 137
History mode ... 138
Storage Toolkit for Realtime Dataset Metrics ... 138
Create Batch Job template dialog box ... 139
New Command template dialog box ... 139
Session Jobs ... 139
User actions in Storage Toolkit for Realtime Dataset Metrics only ... 140
Create Batch Job template dialog box ... 140
New Command template dialog box ... 140
All Datasets feature ... 141
Enabling the All Datasets feature ... 141
Collection control and History mode ... 141
Storage Toolkit ... 143
Storage Toolkit menus and dialog boxes ... 143
How Storage Toolkit processes action requests ... 145

Using the Storage Toolkit....146
Scenario for using the Storage Toolkit....146
Process for using actions in the Storage Toolkit....147
Reusing action requests....147
Action requests in workspaces that you create....148
About results for action requests....148
Checking for the completion of action requests....149
Storage Toolkit limitations and hints....149
Using substitution variables in the Storage Toolkit....154
Setting substitution variables for toolkit commands....154
Standard substitution variables in the toolkit....155
How commands run when substitution variables are defined....156
Examples of substitution variables....156
About the "Command Variable" (or "Variable") column....156
About the "Attribute or String" column....157
Standard tabs in dialog boxes of Storage Toolkit....157
Options tab....158
General tab....158
Scheduling tab....159
JCL tab....159
Command tab....159
Authentication and authorization for running Storage Toolkit commands....160
Data View dialog box....161
Custom actions in Storage Toolkit....161
Create Batch Job dialog box....162
Edit Dataset dialog box....165
Issue Command dialog box....165
Submit Command or Job dialog box....166
Volume actions in Storage Toolkit....167
Volume Backup dialog box....167
Volume Conversion dialog box....167
Volume Migration dialog box....168
Volume Space Release dialog box....168
Compress dialog box....168
Defrag dialog box....169
Dataset actions in Storage Toolkit....169
Dataset Backup dialog box....170
Dataset Attribute Group Extract dialog box....170
Dataset Migration dialog box....171
Dataset Copy and Move dialog box....171
Dataset Recall dialog box....172
Dataset Recover dialog box....172
Dataset Space Release dialog box....172
Dataset Compress dialog box....173
Dataset Allocate dialog box....173
Dataset Alter dialog box....174
Dataset Define Cluster dialog box....175
Dataset Catalog dialog box....179
Dataset Delete dialog box....179
Dataset Listcat dialog box....180
Dataset Rename dialog box....182
Dataset Repro dialog box....182
Dataset Uncatalog dialog box....184
Dataset Verify dialog box....184
Print dialog box....184
Cancel HSM Request dialog box....186
Accessing the Cancel HSM Request feature....186
Warning about cancellation of HSM requests....186
DFSMSrmm actions in the Storage Toolkit....187
Add Dataset dialog box....187
Change Dataset dialog box....187
Delete Dataset dialog box....188
Add Owner dialog box....188
Change Owner dialog box....188
Delete Owner dialog box....188
Add Volume dialog box....188
Change Volume dialog box....189
Confirm Actions for All Volumes dialog box....189
Delete Volume dialog box....189
Add VRS dialog box....189
Change VRS dialog box....190
Delete VRS dialog box....190
Situation Editor - Storage Toolkit Action Extension....190
Associate Storage Toolkit Command....190
Overview of common action commands....191
Managing Actions....192
Associating Storage Toolkit commands with situations....192

User actions in Storage Toolkit for Realtime Dataset Metrics only 193
Create Batch Job template dialog box 194
New Command template dialog box 194
Dataset Group feature 195
The Dataset Group feature versus the Dataset Attributes feature 195
Working with dataset groups 195
Dataset Group Properties dialog box 195
Dataset Group Data Collection dialog box 196
Adding a dataset group 196
Setting the collection interval 196
Dataset group masking rules 197
Using dataset groups 198
Obtaining an overview of space usage 199
Identifying datasets with excess space 199
Create an alert when response time exceeds specific thresholds 199
Identify VSAM datasets with excessive CA/CI splits 200
Cataloged and not-allocated datasets 201
Notify the Production Scheduling group when a critical dataset exceeds a specified number of extents 202
Automatically migrate infrequently used datasets 202
Dataset Attributes Database 204
Benefits 204
Types of Dataset Attribute collection 204
Full collection 204
Partial Sweep collection 204
Incremental collection 204
Optimization 205
Attribute derivation and definition 205
Use cases 206
Configuring the Dataset Attribute Database feature 207
Optimizing data collection for the Dataset Attributes Database feature 209
Command-line option 209
Custom queries 210
Working with dataset attribute groups 211
Adding a dataset attribute group 211
Editing a dataset attribute group 212
Deleting a dataset attribute group 212
Additional dataset attribute group considerations 212
How properties and attributes are matched 212
Dataset name masking rules 213
Dataset attribute groups and the Dataset Attribute Database 213
Optimizing dataset attribute groups 213
Dataset Attributes Group Details dataset limit 213
Dataset attribute groups VSAM criteria 214
Application Launch and TotalStorage Productivity Center 216
IBM TotalStorage Productivity Center 216
OMEGAMON for Storage launch feature 216
Launching the IBM TotalStorage Productivity Center 216
Workspaces that provide the launch feature 217
Configuring the launch feature 217
DFSMSrmm Status feature 219
Benefits 219
Optimization 219
Attribute derivation and definition 220
Scenarios for using the DFSMSrmm feature 221
Optimizing data collection for the DFSMSrmm Status feature 221
Command-line option 221
Custom queries 222
Dialog boxes 223
Dialog boxes 223
RMM Collection Control dialog box 223
CA 1 Collection Control dialog box 224
Copy Group dialog box 226
Copy Services Collection Control dialog box 227
RLS Collection Control dialog box 228
ADABAS Collection Control dialog box 228
Dataset Group Properties dialog box 229
Dataset Group Data Collection dialog box 229
Attributes 230

List of attribute groups and attributes 230
ADABAS Databases 233
ADABAS Files 234
ADABAS Volumes 236
Application Datasets 236
Application Dataset Details 236
Application Dataset Monitor 237
Application Monitoring 237
Application RLS Datasets 238
Application Volume Performance 238
Application Volume Space 240
Cache Control Unit 241
Cache Devices 242
Cache Raid Rank 245
Catalog Dataset Attributes 245
CA1 Configuration 246
CA1 Controls 246
CA1 Dataset 247
CA1 Job Summary 248
CA1 Permanent Error Summary 248
CA1 Program Summary 248
CA1 Summary 249
CA1 Temporary Error Summary 250
CA1 Volume 250
Channel Path 253
Cross System Volume Summary 255
Cross System Volumes 256
DASD Group Vol Perform 257
DASD Group Vol Space 258
DASD Physical Group 260
DASD Volume Performance 262
DASD Volume Space 264
Dataset Attributes Blocksize Summary 265
Dataset Attributes CA Split Summary 266
Dataset Attributes Catalog Summary 266
Dataset Attributes CI Split Summary 267
Dataset Attributes Creation Date Summary 267
Dataset Attributes Defs 268
Dataset Attributes Dataset Detail 268
Dataset Attributes DSORG Detail 272
Dataset Attributes DSORG Summary 273
Dataset Attributes Extent Summary 274
Dataset Attributes High Level Qualifier Summary 274
Dataset Attributes Inefficient Blocksize Summary 275
Dataset Attributes Mature Dataset Summary 276
Dataset Attributes Never Referenced Dataset Summary 276
Dataset Attributes Newborn Dataset Summary 277
Dataset Attributes Reference Date Summary 277
Dataset Attributes SMS Class Detail 278
Dataset Attributes SMS Class Summary 279
Dataset Attributes Space Allocated Summary 279
Dataset Attributes Space Unused Summary 280
Dataset Attributes System Summary 280
Dataset Metrics 282
Dataset Attributes Uncataloged Detail 282
Dataset Attributes Uncataloged Summary 283
DSN Attr Group Detail 283
DSN Attr Group Summary 286
Dataset Group Defs 288
Dataset Group Details 288
Dataset Group Detail Attrs 291
Dataset Group Details Volume 291
Dataset Group Summary 293
Dataset Performance 294

Dataset SMS....295
Dataset Space....295
Dataset Users....297
FICON Director Control Units....297
FICON Director Ports....297
FICON Director Summary....298
S3 App Defs....298
S3 Dag Defs....299
S3 DASD Defs....299
S3 Group Defs....300
S3 Device Class Details....300
S3 Device Summary....301
S3 Storage Group Configuration....301
S3 User Storage Groups....301
S3 Volume Event....303
S3 Volume Model....303
S3 Volume Pool....303
Hitachi Data Systems 2107 Array Configuration....304
Hitachi Data Systems 2107 Configuration....304
Hitachi Data Systems 2107 Parity Group....305
HSM CDS....305
HSM Common Storage....306
HSM Cross System CRQ Hosts....306
HSM Cross System CRQplex....307
HSM CRQ Hosts....308
HSM CRQ Requests....308
HSM CRQplex....309
HSM Dataset Recover Version....310
HSM Function Details....310
HSM Function Statistics....311
HSM Function Summary....312
HSM Host Function Status....312
HSM Host Request Summary....313
HSM Host Status....314
HSM Private Storage....315
HSM Requests....315
HSM Status....316
Logical Control Unit....317
Logical Control Unit Channel Paths....318
LCU Control Unit....319
RLS Appl Dataset Details....319
RLS Buffer LSU Summary....322
RLS Buffer Pool Detail....323
RLS Collection Interval....323
RLS Dataset Details....324
Dataset Collection Interval....326
RLS Dataset Group Details....326
RLS Dataset Group Summary....328
RLS Lock Structure....329
RLS Lock Structure CF Detail....331
RLS Performance Overview....331
RLS Storage Class....332
RMM Configuration....334
RMM Control....336
RMM Control Dataset....336
RMM Dataset Details....337
RMM Location Details....338
RMM Owner Details....339
RMM Pending Actions....339
RMM System Summary....339
RMM Volume Details....342
RMM VRS Details....347
S3 PPRC Connection....348
S3 PPRC Volume....349

S3 RLS Cache Structure Summary 349
S3 XRC Session 350
S3 XRC Volume Pair 351
S3 Flash Copy Relation 351
S3 Flash Copy Volume 352
S3 Link Statistics 353
S3 VTS Cluster Worst Statistics Summary 355
SMS Cache Sets 356
SMS Configuration 356
SMS Data Class 356
SMS Management Class 357
SMS Storage Class 359
SMS Storage Group 359
SMS Storage Group Status 361
SMS Systems 361
SMSplex Navigation 361
SMSplex Summary 362
Storage Toolkit Requests 363
Storage Toolkit Result Summary 364
Storage Toolkit Result Detail 365
Storage Toolkit Result Resources 366
Symmetrix Configuration 366
Symmetrix Devices 366
Symmetrix Director Summary 367
Symmetrix Directors 367
Symmetrix SSIDs 369
System Summary 369
Tape Device 371
Tape Group 371
TotalStorageDS Array Configuration 372
TotalStorageDS Configuration 373
TotalStorageDS Extent Pool 373
TotalStorageDS Rank 374
TotalStorageDS SSIDs 375
UDG Cache Devices 375
Volume Group Summary 377
Volume Status 380
Volume Users 380
VTS Cache 381
VTS Capacity 382
VTS Composite Library 383
VTS Overview 384
VTS Physical Devices 384
VTS TS7700-series Cache Container 385
VTS TS7700-series Cache Partition Container 386
VTS TS7700-series Cache Preference Group 387
VTS TS7700-series Cluster 388
VTS TS7700-series Cluster Replication 389
VTS Virtual Devices 389
VTS Virtual Devices Summary 390
VTSS Channels 391
VTSS Configuration 391
VTSS RTD Queue 392
VTSS RTD Status 392
VTSS Status 393
VTSS Virtual Volumes 393
Situations 395
Overview of situations 395
Predefined situations descriptions 396
KS3_Applic_Resp_Time_Critical 396
KS3_Applic_Resp_Time_Warning 396
KS3_Cachecu_Cache_Stat_Critical 396
KS3_Cachecu_DFW_Retry_Critical 396
KS3_Cachecu_DFW_Retry_Warning 396
KS3_Cachecu_Inact_Vols_Critical 396

KS3_Cachecu_Inact_Vols_Warning 397
KS3_Cachecu_NVS_Stat_Critical 397
KS3_Cachecu_Read_HitP_Critical 397
KS3_Cachecu_Read_HitP_Warning 397
KS3_Cachecu_Trk_Dstg_Critical 397
KS3_Cachecu_Trk_Dstg_Warning 397
KS3_Cachecu_Trk_Stag_Critical 397
KS3_Cachecu_Trk_Stag_Warning 398
KS3_Cachecu_Write_HitP_Critical 398
KS3_Cachecu_Write_HitP_Warning 398
KS3_Channel_Busy_Pct_Critical 398
KS3_Channel_Busy_Pct_Warning 398
KS3_FICON_Port_Frame_Pacing_W 398
KS3_FICON_Switch_Frame_Pacing_W 398
KS3_HSM_Backup_Held_Critical 399
KS3_HSM_Backup_Queue_Critical 399
KS3_HSM_Backup_Queue_Warning 399
KS3_HSM_CRQ_Element_Full_Warn 399
KS3_HSM_CRQ_Entry_Full_Warning 399
KS3_HSM_CRQ_Host_Critical 399
KS3_HSM_CRQ_Host_Disconn_Crit 400
KS3_HSM_CRQ_Host_Held_Critical 400
KS3_HSM_CRQ_Host_Place_Crit 400
KS3_HSM_CRQ_Host_Recall_Crit 400
KS3_HSM_CRQ_Host_Select_Crit 400
KS3_HSM_Dump_Held_Critical 400
KS3_HSM_Dump_Queue_Critical 401
KS3_HSM_Dump_Queue_Warning 401
KS3_HSM_Inactive_Host_Warning 401
KS3_HSM_Migrate_Held_Critical 401
KS3_HSM_Migrate_Queue_Critical 401
KS3_HSM_Migrate_Queue_Warning 401
KS3_HSM_Recall_Held_Critical 402
KS3_HSM_Recall_Queue_Critical 402
KS3_HSM_Recall_Queue_Warning 402
KS3_HSM_Recovery_Held_Critical 402
KS3_HSM_Recovery_Queue_Critical 402
KS3_HSM_Recovery_Queue_Warning 402
KS3_HSM_Status_Inactive_Crit 403
KS3_LCU_Av_Delay_Q_Critical 403
KS3_LCU_Av_Delay_Q_Warning 403
KS3_LCU_Cont_Rate_Critical 403
KS3_LCU_Cont_Rate_Warning 403
KS3_LCU_IO_Rate_Sec_Critical 403
KS3_LCU_IO_Rate_Sec_Warning 403
KS3_RLS_Accelerated_Mode_Warn 404
KS3_RLS_Dataset_Avg_Resp_Time_C 404
KS3_RLS_Dataset_Avg_Resp_Time_W 404
KS3_RLS_DSG_Max_Avg_Resp_Time_C 404
KS3_RLS_DSG_Max_Avg_Resp_Time_W 404
KS3_RLS_Panic_Mode_Critical 404
KS3_RLS_Panic_Mode_Warning 404
KS3_RLS_StorCls_Avg_Resp_Time_C 404
KS3_RLS_StorCls_Avg_Resp_Time_W 405
KS3_RMM_CDS_Backup_Critical 405
KS3_RMM_CDS_Backup_Warning 405
KS3_RMM_CDS_Space_Critical 405
KS3_RMM_CDS_Space_Warning 405
KS3_RMM_Exit_Status_Critical 405
KS3_RMM_Journal_Status_Critical 405
KS3_RMM_Operating_Mode_Warning 405
KS3_RMM_Scratch_Tape_Critical 406
KS3_RMM_Scratch_Tape_Warning 406
KS3_RMM_Inactive_Critical 406
KS3_Stg_Toolkit_Result_Critical 406
KS3_Stg_Toolkit_Result_Warning 406
KS3_Storage_Gr_Pct_Free_Crit 406
KS3_Storage_Gr_Pct_Free_Warning 407
KS3_TDS_Array_Degraded_Crit 407
KS3_TDS_Array_Prob_Crit 407
KS3_TDS_Array_RPM_Crit 407
KS3_TDS_Array_Throttled_Crit 407
KS3_TDS_ExtPool_Array_Prob_Crit 407
KS3_TDS_Rank_Array_Prob_Crit 407
KS3_Vol_Cache_DFW_Retry_Critical 408
KS3_Vol_Cache_DFW_Retry_Warning 408
KS3_Vol_Cache_Read_HitP_Critical 408
KS3_Vol_Cache_Read_HitP_Warning 408
KS3_Vol_Cache_Writ_HitP_Critical 408
KS3_Vol_Cache_Writ_HitP_Warning 409

KS3_Vol_Disabled_VTOC_Critical 409
KS3_Vol_EAV_Fragment_Index_Crit 409
KS3_Vol_EAV_Fragment_Index_Warn 409
KS3_Vol_EAV_Free_Space_Pct_Crit 409
KS3_Vol_EAV_Free_Space_Pct_Warn 409
KS3_Vol_Fragment_Index_Critical 410
KS3_Vol_Fragment_Index_Warning 410
KS3_Vol_Free_Space_Pct_Critical 410
KS3_Vol_Free_Space_Pct_Warning 410
KS3_Vol_Perf_Resp_Time_Critical 410
KS3_Vol_Perf_Resp_Time_Warning 410
KS3_VTS_Disconnect_Time_Crit 411
KS3_VTS_Host_GB_Warning 411
KS3_VTS_Pct_Copy_Throt_Warn 411
KS3_VTS_Pct_Wr_Over_Throt_Warn 411
KS3_VTS_Recall_Pct_Warning 411
KS3_VTS_Virt_MtPend_Av_Warning 411
KS3_VTS_Virt_MtPend_Mx_Warning 411
Take action commands 412
Take Action commands 412
Accessibility 413
Interface information 413
Related accessibility information 413
IBM® and accessibility 413
Notices 414
Trademarks 415

User's guide

The document is intended for anyone who uses the features and functionality of IBM® Z OMEGAMON® AI for Storage®.

Workspaces

In this section, you can find information about predefined workspaces of OMEGAMON AI for Storage.

When using Tivoli® Enterprise Portal, information is displayed in workspaces. Within a given workspace, information is displayed in tabular form. Tivoli® Enterprise Portal refers to this tabular format for information as a table view.

Each workspace contains a table view that provides information you can use to monitor a particular resource in your system. Column headings in the table view correspond to attributes that you can use to create situations. Each of the attributes in a specific table view belongs to a single attribute group.

Information can also be displayed in the workspace as a chart, graph, or other format you can specify.

Information can be presented to you in any of the following views:

- Table view
- Pie chart view
- Bar chart view
- Plot chart view
- Needle gauge view
- Thermometer gauge view
- Notepad view
- Situation Event console view, which shows the status of the situations associated with the system.
- Take Action view, which is used to send a command to the system.
- Terminal view, which enables you to start a 3270 or 5250 work session.
- Browser view, which permits you to open a browser to see HTML pages and Web sites.

Use the Export function of the Tivoli® Enterprise Portal to store product data as text or as comma separated values (CSV format). For example, you can right-click on a table view and select **Export** to export the data in the view.

A workspace is the working area of the Tivoli® Enterprise Portal application window. At the left of the workspace is a Navigator that permits you to select the workspace you want to display. As part of the application window, the right side of the status bar shows the Tivoli® Enterprise Portal server name to which the displayed information applies, as well as the ID of the current user. Some key points about workspaces are as follows:

- As you select items in the Navigator, the workspace presents views pertinent to your selection. Each workspace has at least one view.
- Every workspace has a set of properties associated with it. You can customize the workspace by working in the Properties editor to change the style and content of each view.
- Another way to customize the workspace is to change the type of view or to add views to the workspace.
- Always save your changes before moving to a new workspace. Otherwise, the changes you make to the workspace are lost.
- Some workspaces, such as those dedicated to tracking the metrics of a DS8000® storage device, are not available in selection menus and lists when DS8000 storage devices do not exist in the monitoring environment.

This document has a section for each of the nodes that you see in the Navigator. These nodes are listed in the following section. Each node links to a workspace that has the same name. Each workspace has links to related workspaces. Each link goes to a section that describes the corresponding workspace. The sections also describe all linked workspaces.

Nodes in the Navigator that are linked to workspaces

> **Note:** In the Navigator and in this document, nodes are listed in functional order, not alphabetically.

Enterprise
z/OS Systems
System_name
Storage Subsystem
SYSTEM:SUBSYSTEM:AGENT_NAME
Application Summary
Channel Path
Cache CU Performance
Cache CU Status
Logical Control Unit
Tape Group
Virtual Tape Subsystems
SMS Storage Groups Performance
SMS Storage Groups Space
User DASD Groups Performance
User DASD Groups Space
User Groups for Storage Groups
DFSMShsm™ Status
DFSMSrmm™ Status
Dataset Attributes System Summary
Dataset Group Summary
SMS Configuration
Storage Toolkit Action Requests
Monitoring Copy Services
Monitoring ADABAS Database Details

This document does not attempt to describe all details for a workspace. For details regarding the column headings (which correspond to attribute names) in a given table view, see "Attributes" on page 230. For additional details regarding any workspace, go to the workspace itself and inspect the tables, views, and other features.

Accessing workspaces

This section describes the security features and tells how to access workspaces in the Tivoli® Enterprise Portal.

Security features for user accounts

Security features such as the following regulate access to workspaces or access to the functionality in workspaces:

- The typical user of the Tivoli® Enterprise Portal has a standard level of privileges to view and modify the portal. For example, the **sysadmin** default user can see the views and modify the portal. Be aware that if you use a user account that has the **Workspace administration** level of privileges, the views that you see differ, as in these examples:
  - The title and status bars display `*ADMIN MODE*` after the user name to indicate that the user ID is in workspace administration mode.
  - Comprehensive lists of workspace links are displayed in some pop-up menus, instead of the standard menus that are nested within specific workspaces.

  The **Workspace administration** mode has additional privileges and purposes that are described in the *IBM® Tivoli® Monitoring: User's Guide*.
- The Standard tabs in the dialog boxes of the Storage Toolkit section describes the authorization that is required for the action requests of the Storage Toolkit.

Accessing IBM® Z OMEGAMON® AI for Storage® in the portal

Access IBM® Z OMEGAMON® AI for Storage® in the portal through the Navigator, as in the following example.

1. Click the plus sign (+) beside the **z/OS® Systems** icon to see available systems.

2. Click the plus sign (+) beside the name of the name of the z/OS® system for which you want information, for example, **SYSG**. The available subsystems are displayed.
3. Click the plus sign (+) beside the **Storage Subsystem** icon.
4. Click the plus sign (+) beside the **STORAGE** managed system for which you want information, for example, **VTCTH@@L:SYSG:STORAGE**. A list of product nodes is displayed.
5. Click the node of a workspace, for example, **Application Summary**. The portal displays the Application Summary workspace.

The Nodes in the Navigator that are linked to workspace section provides information on the product nodes within the hierarchy of a typical Navigator.

Accessing the workspace that corresponds to a node in the Navigator

To access a workspace that corresponds to a node in the Navigator, click the node. The workspace is displayed in the portal.

Accessing workspace links from a Navigator node

1. Click the node. The node is activated and the corresponding workspace is displayed in the portal.
2. Right-click the node to display the pop-up menu.
3. Select the workspace you want in the **Workspace** submenu.

> **Note:** Some workspace links that are mentioned in this document are not always visible. The Tivoli® Enterprise Portal displays some workspace links only when the target workspaces are relevant in the monitoring environment. For example, if you do not monitor DS8000® storage facilities, a link is not available for a workspace that tracks attributes of DS8000® storage facilities.

Accessing workspace links from views in a workspace

1. You can access related workspaces from table views and bar charts as follows:
   - **Table views:** Right-click on the link icon ( ) located at the left of a row in a table view.
   - **Bar charts:** Right-click a bar in a bar chart view and select **Link To** in the menu to access the list of links.
2. Select the name of a workspace in the pop-up menu.

Accessing cross-product workspace links

You access cross-product workspace links like any other workspace link. You use this type of link to go to workspaces in other Tivoli® products as described in the Integration with other products section in this topic.

Standard path and method for navigating workspaces

This document presents the hierarchy of linked workspaces for each workspace node in IBM® Z OMEGAMON® AI for Storage®. For example, the list in "Monitoring applications" on page 21 shows the hierarchy for the **Application Summary** workspace node in the Navigator of the portal. The bulleted list illustrates the standard path for navigating the workspaces. The standard method of navigating workspaces is as follows:

1. Right-click the link icon ( ) that is located at the left side of each row in a table view. (Most workspaces have at least one table view.)
2. Select a new workspace to go to in the pop-up menu.

> **Note:** Some link names in the pop-up menu do not exactly match the name that you see in the title bar of the target workspace.

Keep the following tips in mind as you a navigate workspaces:

- When no link icon is displayed in the table view of a workspace, you have reached the end of the standard flow of navigation for a set of workspaces.
- Most workspace nodes in IBM® Z OMEGAMON® AI for Storage® have multiple subpaths.
- Some subpaths terminate in the navigation tree structure of a different, but related, workspace node of IBM® Z OMEGAMON® AI for Storage®. For example, many nodes ultimately link to the workspaces of the SMS Storage Groups Performance node, where relevant performance data is available.
- The **Back** button in the browser client or the **Back** arrow icon ( ) in the Desktop client are essential navigation tools as you traverse multiple subpaths. You move backward to leave one subpath and enter another. Alternatively, you can click the workspace node in the Navigator to return to the first workspace of the node in which you are navigating.

Historical navigation

Several workspaces include historical navigation links. Some of these workspaces also include the option of using the Time Span tool to change a real-time workspace to a historical view. When a workspace has a historical navigation link, use that link to view historical data rather than using the Time Span tool that is visible in each view of the workspace.

By using the historical navigation feature, the workspace view is changed from a real-time view to a trend view of a single object in the real-time view. From the trend-view, you can navigate to a historical point-in-time view of the items within the object. For instance, rather than using the Time Span tool, use the following navigation path:

SMS Storage Groups Performance workspace > Storage Group Performance Trend

The SMS Storage Groups Performance workspace provides real-time data for all of the defined SMS performance groups. The Storage Group Performance Trend workspace provides trending information for a single storage group.

Storage Group Performance Trend > SMS Storage Group Historical Volume Performance

The SMS Storage Group Historical Volume Performance workspace shows the performance data for all of the volumes in the storage group at the point in time that is selected from the row of data in the Storage Group Performance Trend workspace.

Similarly, navigation is available from the SMS Storage Group Historical Volume Performance workspace to the Dataset Historical Performance Summary workspace, where a point-in-time view of the dataset activity for the selected volume is displayed. Similar navigation capability is provided for user DASD groups, channels, cache CUs, logical control units (LCUs), and TotalStorage facilities.

When viewing the point-in-time historical workspaces, the Time Span tool is visible on each view of the workspace. The Time Span tool does not function correctly on these workspaces. These workspaces have been constructed to display a point-in-time historical view that cannot be altered by using the Time Span tool.

For best results, view the Tivoli® Data Warehouse summarization data from an object trend workspace. By using the preceding example, the summarization display is most useful when invoked from the Storage Group Performance Trend workspace. If you want a volume performance summarization, access the Volume Performance Trend workspace. The following trend workspaces are provided by the product and supported by Tivoli® Data Warehouse summarization. The Time Span tool functions correctly in these workspaces:

- Cache CU DFW Retry Trend
- Cache CU Destaging Trend
- Cache CU Performance Trend
- Cache CU Read Hit Percent Trend
- Cache CU Staging Trend
- Cache CU Status Trend
- Cache CU Write Hit Percent Trend
- Channel Path Trend
- Logical Control Unit Trend
- SMS Storage Group Performance Trend

- SMS Storage Group Space Trend
- Storage Group Percent Free Trend
- User DASD Group Performance Trend
- User DASD Group Space Trend
- Volume Cache Trend
- Volume Destaging Trend
- Volume Performance Trend
- Volume Space Trend
- Volume Staging Trend

The following point-in-time historical workspaces are included with the product. The Time Span tool has unsupported results. Do not use the Time Span tool on the following workspaces:

- Cache CU Volume DFW Retry History
- Cache CU Volume Read Hit Percent History
- Cache CU Volume Write Hit Percent History
- User DASD Group Volume Historical Cache
- User DASD SMS Storage Group Historical Volume Performance
- User DASD Group Volume Historical Space
- Volume Cache History
- Volume Destaging History
- Volume Group Space History
- Volume Staging History

Integration with other products

IBM® Tivoli® Monitoring provides the dynamic workspace linking feature to enable you to quickly access workspaces in multiple products. IBM® Z OMEGAMON® AI for Storage® uses that feature to provide predefined links that go to storage-related workspaces in Tivoli® OMEGAMON® on z/OS® and in HSM Advanced Reporter. For example, the Application Summary workspace provides a predefined link to the Address Space Bottleneck Detail workspace in Tivoli® OMEGAMON® on z/OS®.

> **Note:** You can also modify the links that are provided by the product by default. See *Tivoli Enterprise™ Portal Help* for information on the Link Wizard that helps you modify the default linkages provided in the popup menus of rows in a workspace.

> **Note:** Dynamic linking is available only for instances of OMEGAMON® products that report to the same Tivoli® Enterprise Monitoring Server as IBM® Z OMEGAMON® AI for Storage® and that run in the same logical partition.

Links to workspaces of related products

This section lists the links that are available from the workspaces of IBM® Z OMEGAMON® AI for Storage® to the workspaces of related products. The Accessing workspaces section provides general information about navigating workspaces.

Links to workspaces of related products		
IBM® Z OMEGAMON® AI for Storage® linked product	**Workspace link location in IBM® Z OMEGAMON® AI for Storage®**	**Name of the link in pop-up menu**
IBM® Tivoli® Advanced Audit for DFSMShsm™	DFSMShsm™ Status	Backup Control Dataset Error Summary
		Migration Control Dataset Error Summary
		Offline Control Dataset Error Summary
IBM® Tivoli® Advanced Reporter for DFSMShsm™	DFSMShsm™ Status	BCDS Volume Summary
		BCDS Details
		Daily Space Management Detail
		Daily Space Management Summary
		DFSMShsm™ Return Code Summary
		Extent Reduction Summary
		Log Migration Detail
		Log Recall Detail
		MCDS Details
		MCDS Volume Summary
	SMS Configuration	BCDS Details
		MCDS Details
IBM® Tivoli® Advanced Catalog Management for z/OS®	Catalog Summary	Catalog Dataset Attributes
		Catalog Details Link
		Last Catalog Backup
	Dataset Group Details Primary	Catalog Dataset Attributes
		Catalog Details
		Last Catalog Backup
	Dataset Group Details VSAM	Catalog Dataset Attributes
		Catalog Details
		Last Catalog Backup
	Dataset Group VSAM Attributes	Catalog Dataset Attributes
		Catalog Details
		Last Catalog Backup
IBM® Tivoli® OMEGAMON® on z/OS®	Application Summary	Address Space CPU Usage Details
		Address Space CPU Usage Enclaves
		Address Space Impact Analysis
		Address Space Storage for Job
		Owned Enclaves
	DFSMShsm™ Function Details	Address Space Bottleneck Details
		Address Space CPU Usage Details
		Address Space CPU Usage Enclaves
		Address Space Impact Analysis
		Address Space Storage for Job
		Owned Enclaves

IBM® Z OMEGAMON® AI for Storage® linked product	Workspace link location in IBM® Z OMEGAMON® AI for Storage®	Name of the link in pop-up menu
	DFSMShsm™ Host Details	Address Space Bottleneck Details
		Address Space CPU Usage Details
		Address Space CPU Usage Enclaves
		Address Space Impact Analysis
		Address Space Storage for Job
		Owned Enclaves
	DFSMShsm™ Status	Address Space Bottleneck Details
		Address Space CPU Usage Details
		Address Space CPU Usage Enclaves
		Address Space Impact Analysis
		Address Space Storage for Job
		Owned Enclaves
	DFSMShsm™ Storage Usage	Address Space Storage

About dynamic linking to cross-product workspaces

Dynamic workspace linking (sometimes referred to as DWL) permits you to easily navigate between workspaces that are provided by multiple products. This feature aids problem determination and improves integration across the monitoring products, permitting you to quickly determine the root cause of a problem. Predefined cross-product links provided by the OMEGAMON® products permits 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, the list of links is displayed. This list might contain links to workspaces provided by one or more monitoring products. The product you are linking to must be installed and configured.

> **Note:** Your Tivoli Enterprise™ Portal user ID must be authorized to access the target product. Otherwise links to workspaces in the targeted product are not included in the list.

Choose a workspace from the list to navigate to that workspace. By linking to the target workspace in context, you receive additional information that is related to the system, subsdsystem, or resource you are currently viewing.

If you choose a workspace from the list and the target workspace is not available, you receive message KFWITM081E. Refer to the *Problem Determination Guide* for more information.

Associating workspaces with attributes

There is a direct relationship between attributes and workspaces. An attribute group typically corresponds to a table view within a named workspace and attribute items correspond to columns in the table view.

Each Tivoli® Enterprise Portal workspace displays real-time information or historical data for many of the attributes. The information is available to you, independent of whether you are using IBM® Z OMEGAMON® AI for Storage® to monitor situations.

Using IBM Z OMEGAMON AI for Storage information

You can view information about each managed system that you are monitoring. Use this information to perform the following tasks:

- Monitor the performance of each managed system, helping you to identify bottlenecks and evaluate tuning decisions
- Select the most effective threshold values for situations you create
- Review status information when a change in the state of a given resource occurs, such as from OK to Warning or Critical

Defining workspace properties

Every workspace has a set of properties associated with it. You can customize the workspace by working in the Properties editor to change the style and content of each view. Changes you make to workspace properties, such as adding or editing a view, are only temporary. They are lost when you exit IBM® Z OMEGAMON® AI for Storage® unless you save the workspace.

The properties of a workspace can include some or all of the following:

- Query: Specify what data must be specified in the chart or table
- Filters: Refine the view by filtering out unwanted data from the chart or table
- Thresholds: Establish threshold values and color indicators for a table view
- Configuration: Specify the script to run or the connection to make whenever you open the terminal view
- Style: Change the behavior and appearance of the view

Investigating a situation event

When the conditions of a situation have been met, the situation evaluates to True, causing a situation event indicator to be displayed in the Navigator. You can investigate the situation event by opening its workspace.

The situation event workspace shows two table views, one with the values of the attributes when the situation evaluated to True, and the other with the attributes' current values.

The situation event workspace can also display a text view with any expert advice written by the author of the situation, and the Take Action view so you can send a command to the application started on that system.

Filtering information

To manually set up filtering for a given table view, place the cursor on the table view, press the right mouse button, and select **Properties**. (The equivalent keystroke is **Ctrl+R**.) From the displayed dialog box, select the **Filters** tab. Here you can select the columns to display as well as set up the criteria for which rows to display. To save your filtering specifications, you must save the workspace before exiting.

Sorting information

Sorting is handled by simply clicking on a column heading. Click once and the report is sorted in ascending order. Click a second time to resort the report into descending order. A third click returns you to the report's default sort order.

Adding a workspace to your favorites

When using IBM® Z OMEGAMON® AI for Storage® in browser mode, you can start it from any workstation by entering the URL of the Tivoli® Enterprise Portal server where the browser mode client is installed. Each Tivoli® Enterprise Portal workspace also has a URL, so that you can save the workspace to your Favorites list or specify it as your home page.

Monitoring applications

You can use some of the workspaces to monitor the status of applications.

You access these workspaces through the Application Summary node in the Tivoli® Enterprise Portal. With IBM® Z OMEGAMON® AI for Storage®, you can manage application definitions within the Manage Application Definition dialog box of the Tivoli® Enterprise Portal. This topic also describes the dialog boxes that you use to manage, add, edit, and delete application definitions.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

Navigating the workspaces of the Application Summary node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the Application Summary node to navigate to the workspaces that you see in the following hierarchy. You can access the workspaces in this table using the standard path and method for navigating workspaces.

Application Summary node: bulleted list showing the hierarchy of linked workspaces

- Application Dataset Performance
  - Dataset Details
- Application Volume Performance
  - Application Volume Dataset Performance
  - Physical Volume Group Performance
  - Device Performance Details
  - Cross System Volume Performance
  - Volume Performance Trend
    - Dataset Performance Summary History
      - Dataset Performance Historical Detail
- Application Volume Cache
  - Application Volume Dataset Performance
  - Physical Volume Group Cache
  - Device Performance Details
  - Cross System Volume Performance
  - Volume Cache Trend
- Application Dataset Space Usage
  - Dataset Details
- Application Volume Space Usage
  - Application Dataset Volume Space Utilization
  - Device Space Details
- Application RLS Datasets
  - Application RLS Dataset Performance

Application Summary workspace

This topic describes the main components of the principal workspace for this node.

The following sections describe various components of this workspace. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help, including in the *hover help* (a form of online help that can be viewed by moving a cursor over a GUI item such as an icon or a field). The direct situation analysis feature is available for this workspace. The feature can help you analyze and understand a problem that is displayed in this workspace. See "Directly monitoring situations" on page 133.

Views

The workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading. For information about related workspaces, see "Summary information for the workspaces in this node" on page 24.

Application table view

The table view shows one row of data for each active address space for which monitoring has been requested. The view contains the following information the following information:

- The name of the job or started task
- The decimal ASID number
- The I/O rate for the job
- The number of datasets allocated for the job
- The worst response time suffered by any of those datasets
- The name of the dataset suffering the worst response time and the volume on which it is located
- The number of volumes in use by the job and the worst response time on any of those volumes along with the name of the volume incurring that response time

Applications with Dataset MSR > 40 bar graph

This graph contains a bar for each monitored application that has a dataset with an average millisecond response time since the beginning of the RMF™ interval exceeded 40 milliseconds. The height of the bar indicates the worst dataset response time for the application.

Applications with volume MSR > 25 bar graph

This graph contains a bar for each monitored application that has a volume with an average millisecond response time in the previous RMF™ interval exceeding 25 milliseconds. The height of the bar indicates the worst volume response time for the application.

Defining applications to monitor

You define the list of monitored applications by setting up a situation. The situation can be named anything, however, it must be written against the Application Monitoring group. The attributes that must be set in the situation to control monitoring of the application include the attributes and values shown in the following table.

> **Note:** Monitored applications defined by setting up situations are not displayed in the Manage Application Definition dialog box. However, the applications are displayed in the application workspaces and continue to be monitored using those situations. IBM® recommends that you define applications to monitor using the Manage Application Definition dialog box.

Attribute	Value
Application	Name of the application for which monitoring is to be started or stopped. To specify a group of applications to monitor whose names all begin with a common string (for example, `CICS`*xxxx*), use the **Value of Expression** function in the Situation Editor. Type the application name with an asterisk as the final character, for example, `CICS*`. This specification enables monitoring for all applications that start with the string that precedes the asterisk.
Monitor Status	Y = Begin monitoring this application. N = Stop monitoring this application.
I/O Monitor Status	Y = Turn ON dataset-level I/O monitoring for this application. N = Turn OFF dataset-level I/O monitoring for this application.

To collect dataset performance statistics for all datasets accessed by an application, set the 'I/O Monitor Status' attribute to 'Y' when defining the situation that starts monitoring for the application. The following additional guidelines can help you customize data collection:

- Dataset performance statistics are gathered for all datasets accessed by the application regardless of the volume on which the datasets are located and the dataset performance options in effect for that volume.
- To discontinue the collection of dataset performance statistics for the application, you change the value of the 'I/O Monitor Status' attribute to 'N'. Depending on the dataset collection status of the volumes being accessed by the application, some dataset performance statistics continue to be collected.

- This workspace is one of two places that display statistics from the level of datasets. These statistics are available only when the Tivoli® Enterprise subsystem is running.

Workspaces that are linked to this workspace

See “Navigating the workspaces of the Application Summary node” on page 22.
Learn how to navigate among workspaces here: “Workspaces” on page 14.

See “Workspaces” on page 14 for a list of all nodes provided in IBM® Z OMEGAMON® AI for Storage®.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the Application Summary node.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

See “Application Summary workspace” on page 22 for more information on the primary workspace for this node.

Application Summary workspace		
View	**View type**	**Attribute group**
Application Summary Report	TABLE	Application Monitoring
Applications with Dataset MSR > 40	BAR	Application Monitoring
Applications with Volume MSR > 25	BAR	Application Monitoring

Application Dataset Performance workspace		
View	**View type**	**Attribute group**
Dataset Performance Report	TABLE	Dataset Performance
Datasets with MSR > Storage Class Objective	TABLE	Dataset Performance

Application Volume Performance workspace		
View	**View type**	**Attribute group**
Volume Performance Report	TABLE	Application Volume Performance
Volume Response Time > 25	BAR	Application Volume Performance

Application Volume Dataset Performance workspace		
View	**View type**	**Attribute group**
Dataset Performance Report	TABLE	Dataset Performance
Datasets with MSR > Storage Class Objective	TABLE	Dataset Performance

Application Volume Cache workspace		
View	**View type**	**Attribute group**
Volume Cache Report	TABLE	Application Volume Performance
Read Hit Percent < 55	BAR	Application Volume Performance

View	View type	Attribute group
Write Hit Percent < 30	BAR	Application Volume Performance

Application Dataset Space Usage workspace		
View	**View type**	**Attribute group**
Dataset Space Report	TABLE	Application Dataset Details
Volumes with Free Space % < 10	BAR	Application Volume Space

Application Volume Space Usage workspace		
View	**View type**	**Attribute group**
Volumes with Free Space % < 10	BAR	Application Volume Space
Volumes with Fragmentation Index > 650	BAR	Application Volume Space
Volume Space Report	TABLE	Application Volume Space

Application Dataset Volume Space Usage workspace		
View	**View type**	**Attribute group**
Dataset Space Report	TABLE	Application Dataset Details
Volume Space Allocation	PIE	DASD Volume Space

Application Dataset Response Time workspace		
View	**View type**	**Attribute group**
Dataset Performance Report	TABLE	Dataset Performance
Datasets with MSR > Storage Class Objective	TABLE	Dataset Performance

Application RLS Datasets workspace		
View	**View type**	**Attribute group**
Sysplex Level Average Response Time	BAR	RLS Application Dataset Details
Sysplex Level DIWA Lock Requests per Minute	BAR	RLS Application Dataset Details
Sysplex Level Application Dataset RLS Performance	TABLE	RLS Application Dataset Details

Application RLS Dataset Performance workspace		
View	**View type**	**Attribute group**
System Level Data Component Average Response Time	BAR	RLS Application Dataset Details
System Level Index Component Average Response Time	BAR	RLS Application Dataset Details
Sysplex Level RLS Dataset Performance	TABLE	RLS Application Dataset Details
System Level RLS Dataset Performance	TABLE	RLS Application Dataset Details

Managing application definitions

Use the **Manage Application Definition** dialog box to manage application definitions.

Application Definition

The unique name for the application definition you are adding. Names can consist of any combination of characters and must not exceed 30 characters in length. Names are case-sensitive.

Description

An optional description for the application definition. Descriptions can consist of any combination of characters and must not exceed 50 characters in length.

Persist

This check box indicates that you want the application monitoring definition to persist or be saved across recycles of the Tivoli® Enterprise Monitoring Server. Application monitoring definitions are discarded during a recycle of the monitoring server if the **Persist** field is not selected.

Application Mask

The application mask associated with the application definition. You can specify a partial name pattern and place the wildcard character (*) at the end. For example, a mask of CICS* includes all applications with names that begin with "CICS".

I/O Monitor Status

The I/O monitoring status from the drop-down list to indicate whether the I/O monitoring status should be started or stopped. This selection controls the collection of dataset level performance statistics for all datasets accessed by the application.

Adding an application definition

Use the **Add Application Definition** dialog box to add a new application definition.

To access this dialog box, right-click on a row of the Application Summary Report table view and select **Manage Application Definition** from the menu. If no applications are defined, the Add Application Definition dialog box displays. If one or more applications are defined in the Manage Application Definition list, right-click on any of the rows and select **Add Application Definition** from the menu. The Add Application Definition dialog box displays.

Application Definition

Enter a unique name for the application definition you are adding. Names may consist of any combination of characters and must not exceed 30 characters in length. Names are case-sensitive.

Description

Enter an optional description for the application definition. Descriptions may consist of any combination of characters and must not exceed 50 characters in length.

Application Mask

Enter an application mask. You can specify a partial name pattern and place the wildcard character (*) at the end. For example, a mask of CICS* includes all applications with names that begin with "CICS".

I/O Monitor Status

Select the I/O monitoring status from the drop-down list to indicate whether the I/O monitoring status should be started or stopped. This selection controls the collection of dataset level performance statistics for all datasets accessed by the application. By default, the I/O Monitor Status is "Stop".

Persist

Click this check box to indicate that you want the application monitoring definition to persist (or be saved) across recycles of the Tivoli Enterprise Monitoring Server (monitoring server or TEMS). Application monitoring definitions are discarded during a recycle of the monitoring server if the **Persist** field is not selected. By default, the application monitoring definition persists.

Editing an application definition

Use the **Edit Application Definition** dialog box to edit an application definition.

To access this dialog box, navigate to the Manage Application Definition dialog box, right-click on the table row for the application definition you want to edit, and select **Edit Application Definition** from the menu.

For descriptions of the fields in the dialog box, see "Adding an application definition" on page 26.

Deleting an application definition

Use the **Delete Application Definition** dialog box to confirm the deletion of the application definition.

You can delete an application definition by navigating to the Manage Application Definition dialog box, right-clicking on the table row for the application definition you want to delete, and selecting **Delete Application Definition** from the menu.

Click **Yes** to confirm the deletion of the application definition or click **No** to cancel the deletion.

Monitoring channel paths

You can use some workspaces to monitor the status of channel paths. You access these workspaces through the Channel Path node in the Tivoli® Enterprise Portal.

Navigating the workspaces of the Channel Path node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the Channel Path node to navigate to the workspaces that you see in the following hierarchy. You can access the workspaces in this table using the standard path and method for navigating workspaces.

Channel Path node: bulleted list showing the hierarchy of workspace links

- FICON® Director Summary
  - FICON® Director Ports
    - FICON® Director Control Units
- Channel Volume Performance (restricted to channels with hard disk drives attached)
  - Dataset Performance Summary
    - Dataset Performance Detail
      - Dataset Details
      - Dataset Performance Detail Trend
    - Dataset Details
    - Dataset Performance Summary Trend
      - Dataset Performance Historical Detail
  - Device Performance Details
  - Cross System Volume Performance
  - Volume Performance Trend
    - Dataset Historic Performance Summary
      - Dataset Performance Historical Detail
- Channel Volume Cache (restricted to channels with hard disk drives attached)
  - Dataset Performance Summary
    - (as above)
  - Device Performance Details
  - Cross System Volume Performance
  - Volume Cache Trend
- Channel Path Trend
  - Channel Path Historical Volume Performance
    - Dataset Historic Performance Summary

        - Dataset Performance Historical Detail
    - Channel Path Historical Volume Cache Performance
        - Dataset Historic Performance Summary
            - Dataset Performance Historical Detail

    The following workspaces display the volumes that have the least favorable readings for the attribute within the SSID selected for navigation:

- Highest Volumes Response Time (Top 15)
- Highest Volumes MPL (Top 15)
- Highest Volumes I/O Rate (Top 15)
- Highest Volumes Busy Percent (Top 15)
- Lowest Volumes Connect Percent (Top 15)

Channel Path workspace

This topic describes the main components of the principal workspace for this node.

The direct situation analysis feature is available for this workspace. The feature can help you analyze and understand a problem that is displayed in this workspace. See "Directly monitoring situations" on page 133.

Views

This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

Channel path table view

The table view displays one row of data for each channel path. The view contains the following information:

- The channel path ID
- Its status
- Whether or not it is a shared path
- Utilization percentage across the complex
- Utilization percentage within this LPAR
- The type of connection it is
- Whether or not it is DCM managed
- The BUS utilization percentage
- Various read and write metrics for both the complex and the reporting LPAR
- The number of SMS volumes, NON-SMS volumes and the total number of volumes configured in the LPAR
- The channel path management facility group number for the path

Complex % utilization bar graph

This graph contains a bar for each channel path that had any utilization since the start of the RMF™ interval. The height of the bar indicates the percent to which the path was utilized in the complex.

LPAR % utilization DASD bar graph

This graph contains a bar for each channel path for which there are any volumes configured in the reporting LPAR and for which there was any utilization within the reporting LPAR since the beginning of the RMF™ interval.

LPAR MB/Sec WRITE - DASD bar graph

This graph contains a bar for each channel path for which there are any volumes configured in the reporting LPAR and for which there were any write operations that were performed to DASD in the reporting LPAR since the beginning of the RMF™ interval.

LPAR MB/Sec READ - DASD bar graph

This graph contains a bar for each channel path for which there are any volumes configured in the reporting LPAR and for which there were any read operations to DASD that were performed in the reporting LPAR since the beginning of the RMF™ interval.

Workspaces that are linked to this workspace

See "Navigating the workspaces of the Channel Path node" on page 27.
Learn how to navigate to workspaces here: "Workspaces" on page 14.

See "Workspaces" on page 14 for a list of all nodes provided in IBM® Z OMEGAMON® AI for Storage®.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the Channel Path node.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

See "Channel Path workspace" on page 28 for more information on the primary workspace for this node.

Channel Path workspace		
View	**View type**	**Attribute group**
Complex % Utilization	BAR	Channel Path
LPAR % Utilization - DASD	BAR	Channel Path
LPAR MB/Sec Written - DASD	BAR	Channel Path
LPAR MB/Sec Read - DASD	BAR	Channel Path
Channel Path Report	TABLE	Channel Path

Channel Volume Performance workspace		
View	**View type**	**Attribute group**
Volume Response Time > 25	BAR	DASD Volume Performance
Volume Performance Report	TABLE	DASD Volume Performance

Channel Volume Cache workspace		
View	**View type**	**Attribute group**
Volume Cache Report	TABLE	Cache Devices
Read Hit Percent < 55	BAR	Cache Devices
Write Hit Percent < 30	BAR	Cache Devices

Channel Path Trend workspace		
View	**View type**	**Attribute group**
Channel Utilization	PLOT	Channel Path
Channel Path Trend	TABLE	Channel Path

Channel Volume Response Time workspace		
View	**View type**	**Attribute group**
Volume Performance Report	TABLE	DASD Volume Performance

View	View type	Attribute group
Volume Response Time > 25	BAR	DASD Volume Performance

Monitoring cache CU performance

This topic describes the workspaces that monitor the status of the performance of cache CU. You access these workspaces through the Cache CU Performance node in the Tivoli® Enterprise Portal.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

Navigating the workspaces of the Cache CU Performance node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the Cache CU Performance node to navigate to the workspaces. You can access the workspaces in this table using the standard path and method for navigating workspaces.

Cache CU Performance node: bulleted list showing the hierarchy of linked workspaces

- Cache CU Volume Cache
  - Dataset Performance Summary [1]
  - Physical Volume Group Cache
    - Dataset Performance Summary [1]
  - Device Performance Details
  - Cross System Volume Performance
  - Volume Cache Trend
- Cache CU Volume Performance
  - Dataset Performance Summary [1]
  - Physical Volume Group Performance
    - Dataset Performance Summary [1]
    - Volume Performance Trend
      - Dataset Performance Summary History
  - Device Performance Details
  - Cross System Volume Performance
  - Volume Performance Trend
    - Dataset Historic Performance Summary
      - Dataset Performance Historical Detail
- Symmetrix Configuration.

> **Note:** Obtaining EMC Symmetric data can only be done in conjunction with the applicable EMC hardware and software. The hardware required is an EMC Symmetrix at the appropriate

> microcode level. Also required are the EMC Symmetrix Control Facility runtime libraries available from EMC.

- Cache CU Raid Rank
  - Physical Volume Group Cache
    - Dataset Performance Summary [1]
  - Cache CU Raid Rank Trend
- Cache CU Performance Trend
  - Cache CU Volume Historical Performance
    - Dataset Performance Summary History
      - Dataset History Performance Detail
  - CCU Volume Cache History
    - Dataset Historic Performance Summary
      - Dataset Performance Historical Detail
  - Cache CU Raid Rank Historical Data
- Hitachi Data Systems 2107 Configuration
  - Hitachi Data Systems 2107 Parity Groups
    - Hitachi Data Systems 2107 Array Configuration
    - Hitachi Data Systems 2107 Parity Group Trend
  - Hitachi Data Systems 2107 Extent Pool Volumes
    - Hitachi Data Systems 2107 Volume Trend
  - Hitachi Data Systems 2107 Extent Pool Trend
    - Hitachi Data Systems 2107 Parity Group History
    - Hitachi Data Systems 2107 Extent Pool Volume History
- Hitachi Data Systems 2107 CU Volumes
  - Hitachi Data Systems 2107 Volume Trend
- TotalStorage Configuration
  - TotalStorage Ranks
    - TotalStorage Array Configuration
    - TotalStorage Ranks Trend
  - TotalStorage Extent Pool Volumes
    - TotalStorage Volume Trend
  - TotalStorage Extent Pool Trend
    - TotalStorage Rank History
    - TotalStorage Volume History
- TotalStorage CU Volumes
  - TotalStorage Volume Trend

The following workspaces display the volumes that have the least favorable readings for the attribute within the SSID selected for navigation:

- Lowest Volumes CU Read Hit Percent (Top 10) [2]
- Lowest Volumes CU Write Hit Percent (Top 10) [2]
- Lowest Volumes DFW Hit Percent (Top 10) [2]

- Highest Volumes Destaging Rate (Top 10) [2]
- Highest Volumes CU Staging Rate (Top 10) [2]

Note:

1. This group of workspaces is organized in the following nested pattern:
   - Dataset Performance Summary
     - Dataset Performance Detail
       - Dataset Details
       - Dataset Performance Detail Trend
     - Dataset Details
     - Dataset Performance Summary Trend
       - Dataset Performance Historical Detail
2. This workspace links to the following set of related workspaces:
   - Dataset Performance Summary (See the preceding note ([1]) for details.)
   - Physical Volume Group Performance
     - Dataset Performance Summary (See the preceding note ([1]) for details.)
     - Volume Performance Trend
       - Dataset Performance Summary History
   - Device Performance Details
   - Cross System Volume Performance
   - Volume Cache Trend
     - Dataset Historic Performance Summary
       - Dataset Performance Historical Detail

Access the following linked workspaces by selecting the Cache CU Performance node in the Navigator. Right-click the node. Select the workspace that you want to access in the pop-up menu.

These workspaces list the volumes with the least favorable readings for the attribute across the z/OS® image.

- Lowest Volume Read Hit Percent
- Lowest Volume Write Hit Percent
- Lowest Volume DFW Hit Percent
- Highest Volume Destaging Rate
- Highest Volume Staging Rate

Cache CU Performance workspace

This topic describes the main components of the principal workspace for this node.

The direct situation analysis feature is available for this workspace. The feature can help you analyze and understand a problem that is displayed in this workspace. See “Directly monitoring situations” on page 133.

Views

The workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading. For information about related workspaces, see “Summary information for the workspaces in this node” on page 33.

Cache performance table view

The table view displays one row of data for each cache control unit configured to the reporting LPAR. The view contains the following information:

- The subsystem ID of the control unit
- The control unit type
- The read and write hit percentages for the control unit
- The read I/O percent
- The bypass and inhibit percents
- The DASD Fast Write hit and retry percents
- The Cache Fast Write read and write percents
- The track destaging rate
- The track staging rate
- Storage facility series

Read Hit Percent bar graph

This graph contains a bar for each cache control unit that had a read hit percent greater than zero. The height of the bar indicates the read hit percentage.

Write Hit Percent bar graph

This graph contains a bar for each cache control unit that had a write hit percent greater than zero. The height of the bar indicates the write hit percentage.

DFW Hit Percent bar graph

This graph contains a bar for each cache control unit with a DASD Fast Write hit percent greater than zero. The height of the bar indicates the DASD Fast Write hit percentage.

DFW Retry Percent bar graph

This graph contains a bar for each cache control unit with a DASD Fast Write retry percent greater than zero. The height of the bar indicates the DASD Fast Write retry percentage.

Track Destaging Rate bar graph

This graph contains a bar for each cache control unit with a track destaging rate greater than zero. The height of the bar indicates the track destaging rate.

Workspaces that are linked to this workspace

See “Navigating the workspaces of the Cache CU Performance node” on page 30.
Learn how to navigate to workspaces here: “Workspaces” on page 14.

See “Workspaces” on page 14 for a list of all nodes provided in IBM® Z OMEGAMON® AI for Storage®.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the Cache CU Performance node.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

See “Cache CU Performance workspace” on page 32 for more information on the primary workspace for this node.

Cache CU Performance workspace		
View	**View type**	**Attribute group**
Cache Control Unit Performance Report	TABLE	Cache Control Unit
DFW Hit Percent	BAR	Cache Control Unit
DFW Retry Percent	BAR	Cache Control Unit
Track Destaging Rate	BAR	Cache Control Unit
Read Hit Percent	BAR	Cache Control Unit
Write Hit Percent	BAR	Cache Control Unit

Cache CU Volume Cache workspace		
View	**View type**	**Attribute group**
Read Hit Percent < 55	BAR	Cache Devices
Write Hit Percent < 30	BAR	Cache Devices
Volume Cache Report	TABLE	Cache Devices

Cache CU Volume Performance workspace		
View	**View type**	**Attribute group**
Volume Performance Report	TABLE	DASD Volume Performance
Volume Response Time > 25	BAR	DASD Volume Performance

Cache CU Raid Rank workspace		
View	**View type**	**Attribute group**
Raid Rank Report	TABLE	Cache Raid Rank
Read Response Time	BAR	Cache Raid Rank
Write Response Time	BAR	Cache Raid Rank

Symmetrix Configuration workspace		
View	**View type**	**Attribute group**
Subsystem IDs	TABLE	Symmetrix SSIDs
Symmetrix Configuration Report	TABLE	Symmetrix Configuration
Symmetrix Directors Report	TABLE	Symmetrix Directors

Symmetrix Disk Director Summary workspace		
View	**View type**	**Attribute group**
Logical Volumes Per Disk	BAR	Symmetrix Director Summary
Cache Efficiency	BAR	Symmetrix Directors

Symmetrix Devices workspace		
View	**View type**	**Attribute group**
Devices with Mirrors in Error	TABLE	Symmetrix Devices
Symmetrix Devices	TABLE	Symmetrix Devices

Cache CU Performance Trend workspace		
View	**View type**	**Attribute group**
Cache CU Performance	PLOT	Cache Control Unit
Cache Control Unit Performance Trend	TABLE	Cache Control Unit

Volume Cache Trend workspace		
View	**View type**	**Attribute group**
Volume Cache Trend	TABLE	Cache Devices
Cache Hit Percent	PLOT	Cache Devices

Volume Cache History workspace		
View	**View type**	**Attribute group**
Volume Cache Report	TABLE	Cache Devices
Read Hit Percent < 55	BAR	Cache Devices
Write Hit Percent < 30	BAR	Cache Devices

Cache CU Historic Raid Rank workspace		
View	**View type**	**Attribute group**
Raid Rank Report	TABLE	Cache Raid Rank
Read Response Time	BAR	Cache Raid Rank
Write Response Time	BAR	Cache Raid Rank

Cache CU Raid Rank Trend workspace		
View	**View type**	**Attribute group**
Raid Rank Report	TABLE	Cache Raid Rank
Response Time	PLOT	Cache Raid Rank

TotalStorage® Ranks workspace		
View	**View type**	**Attribute group**
TotalStorage® Ranks	TABLE	TotalStorageDS Rank
Operations Per Second	BAR	TotalStorageDS Rank
Average Response Time	BAR	TotalStorageDS Rank

TotalStorage® Array Configuration workspace		
View	**View type**	**Attribute group**
DDM Capacity	BAR	TotalStorageDS Array
TotalStorage® Arrays	TABLE	TotalStorageDS Array
DDM Speed	BAR	TotalStorageDS Array

TotalStorage® CU Volumes workspace		
View	**View type**	**Attribute group**
TotalStorage® CU Volumes	TABLE	Cache Devices
Physical I/O Rate	BAR	Cache Devices
Physical Response Time	BAR	Cache Devices

TotalStorage® Extent Pool Volumes workspace		
View	**View type**	**Attribute group**
TotalStorage® Extent Pool Volumes	TABLE	Cache Devices
Physical I/O Rate	BAR	Cache Devices
Average Physical Response Time	BAR	Cache Devices

TotalStorage® Extent Pool Trend workspace		
View	**View type**	**Attribute group**
TotalStorage® Extent Pool Trend	TABLE	TotalStorageDS Extent Pool
Response Time	PLOT	TotalStorageDS Extent Pool

TotalStorage® Ranks Trend workspace		
View	**View type**	**Attribute group**
TotalStorage® Ranks Trend	TABLE	TotalStorageDS Rank
Average Response Time	PLOT	TotalStorageDS Rank

TotalStorage® Volume Trend workspace		
View	**View type**	**Attribute group**
TotalStorage® Volume Trend	TABLE	Cache Devices
Physical Response Time	PLOT	Cache Devices

TotalStorage® Rank History workspace		
View	**View type**	**Attribute group**
TotalStorage® Ranks	TABLE	TotalStorageDS Rank
Operations per Second	BAR	TotalStorageDS Rank
Average Response Time	BAR	TotalStorageDS Rank

TotalStorage® Volume History workspace		
View	**View type**	**Attribute group**
TotalStorage® Volumes	TABLE	Cache Devices
Physical I/O Rate	BAR	Cache Devices
Physical Response Time	BAR	Cache Devices

TotalStorage® Configuration workspace		
View	**View type**	**Attribute group**
Average Read Response Time	BAR	TotalStorageDS Extent Pool
Average Write Response Time	BAR	TotalStorageDS Extent Pool
TotalStorage® Extent Pools	TABLE	TotalStorageDS Extent Pool
TotalStorage® SSIDs	TABLE	TotalStorageDS SSIDs
TotalStorage® Configuration	TABLE	TotalStorageDS Configuration

Cache CU Destaging Trend workspace		
View	**View type**	**Attribute group**
Cache Control Unit Performance Trend	TABLE	Cache Control Unit

View	View type	Attribute group
Cache CU Physical Activity	PLOT	Cache Control Unit

Cache CU Volume Destaging workspace		
View	**View type**	**Attribute group**
Physical Volume Activity	BAR	Cache Devices
Volume Cache Destaging Report	TABLE	Cache Devices

Volume Destaging Trend workspace		
View	**View type**	**Attribute group**
Volume Destaging Trend	TABLE	Cache Devices
Physical Volume Activity	PLOT	Cache Devices

Volume Destaging History workspace		
View	**View type**	**Attribute group**
Physical Volume Activity	BAR	Cache Devices
Volume Historical Destaging Report	TABLE	Cache Devices

Cache CU Staging Trend workspace		
View	**View type**	**Attribute group**
Cache CU Physical Activity	PLOT	Cache Control Unit
Cache Control Unit Performance Trend	TABLE	Cache Control Unit

Cache CU Volume Staging workspace		
View	**View type**	**Attribute group**
Volume Cache Staging Report	TABLE	Cache Devices
Physical Volume Activity	BAR	Cache Devices

Volume Staging Trend workspace		
View	**View type**	**Attribute group**
Volume Staging Trend	TABLE	Cache Devices
Physical Volume Activity	PLOT	Cache Devices

Volume Staging History workspace		
View	**View type**	**Attribute group**
Volume Historical Staging Report	TABLE	Cache Devices
Physical Volume Activity	BAR	Cache Devices

Cache CU Volume DFW Retry workspace		
View	**View type**	**Attribute group**
Volume Cache DFW Retry Report	TABLE	Cache Devices
DFW Retry %	BAR	Cache Devices
DFW Hit %	BAR	Cache Devices

Cache CU DFW Retry Trend workspace		
View	**View type**	**Attribute group**
DFW Retry %	PLOT	Cache Control Unit
Cache Control Unit DFW Retry Trend	TABLE	Cache Control Unit

Cache CU Volume DFW Retry History workspace		
View	**View type**	**Attribute group**
Volume Cache Report	TABLE	Cache Devices
DFW Retry %	BAR	Cache Devices
DFW Hit %	BAR	Cache Devices

Cache CU Volume Read Hit Percent workspace		
View	**View type**	**Attribute group**
Volume Cache Report	TABLE	Cache Devices
Read Hit Percent	BAR	Cache Devices
Write Hit Percent	BAR	Cache Devices

Cache CU Read Hit Percent Trend workspace		
View	**View type**	**Attribute group**
Cache CU Read Hit Percent	PLOT	Cache Control Unit
Cache Control Unit Performance Trend	TABLE	Cache Control Unit

Cache CU Volume Read Hit Percent History workspace		
View	**View type**	**Attribute group**
Volume Cache Report	TABLE	Cache Devices
Read Hit Percent	BAR	Cache Devices
Write Hit Percent	BAR	Cache Devices

Cache CU Volume Write Hit Percent workspace		
View	**View type**	**Attribute group**
Volume Cache Report	TABLE	Cache Devices
Write Hit Percent	BAR	Cache Devices
Read Hit Percent	BAR	Cache Devices

Cache CU Write Hit Percent Trend workspace		
View	**View type**	**Attribute group**
Cache CU Write Hit Percent	PLOT	Cache Control Unit
Cache Control Unit Performance Trend	TABLE	Cache Control Unit

Cache CU Volume Write Hit Percent History workspace		
View	**View type**	**Attribute group**
Volume Cache Report	TABLE	Cache Devices
Write Hit Percent	BAR	Cache Devices

View	View type	Attribute group
Read Hit Percent	BAR	Cache Devices

Cache CU Volume Inactive Status workspace		
View	**View type**	**Attribute group**
Volume Cache Report	TABLE	Cache Devices
I/O Count	BAR	Cache Devices

Lowest Volume Read Hit Percent workspace		
View	**View type**	**Attribute group**
Volumes With Lowest Read Hit Percent	TABLE	Cache Devices
Lowest Read Hit Percent	BAR	Cache Devices
Staging Rate	BAR	Cache Devices

Lowest Volume Write Hit Percent workspace		
View	**View type**	**Attribute group**
Volumes With Lowest Write Hit Percent	TABLE	Cache Devices
Lowest Write Hit Percent	BAR	Cache Devices
Staging Rate	BAR	Cache Devices

Lowest Volume DFW Hit Percent workspace		
View	**View type**	**Attribute group**
Volumes With Lowest DFW Hit Percent	TABLE	Cache Devices
Lowest DFW Hit Percent	BAR	Cache Devices
DFW Retry Percent	BAR	Cache Devices

Highest Volume Destaging Rate workspace		
View	**View type**	**Attribute group**
Volumes With Highest Destaging Rate	TABLE	Cache Devices
Highest Destaging Rate	BAR	Cache Devices
Volume Staging Rate	BAR	Cache Devices

Highest Volume Staging Rate workspace		
View	**View type**	**Attribute group**
Volumes With Highest Staging Rate	TABLE	Cache Devices
Highest Staging Rate	BAR	Cache Devices
Volume Destaging Rate	BAR	Cache Devices

Lowest Volume CU Read Hit Percent workspace		
View	**View type**	**Attribute group**
Lowest Volume Read Hit Percent Report	TABLE	Cache Devices

View	View type	Attribute group
Lowest Read Hit Percent	BAR	Cache Devices
Staging Rate	BAR	Cache Devices

Lowest Volume CU Write Hit Percent workspace		
View	**View type**	**Attribute group**
Lowest Volume Write Hit Percent Report	TABLE	Cache Devices
Lowest Write Hit Percent	BAR	Cache Devices
Staging Rate	BAR	Cache Devices

Lowest Volume CU DFW Hit Percent workspace		
View	**View type**	**Attribute group**
Lowest Volume DFW Hit Percent Report	TABLE	Cache Devices
Lowest DFW Hit Percent	BAR	Cache Devices
DFW Retry Percent > 0	BAR	Cache Devices

Highest Volume CU Destaging Rate workspace		
View	**View type**	**Attribute group**
Highest Volume Destaging Rate Report	TABLE	Cache Devices
Highest Destaging Rate	BAR	Cache Devices
Volume Staging Rate > 0	BAR	Cache Devices

Highest Volume CU Staging Rate workspace		
View	**View type**	**Attribute group**
Highest Volume Staging Rate Report	TABLE	Cache Devices
Highest Staging Rate	BAR	Cache Devices
Volume Destaging Rate > 0	BAR	Cache Devices

Hitachi Data Systems 2107 Parity Groups workspace		
View	**View type**	**Attribute group**
HDS 2107 Parity Groups	TABLE	TotalStorageDS Rank
Operations Per Second	BAR	TotalStorageDS Rank
Average Response Time	BAR	TotalStorageDS Rank

Hitachi Data Systems 2107 Array Configuration workspace		
View	**View type**	**Attribute group**
DDM Capacity	BAR	TotalStorageDS Array
HDS 2107 Arrays	TABLE	TotalStorageDS Array
DDM Speed	BAR	TotalStorageDS Array

Hitachi Data Systems 2107 CU Volumes workspace		
View	**View type**	**Attribute group**
HDS 2107 CU Volumes	TABLE	Cache Devices
Physical I/O Rate	BAR	Cache Devices
Physical Response Time	BAR	Cache Devices

Hitachi Data Systems 2107 Extent Pool Volumes workspace		
View	**View type**	**Attribute group**
HDS 2107 Extent Pool Volumes	TABLE	Cache Devices
Physical I/O Rate	BAR	Cache Devices
Average Physical Response Time	BAR	Cache Devices

Hitachi Data Systems 2107 Extent Pool Trend workspace		
View	**View type**	**Attribute group**
HDS 2107 Extent Pool Trend	TABLE	TotalStorageDS Extent Pool
Response Time	PLOT	TotalStorageDS Extent Pool

Hitachi Data Systems 2107 Parity Group Trend workspace		
View	**View type**	**Attribute group**
HDS 2107 Parity Group Trend	TABLE	TotalStorageDS Rank
Average Response Time	PLOT	TotalStorageDS Rank

Hitachi Data Systems 2107 Volume Trend workspace		
View	**View type**	**Attribute group**
HDS 2107 Volume Trend	TABLE	Cache Devices
Physical Response Time	PLOT	Cache Devices

Hitachi Data Systems 2107 Parity Group History workspace		
View	**View type**	**Attribute group**
HDS 2107 Parity Groups	TABLE	TotalStorageDS Rank
Operations per Second	BAR	TotalStorageDS Rank
Average Response Time	BAR	TotalStorageDS Rank

Hitachi Data Systems 2107 Extent Pool Volume History workspace		
View	**View type**	**Attribute group**
HDS 2107 Volumes	TABLE	Cache Devices
Physical I/O Rate	BAR	Cache Devices
Physical Response Time	BAR	Cache Devices

Hitachi Data Systems 2107 Configuration workspace		
View	**View type**	**Attribute group**
Average Read Response Time	BAR	TotalStorageDS Extent Pool
Average Write Response Time	BAR	TotalStorageDS Extent Pool
HDS 2107 Extent Pools	TABLE	TotalStorageDS Extent Pool
HDS 2107 SSIDs	TABLE	TotalStorageDS SSIDs

View	View type	Attribute group
HDS 2107 Configuration	TABLE	TotalStorageDS Configuration

Monitoring cache CU status

This topic describes the workspaces that monitor the status of cache CU. You access these workspaces through the Cache CU Status node in the Tivoli® Enterprise Portal.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

Navigating the workspaces of the Cache CU Status node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the Cache CU Status node to navigate to the workspaces that you see in the following bulleted list. You can access the workspaces in this table using the standard path and method for navigating workspaces.

Cache CU Status node: bulleted list showing the hierarchy of linked workspaces

- Cache CU Volume Cache Performance
  - Dataset Performance Summary [1]
  - Physical Volume Group Cache
    - Dataset Performance Summary [1]
  - Device Performance Details
  - Cross System Volume Performance
    - Dataset Performance Summary [1]
  - Volume Cache Trend
- Cache CU Volume Performance
  - Dataset Performance Summary [1]
  - Physical Volume Group Performance
    - Dataset Performance Summary [1]
  - Volume Performance Trend
    - Dataset Performance Summary History
  - Device Performance Details
  - Cross System Volume Performance
    - Dataset Performance Summary [1]
  - Volume Performance Trend
    - Dataset Performance Summary History
      - Dataset Performance Historical Detail
- Symmetrix Configuration **Note:** Obtaining EMC Symmetric data can only be done in conjunction with the applicable EMC hardware and software. The hardware required is an EMC Symmetrix at the appropriate microcode level. Also required are the EMC Symmetrix Control Facility runtime libraries available from EMC.

- Cache CU Raid Rank
  - Physical Volume Group Cache
    - Dataset Performance Summary [1]
  - Cache CU Raid Rank Trend
- Cache CU Status Trend
  - CCU Volume Cache History
    - Dataset Historic Performance
      - Dataset Performance Historical Detail
- Hitachi Data Systems 2107 Configuration
  - Hitachi Data Systems 2107 Parity Groups
    - Hitachi Data Systems 2107 Array Configuration
    - Hitachi Data Systems 2107 Parity Group Trend
  - Hitachi Data Systems 2107 Extent Pool Volumes
    - Hitachi Data Systems 2107 Volume Trend
  - Hitachi Data Systems 2107 Extent Pool Trend
    - Hitachi Data Systems 2107 Parity Group History
    - Hitachi Data Systems 2107 Extent Pool Volume History
- Hitachi Data Systems 2107 CU Volumes
  - Hitachi Data Systems 2107 Volume Trend
- TotalStorage Configuration
  - TotalStorage Ranks
    - TotalStorage Array Configuration
    - TotalStorage Ranks Trend
  - TotalStorage Extent Pool Volumes
    - TotalStorage Volume Trend
  - TotalStorage Extent Pool Trend
    - TotalStorage Rank History
    - TotalStorage Volume History
- TotalStorage CU Volumes
  - TotalStorage Volume Trend

Note:

1. This group of workspaces is organized in the following nested pattern:
   - Dataset Performance Summary
     - Dataset Performance Detail
       - Dataset Details
       - Dataset Performance Detail Trend
     - Dataset Details
     - Dataset Performance Summary Trend
       - Dataset Performance Historical Detail

Cache CU Status workspace

This topic describes the main components of the principal workspace for this node.

The direct situation analysis feature is available for this workspace. The feature can help you analyze and understand a problem that is displayed in this workspace. See "Directly monitoring situations" on page 133.

Views

This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

Cache status table view

The table view displays one row of data for each cache control unit configured to the reporting LPAR. The view contains the following information:

- The subsystem ID of the control unit
- The control unit type
- The number of active volumes
- The number of deactivated volumes
- The cache status
- The amount of cache configured in the control unit and the amount currently available
- The nonvolatile storage status, the amount of NVS configured and the amount pinned
- Storage facility series

Available cache bar graph

This graph contains a bar for each cache control unit. The height of the bar indicates the amount of cache in MB available in the unit.

NVS pinned bar graph

This graph contains a bar for each cache control unit. The height of the bar indicates the amount of NVS (non-volatile storage, in KB) pinned in the unit.

Deactivated volumes bar graph

This graph contains a bar for each cache control unit. The height of the bar indicates the number of deactivated volumes on that control unit.

Workspaces that are linked to this workspace

See "Navigating the workspaces of the Cache CU Status node" on page 42.
Learn how to navigate to workspaces here: "Workspaces" on page 14.

See "Workspaces" on page 14 for a list of all nodes provided in IBM® Z OMEGAMON® AI for Storage®.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the Cache CU Status node.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

Cache CU Status workspace		
View	**View type**	**Attribute group**
Cache Control Unit Status Report	TABLE	Cache Control Unit

View	View type	Attribute group
Available Cache	BAR	Cache Control Unit
NVS Pinned	BAR	Cache Control Unit
Deactivated Volumes	BAR	Cache Control Unit

Cache CU Status Trend workspace		
View	**View type**	**Attribute group**
Cache Control Unit Status Trend	TABLE	Cache Control Unit
NVS Pinned	PLOT	Cache Control Unit
Deactivated Volumes	PLOT	Cache Control Unit
Cache Available	PLOT	Cache Control Unit

TotalStorage® DS Configuration workspace		
View	**View type**	**Attribute group**
Average Read Response Time	BAR	TotalStorageDS Extent Pool
Average Write Response Time	BAR	TotalStorageDS Extent Pool
TotalStorage® Extent Pools	TABLE	TotalStorageDS Extent Pool
TotalStorage® SSIDs	TABLE	TotalStorageDS SSIDs
TotalStorage® Configuration	TABLE	TotalStorageDS Configuration

Hitachi Data Systems 2107 Configuration workspace		
View	**View type**	**Attribute group**
Average Read Response Time	BAR	TotalStorageDS Extent Pool
Average Write Response Time	BAR	TotalStorageDS Extent Pool
HDS 2107 Extent Pools	TABLE	TotalStorageDS Extent Pool
HDS 2107 SSIDs	TABLE	TotalStorageDS SSIDs
HDS 2107 Configuration	TABLE	TotalStorageDS Configuration

Monitoring logical control units

This topic describes the workspaces that monitor the status of logical control units. You access these workspaces through the Logical Control Unit node in the Tivoli® Enterprise Portal.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

Navigating the workspaces of the Logical Control Unit node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the Logical Control Unit node to navigate to the workspaces that you see in in the following hierarchy. You can access the workspaces in this table using the standard path and method for navigating workspaces.

Logical Control Unit node: bulleted list showing the hierarchy of linked workspaces

- LCU Volume Performance
  - Dataset Performance Summary [1]
  - Physical Volume Group Performance
    - Dataset Performance Summary [1]
    - Volume Performance Trend
      - Dataset Performance Summary History
  - Device Performance Details
  - Cross System Volume Performance
    - Dataset Performance Summary [1]
  - Volume Performance Trend
    - Dataset Historic Performance Summary
      - Dataset Performance Historical Detail
- LCU Channel Detail
- LCU Volume Cache
  - Dataset Performance Summary [1]
  - Physical Volume Group Cache
    - Dataset Performance Summary [1]
  - Device Performance Details
  - Cross System Volume Performance
    - Dataset Performance Summary [1]
  - Volume Cache Trend
- LCU DCM Channel Summary
- Logical Control Unit Trend
  - LCU Volume Historical Performance
    - Dataset Historic Performance Summary
      - Dataset Performance Historical Detail
  - LCU Historical Volume Cache Performance
    - Dataset Historic Performance Summary
      - Dataset Performance Historical Detail

  These workspaces display the volumes with the least favorable readings for the attribute within the logical control unit selected for navigation.
- Highest LCU Volumes Response Time (Top 10) [2]
- Highest LCU Volumes MPL (Top 10) [2]
- Highest LCU Volumes I/O Rate (Top 10) [2]
- Highest LCU Volumes Busy Percent (Top 10) [2]
- Lowest LCU Volumes Connect Percent (Top 10)[2]

Note:

1. This group of workspaces is organized in the following nested pattern:

- Dataset Performance Summary
  - Dataset Performance Detail
    - Dataset Details
    - Dataset Performance Detail Trend
  - Dataset Details
  - Dataset Performance Summary Trend
    - Dataset Performance Historical Detail

2. This workspace links to the following set of related workspaces:
   - Dataset Performance Summary (See the preceding note (1) for details.)
   - Physical Volume Group Performance
     - Dataset Performance Summary (See the preceding note (1) for details.)
     - Volume Performance Trend
       - Dataset Performance Summary History
   - Device Performance Details
   - Cross System Volume Performance
     - Dataset Performance Summary [1]
   - Volume Performance Trend
     - Dataset Historic Performance Summary
       - Dataset Performance Historical Detail

Logical Control Unit workspace

This topic describes the main components of the principal workspace for this node.

The direct situation analysis feature is available for this workspace. The feature can help you analyze and understand a problem that is displayed in this workspace. See "Directly monitoring situations" on page 133.

Views

This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

Logical control unit table view

The table view displays one row of data for each logical control unit configured in the reporting LPAR. The view contains the following information:

- The logical control unit number
- The average delay queue
- The contention rate
- The director port busy percentage
- The control unit busy percentage
- The channel path I/O rate
- The number of attached channels and whether there are DCM channels attached to the control unit
- The high volume utilization rate
- The high volume millisecond response time
- The low read hit percentage

- The low write hit percentage
- The high DFW retry percentage
- The number of SMS and non-SMS volumes
- The total number of volumes

Average delay queue bar graph

This graph contains one bar for each logical control unit. The height of the bar indicates the average number of delayed requests on the control unit header.

Average contention rate bar graph

This graph contains one bar for each logical control unit. The height of the bar indicates the contention rate for the unit.

Workspaces that are linked to this workspace

See "Navigating the workspaces of the Logical Control Unit node" on page 45.
Learn how to navigate to workspaces here: "Workspaces" on page 14.

See "Workspaces" on page 14 for a list of all nodes provided in IBM® Z OMEGAMON® AI for Storage®.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the Logical Control Unit node.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

See "Logical Control Unit workspace" on page 47 for more information on the primary workspace for this node.

Logical Control Unit workspace		
View	**View type**	**Attribute group**
Logical Control Unit Report	TABLE	Logical Control Unit
Average Delay Queue	BAR	Logical Control Unit
Average Contention Rate	BAR	Logical Control Unit

LCU Volume Performance workspace		
View	**View type**	**Attribute group**
Volume Response Time > 25	BAR	DASD Volume Performance
Volume Performance Report	TABLE	DASD Volume Performance

LCU Channel Detail workspace		
View	**View type**	**Attribute group**
Control Unit Percent Busy	BAR	Logical Control Unit Channel Paths
Complex Channel Percent Busy	BAR	Logical Control Unit Channel Paths
Director Port Percent Busy	BAR	Logical Control Unit Channel Paths
Logical Control Unit Details	TABLE	Logical Control Unit Channel Paths

LCU Volume Cache workspace		
View	**View type**	**Attribute group**
Volume Cache Report	TABLE	Cache Devices
Read Hit Percent < 55	BAR	Cache Devices
Write Hit Percent < 30	BAR	Cache Devices

LCU DCM Channels Summary workspace		
View	**View type**	**Attribute group**
DCM Configuration	BAR	Logical Control Unit
DCM Channels Percent Busy	BAR	Logical Control Unit
DCM Channels I/O per Second	BAR	Logical Control Unit

Logical Control Unit Trend workspace		
View	**View type**	**Attribute group**
Logical Control Unit Trend Report	TABLE	Logical Control Unit
LCU Trend	PLOT	Logical Control Unit

LCU Volume Event Performance workspace		
View	**View type**	**Attribute group**
Volume Performance Report	TABLE	DASD Volume Performance
Volume Response Time	BAR	DASD Volume Performance

LCU HyperPAV Details workspace		
View	**View type**	**Attribute group**
HyperPAV Report	TABLE	LCU Control Unit
I/Os Queued - Alias Unavailable	BAR	LCU Control Unit
Percent I/Os Queued - Alias Unavailable	BAR	LCU Control Unit

LCU HyperPAV Details Trend workspace		
View	**View type**	**Attribute group**
HyperPAV Report	TABLE	LCU Control Unit
Percent I/Os Queued - Alias Unavailable	PLOT	LCU Control Unit

Monitoring tape groups

This topic describes the workspaces that monitor the status of tape groups. You access these workspaces through the Tape Group node in the Tivoli® Enterprise Portal.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

Navigating the workspaces of the Tape Group node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the Tape Group node to navigate to the workspaces that you see in following node hierarchy. You can access the workspaces in this table using the standard path and method for navigating workspaces.

Tape Group node: bulleted list showing the hierarchy of linked workspaces

- Tape Devices

Tape Group workspace

This topic describes the main components of the principal workspace for this node.

The following sections describe various components of this workspace.

Views

This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

Tape group table view

The table view displays one row of data for each tape group (device or tape library) configured in the reporting LPAR. The view contains the following information:

- The library name or device type
- The library ID number
- The library type (or Non_Library for non-library devices)
- The library device type
- The console name
- The library description
- The number of configured drives, online drives and available drives
- The number of pending mounts and the length of time the current mount has been pending
- The total number of slots in the device and the number currently available
- The number of scratch volumes for 3480 type magnetic tape cartridges (Media1), for 3490E type enhanced capacity magnetic tape cartridges (Media2), for 3590 type enhanced capacity magnetic tape cartridges (Media3), and for 3590 type extended tape cartridges (Media 4)
- The I/O rate
- The number temporary and permanent errors encountered on the device

Mount pend time > 5 minutes bar graph

This graph contains one bar for each tape group which has been waiting more than 5 minutes for a tape mount. The bar height indicates how long the mount has been pending.

Workspaces that are linked to this workspace

See “Navigating the workspaces of the Tape Group node” on page 50.
Learn how to navigate to workspaces here: “Workspaces” on page 14.

See “Workspaces” on page 14 for a list of all nodes provided in IBM® Z OMEGAMON® AI for Storage®.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the Tape Group node.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

See "Tape Group workspace" on page 50 for more information on the primary workspace for this node.

Tape Group workspace		
View	**View type**	**Attribute group**
Tape Group Report	TABLE	Tape Group
Mount Pend Time > 5 Minutes	BAR	Tape Group

Tape Devices workspace		
View	**View type**	**Attribute group**
Tape Device Report	TABLE	Tape Device
Mount Pend Time > 5 Minutes	BAR	Tape Device

Monitoring Virtual Tape Subsystems (VTS)

This topic describes the workspaces that monitor the status of virtual tape subsystems.

You access these workspaces through the Virtual Tape Subsystems node in the Tivoli® Enterprise Portal.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

Navigating the workspaces of the Virtual Tape Subsystems node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the Virtual Tape Subsystems node to navigate to the workspaces that you see in following hierarchy. You can access the workspaces in this table using the standard path and method for navigating workspaces .

Virtual Tape Subsystems node: bulleted list showing the hierarchy of linked workspaces

- VTS Physical Devices
- VTS Virtual Devices
- VTS Capacity
- VTS Tape Volume Cache
- VTSS Status
  - VTSS Real Tape Drives
- TS7700 Series Cluster
  - TS7700 Series Cache
    - TS7700 Series Cache Partition
      - TS7700 Series Cache Preference Group
- TS7700 Series Cluster Healthcheck
  - TS7700 Series Cache Healthcheck

        - TS7700 Series Cache Partition Healthcheck
            - TS7700 Series Cache Preference Group Healthcheck
- VTS Virtual Drive Summary
    - VTS Composite Details

Virtual Tape Subsystems workspace

This workspace is the principal workspace for this node, and it enables you to monitor both types of virtual tape subsystems, 3494 and 7700.

In a table view, when you right-click the row for a 3494 subsystem, you see standard (non-7700) options. When you select a 7700 subsystem, you see only the link to the TS7700 Series Cluster Workspace.

When you click the Virtual Tape Subsystems node in the Navigator of the Tivoli® Enterprise Portal, you see the Virtual Tape Subsystems workspace. The main components of the workspace are as follows:

Views

This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

Virtual tape subsystems table view

The table view contains one row of data for virtual tape subsystem. The view contains the following information:

- The VTS name and library type
- The number of physical drives available
- The number of virtual devices online
- The average and maximum virtual mount pending times
- The virtual disconnect time
- The volume data (in GB) read and written through the host channels to virtual volumes associated with this VTS
- The recall percent
- Number of scratch volumes, library type, model and serial number for this VTS

Volume recall % > 20 bar graph

This graph contains one bar for each VTS for which the percentage of virtual mounts to actual physical mounts exceeds 20%. The bar height indicates the recall percentage.

Average virtual mount pending time > 300 bar graph

This graph contains one bar for each VTS for which the average time spent waiting for a virtual mount exceeded 300 seconds. The height of the bar indicates the actual time spent waiting.

TS7700 Worst Statistics table view

The table view contains one row of data for each TS7700 Node/Cluster. The view contains the following information:

- The hexidecimal field indicates the Node ID
- The hexidecimal field indicates the cluster ID.
- The node's machine type.
- The node's model type.
- The serial number of this node.
- The code level of the Virtualization Engine.
- The Sequence Number of the Grid (Composite) Library.
- The minimum number of physcial drives available on this cluster of the select TS7700 grid for the last 24 hours.

- The minimum average amount of time in seconds that and z/OS system waited for a virtual tape volume to be mounted on this cluster of the selected TS7700 grid for the last 24 hours.
- The maximum total number of GB (read + write) processed by the z/OS host channels to the virtual tape volumes associated with this cluster of the selected TS7700 grid for the last 24 hours.
- The maximum percentage of the virtual mounts that result in a physical mount in this cluster of the selected TS7700 grid for the last 24 hours.
- The minimum number of empty physical cartridges for use in this cluster of the selected TS700 grid for the last 24 hours.
- The Name of this TS7700 Cluster.
- Cluster ID and Machine Serial Number.
- The maximum size of the Cluster's TVC in increments of GB (1024x1024x1000) for the last 24 hours.
- The maximum number of active logical volumes for the last 24 hours.
- The maximum number of GB of the active data for the last 24 hours.
- The maximum number of total mounts for the last 24 hours.
- The maximum number of scratch mounts for the last 24 hours.
- The maximum GB Read by the Channel for the last 24 hours.
- The maximum GB Read by the Channel for the last 24 hours.
- The maximum GB Written by the Channel for the last 24 hours.
- The maximum GB Written by the Channel for the last 24 hours.
- The maximum sum of MB Read and MB Written for the last 24 hours.
- The maximum sum of GB Read and GB Written for the last 24 hours.
- The maximum percentage of 30-second periods where there was at least one throttling value greater than zero and that write overrun was the predominant reason for that throttling for the last 24 hours.
- The maximum average throttle value during the interval for the last 24 hours.
- The maximum percentage of 30-second periods where there was at least one throttling value greater than zero and that copy was the predominant reason for the throttling for the last 24 hours.
- The maximum average copy throttle value during the interval for the last 24 hours.
- The maximum average of all throttle values during the interval for the last 24 hours.
- The maximum percentage of 30-second periods where deferred copy throttle from cache partition 0 of this Hnode was active (non-zero) for the last 24 hours.
- The maximum average deferred copy throttle value during the interval for the last 24 hours.
- The maximum base deferred copy throttle value which would be used if deferred copy throttle were to be applied for the last 24 hours.
- The maximum reason(s) for write overrun (host) throttling during the interval for the last 24 hours.
- The maximum reason(s) for copy throttling during the interval for the last 24 hours.
- The maximum reason(s) for deferred copy throttling during the interval for the last 24 hours.

Workspaces that are linked to this workspace

See “Navigating the workspaces of the Virtual Tape Subsystems node” on page 51.
See standard path and method for navigating workspaces to learn how to navigate among workspaces and to check the list of all nodes provided in IBM® Z OMEGAMON® AI for Storage®.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the Virtual Tape Subsystems node.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

See “Virtual Tape Subsystems workspace” on page 52 for more additional information on the primary workspace for this node.

Virtual Tape Subsystems workspace		
View	**View type**	**Attribute group**
Virtual Tape Subsystems Report	TABLE	VTS Overview
Volume Recall % > 20	BAR	VTS Overview
Average Virtual Mount Pend Time > 300	BAR	VTS Overview
TS7700 Worst Statistics	TABLE	VTS Cluster Worst Statistics Summary

VTS Physical Devices workspace		
View	**View type**	**Attribute group**
Mount Pend Time	BAR	VTS Physical Devices
Concurrent Mounts	BAR	VTS Physical Devices
Physical Drives Status	BAR	VTS Physical Devices
Physical Mount Reason	BAR	VTS Physical Devices

VTS Virtual Devices workspace		
View	**View type**	**Attribute group**
Concurrent Mounts	BAR	VTS Virtual Devices
Mount Pend Time	BAR	VTS Virtual Devices
Mount Use Time	BAR	VTS Virtual Devices
Virtual Device Activity Report	TABLE	VTS Virtual Devices
Recall Mount Time	BAR	VTS Virtual Devices

VTS Capacity workspace		
View	**View type**	**Attribute group**
VTS Capacity Report	TABLE	VTS Capacity
Host Channel Activity	BAR	VTS Capacity
VTS Channel Activity	BAR	VTS Capacity
Physical Tape Usage	BAR	VTS Capacity

VTS Tape Volume Cache workspace		
View	**View type**	**Attribute group**
VTS Tape Volume Cache Report	TABLE	VTS Cache
Throttling Delay	BAR	VTS Cache
TVC Volume Age	BAR	VTS Cache
Virtual Mounts	BAR	VTS Cache
Throttling Percent	PIE	VTS Cache

VTS Composite Details workspace		
View	**View type**	**Attribute group**
VTS Controller Report	TABLE	VTS Composite Library
Mounts by Unit	BAR	VTS Virtual Devices
Logical Volumes to be Copied	BAR	VTS Virtual Devices
Data to Copy	BAR	VTS Virtual Devices

VTS Virtual Drive Summary workspace		
View	**View type**	**Attribute group**
Virtual Tape Subsystems Virtual Tape Drives	TABLE	VTS Virtual Devices Summary
Tape Drive Response Time	BAR	VTS Virtual Devices Summary

VTSS Status workspace		
View	**View type**	**Attribute group**
VTSS Scratch Volumes	TABLE	VTSS Virtual Volumes
VTSS Configuration	TABLE	VTSS Configuration
VTSS Status	TABLE	VTSS Status

VTSS Real Tape Drives workspace		
View	**View type**	**Attribute group**
VTSS RTD Status	BAR	VTSS RTD Status
VTSS RTD Queue Report	TABLE	VTSS RTD Queue
VTSS RTD Status	TABLE	VTSS RTD Status

Workspaces related to TS7700-series VTS

The following attribute groups supply data to the workspaces that are related to TS7700-series VTS:

- "VTS TS7700-series Cache Container" on page 385
- "VTS TS7700-series Cache Partition Container" on page 386
- "VTS TS7700-series Cache Preference Group" on page 387
- "VTS TS7700-series Cluster" on page 388

TS7700 Series Cluster workspace			
View	**View type**	**Attribute group**	**Description**
Cluster Summary	TABLE	VTS TS7700-series Cluster	Contains basic information about a cluster in a TS7700 tape virtualization system, such as ID and machine type and model. The view also contains metrics that are derived from standard sources, such as the virtualization (vNode) or hierarchical data storage management (hNode).
Volume Recall Percent	BAR	VTS TS7700-series Cluster	Charts the following metric for each cluster in the grid: Volume Recall Percent
Average Virtual Mount Pend Time	BAR	VTS TS7700-series Cluster	Charts the following metric for each cluster in the grid: Average Virtual Mount Pend Time.

TS7700 Series Cache workspace			
View	**View type**	**Attribute group**	**Description**
Cache Throttling Percentages	BAR	VTS TS7700-series Cache Container	Charts the following two metrics: Percent Write Overrun Throttle and Percent Copy Throttle.
Cache Throttling Averages	BAR	VTS TS7700-series Cache Container	Charts the following three metrics: Average Write Overrun Throttle, Average Copy Throttle, and Average Overall Throttle.
Tape Volume Cache Summary	TABLE	VTS TS7700-series Cache Container	Contains basic information about a cache in a TS7700 tape virtualization system, based on metrics that are derived from the "HSM - Cache Container" record of the "Hnode HSM Historical Record."

TS7700 Series Cache Partition workspace			
View	**View type**	**Attribute group**	**Description**
Average Cache Miss Mount Times	BAR	VTS TS7700-series Cache Partition Container	Charts the following metric for each of the 8 partitions: Average Cache Miss Mount Time.
Cache Miss Mounts	BAR	VTS TS7700-series Cache Partition Container	Charts the following metric for each of the 8 partitions: Cache Miss Mounts.

View	View type	Attribute group	Description
Cache Partition Summary	TABLE	VTS TS7700-series Cache Partition Container	Contains basic information about a cache partition in a TS7700 tape virtualization system. All attributes in this table are derived from the "HSM - Cache - Partition Container" record of the "Hnode HSM Historical Record."

TS7700 Series Cache Preference Group workspace			
View	**View type**	**Attribute group**	**Description**
Average Cache Age	BAR	VTS TS7700-series Cache Preference Group	Charts the following three metrics: 4 Hour Average Cache Age, 48 Hour Average Cache Age, and 35 Day Average Cache Age.
Number of Volumes Migrated	BAR	VTS TS7700-series Cache Preference Group	Charts the following three metrics: Volumes Migrated Last 4 Hours, Volumes Migrated Last 48 Hours, and Volumes Migrated Last 35 Days.
Cache Preference Group Summary	TABLE	VTS TS7700-series Cache Preference Group	Contains basic information about a cache preference group in a TS7700 tape virtualization system, based on metrics that are derived from the HSM HSM - Cache - Partition - Preference Group Container record of the "Hnode HSM Historical Record."

TS7700 Series Cluster Healthcheck workspace			
View	**View type**	**Attribute group**	**Description**
Cluster Healthcheck Summary	TABLE	VTS CLUSTER	Contains basic information about a cluster in a TS7700 tape virtualization system, such as ID and machine type and model. The view also contains metrics that are derived from standard sources, such as the virtualization (vNode) or hierarchical data storage management (hNode).

TS7700 Series Cache Healthcheck workspace			
View	**View type**	**Attribute group**	**Description**
Tape Volume Cache Healthcheck Summary	TABLE	VTS TPVOLC	Contains basic information about a cache in a TS7700 tape virtualization system, based on metrics that are derived from the "HSM - Cache Container" record of the "Hnode HSM Historical Record."

TS7700 Series Cache Partition Healthcheck workspace			
View	**View type**	**Attribute group**	**Description**
Cache Partition Healthcheck Summary	TABLE	VTS CACHEPC	Contains basic information about a cache partition in a TS7700 tape virtualization system. All attributes in this table are derived from the "HSM - Cache - Partition Container" record of the "Hnode HSM Historical Record."

TS7700 Series Cache Preference Group Healthcheck workspace			
View	**View type**	**Attribute group**	**Description**
Cache Preference Group Healthcheck Summary	TABLE	VTS CHPGRP	Contains basic information about a cache preference group in a TS7700 tape virtualization system, based on metrics that are derived from the HSM HSM - Cache - Partition - Preference Group Container record of the "Hnode HSM Historical Record."

Monitoring SMS storage groups performance

This topic describes the workspaces that monitor the status of SMS storage groups performance. You access these workspaces through the SMS Storage Groups Performance node in the Tivoli® Enterprise Portal.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

You can also issue volume-level Storage Toolkit commands from the SMS Storage Group report level. To access these commands, open the SMS Storage Groups Performance workspace from the Navigator, right-click a row on the table view, and select the **Volume Actions...** menu item, then select the action you want to perform.

Navigating the workspaces of the SMS Storage Groups Performance node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the SMS Storage Groups Performance node to navigate to the workspaces that you see in the following hierarchy. You can access the workspaces in this table using the standard path and method for navigating workspaces.

SMS Storage Groups Performance node: bulleted list showing the hierarchy of linked workspaces

- SMS Storage Group Volume Performance
  - Dataset Performance Summary [1]
  - Volume Performance Trend
    - Dataset Historic Performance Summary
      - Dataset Performance Historical Detail
  - Physical Volume Group Performance
    - Dataset Performance Summary [1]
  - Volume Performance Trend
    - Dataset Historic Performance Summary
      - Dataset Performance Historical Detail
  - Device Performance Details
  - Cross System Volume Performance
    - Dataset Performance Summary [1]
- SMS Storage Group Volume Cache
  - Dataset Performance Summary [1]
  - Physical Volume Group Cache
    - Dataset Performance Summary [1]
  - Device Performance Details
  - Cross System Volume Performance
    - Dataset Performance Summary [1]
  - Volume Cache Trend
- Cross System Group Summary
  - Cross System Volume Performance
    - Dataset Performance Summary [1]
- Storage Group Performance Trend
  - SMS Storage Group Historical Volume Performance
    - Dataset Historic Performance Summary
      - Dataset Performance Historical Detail
  - SMS Storage Group Historic Volume Cache Performance
    - Dataset Historic Performance Summary
      - Dataset Performance Historical Detail

These workspaces list the volumes with the least favorable readings for the attribute within the SMS storage group selected for navigation:

- Highest Volume Response Time (Top 15) [2]
- Highest Volume MPL (Top 15) [2]
- Highest Volume I/O Rate (Top 15) [2]
- Highest Volume Busy Percent (Top 15) [2]
- Lowest Volume Connect Percent (Top 15) [2]

- Highest Volume Group Response Time (Top 15) [2]
- Highest Volume Group MPL (Top 15) [2]
- Highest Volume Group I/O Rate (Top 15) [2]
- Highest Volume Group Busy Percent (Top 15) [2]
- Lowest Volume Group Connect Percent (Top 15) [2]

Note:

1. This group of workspaces is organized in the following nested pattern:
    - Dataset Performance Summary
        - Dataset Performance Detail
            - Dataset Details
            - Dataset Performance Detail Trend
        - Dataset Details
        - Dataset Performance Summary Trend
            - Dataset Performance Historical Detail
2. This workspace links to the following set of related workspaces:
    - Dataset Performance Summary (See the preceding note ([1]) for details.)
    - Physical Volume Group Performance
        - Dataset Performance Summary [1]
    - Device Performance Details
    - Cross System Volume Performance
        - Dataset Performance Summary [1]
    - Volume Performance Trend
        - Dataset Historic Performance Summary
            - Dataset Performance Historical Detail

Access the following linked workspaces by selecting the SMS Storage Groups Performance node in the Navigator. Right-click the node. Select the workspace that you want to access in the pop-up menu.

These workspaces list the volumes with the least favorable readings for the attribute across the z/OS® image.

- Highest Volume Response Time *
- Highest Volume MPL *
- Highest Volume I/O Rate *
- Highest Volume Busy Percent *
- Lowest Volume Connect Percent *

Each of the workspaces in the preceding bulleted list link to the following set of related workspaces:

- Dataset Performance Summary
    - Dataset Performance Detail
        - Dataset Details
        - Dataset Performance Detail Trend
    - Dataset Details

- Dataset Performance Summary Trend
  - Dataset Performance Historical Detail

- Physical Volume Group Performance
  - Dataset Performance Summary [1]
- Device Performance Details
- Cross System Volume Performance
  - Dataset Performance Summary [1]
- Volume Performance Trend
  - Dataset Historic Performance Summary
    - Dataset Performance Historical Detail

SMS Storage Groups Performance workspace

This topic describes the main components of the principal workspace for this node.

The direct situation analysis feature is available for this workspace. The feature can help you analyze and understand a problem that is displayed in this workspace. See “Directly monitoring situations” on page 133.

Views

This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

SMS storage groups performance table view

The table view contains one row of data for each SMS storage group defined in the reporting LPAR. The view contains the following information:

- The name of the group
- The type
- The SMS status
- The number of volumes in the group
- The current highest response time for any volume in the group
- The current high busy percent for any volume in the group
- The device multi-processing level
- The low read and write hit percents among all volumes in the group
- The highest DFW retry percentage in the group

High volume response time > 25 bar graph

This graph contains one bar for each SMS storage group with a highest volume response time that exceeds 25 milliseconds. The height of the bar indicates the actual response time for that volume.

Workspaces that are linked to this workspace

See “Navigating the workspaces of the SMS Storage Groups Performance node” on page 58.
Learn how to navigate to workspaces here: “Workspaces” on page 14.

See “Workspaces” on page 14 for a list of all nodes provided in IBM® Z OMEGAMON® AI for Storage®.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the SMS Storage Groups Performance node.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

See "SMS Storage Groups Performance workspace" on page 61 for more information on the primary workspace for this node.

SMS Storage Groups Performance workspace		
View	**View type**	**Attribute group**
High Volume Response Time > 25	BAR	Volume Group Summary
SMS Storage Group Performance Report	TABLE	Volume Group Summary

SMS Storage Group Volume Performance workspace		
View	**View type**	**Attribute group**
Volume Response Time > 25	BAR	DASD Volume Performance
Volume Performance Report	TABLE	DASD Volume Performance

Dataset Performance Summary workspace		
View	**View type**	**Attribute group**
Dataset Performance Summary Report	TABLE	Dataset Performance
Datasets with MSR > Storage Class Objective	TABLE	Dataset Performance

Dataset Performance Detail workspace		
View	**View type**	**Attribute group**
Dataset Performance Detail Report	TABLE	Dataset Performance
Component Percent of Dataset MSR for All Jobs	PIE	Dataset Performance

SMS Storage Group Volume Cache workspace		
View	**View type**	**Attribute group**
Read Hit Percent < 55	BAR	Cache Devices
Write Hit Percent < 30	BAR	Cache Devices
Volume Cache Report	TABLE	Cache Devices

Volume Performance Trend workspace		
View	**View type**	**Attribute group**
Volume Performance Trend Report	TABLE	DASD Volume Performance
Response Time Analysis	PLOT	DASD Volume Performance

Physical Volume Group Performance workspace		
View	**View type**	**Attribute group**
Volume Response Time > 25	BAR	DASD Physical Group

View	View type	Attribute group
Physical Volume Group Performance Report	TABLE	DASD Physical Group

Physical Volume Group Cache workspace		
View	**View type**	**Attribute group**
Read Hit Percent < 55	BAR	DASD Physical Group
Write Hit Percent < 30	BAR	DASD Physical Group
Physical Volume Group Cache Report	TABLE	DASD Physical Group

Dataset Details workspace		
View	**View type**	**Attribute group**
Dataset Space Attributes	TABLE	Dataset Space
Dataset Users	TABLE	Dataset Users
Dataset SMS Constructs	TABLE	Dataset SMS
Dataset Performance Summary	TABLE	Dataset Performance

Device Performance Details workspace		
View	**View type**	**Attribute group**
Volume Cache Performance Details Report	TABLE	Cache Devices
Volume Status	TABLE	Volume Status
Volume Details	TABLE	DASD Volume Performance
Volume Performance Details Report	TABLE	DASD Volume Performance
Volume Users	TABLE	Volume Users

Device Performance Details workspace		
View	**View type**	**Attribute group**
Volume Cache Performance Details Report	TABLE	Cache Devices
Volume Status	TABLE	Volume Status
Volume Details	TABLE	DASD Volume Performance
Volume Performance Details Report	TABLE	DASD Volume Performance
Volume Users	TABLE	Volume Users

Cross System Group Summary workspace		
View	**View type**	**Attribute group**
Device Contention Index	BAR	Cross System Vol Summary
Cross System Performance Report	TABLE	Cross System Vol Summary

Cross System Volume Performance workspace		
View	**View type**	**Attribute group**
Response Time	BAR	Cross System Volumes
Percent Busy	BAR	Cross System Volumes

View	View type	Attribute group
Cross System Volume Performance Report	TABLE	Cross System Volumes

SMS Storage Group Performance Trend workspace		
View	**View type**	**Attribute group**
High Volume Response	BAR	Volume Group Summary
SMS Storage Group Performance Trend Report	TABLE	Volume Group Summary

Group Volume Historical Performance workspace		
View	**View type**	**Attribute group**
Volume Performance Report	TABLE	DASD Volume Performance
Volumes with Response Time > 25	BAR	DASD Volume Performance

Dataset Performance Summary Trend workspace		
View	**View type**	**Attribute group**
Dataset Performance Summary Trend	TABLE	Dataset Performance
Response Time Analysis	PLOT	Dataset Performance

Dataset Historic Performance Summary workspace		
View	**View type**	**Attribute group**
Dataset Performance Summary History Report	TABLE	Dataset Performance
Datasets with MSR > Storage Class Objective History	TABLE	Dataset Performance

Dataset Performance Historical Detail workspace		
View	**View type**	**Attribute group**
Dataset Performance Detail Report	TABLE	Dataset Performance
Component Percent of Dataset MSR for All Jobs	PIE	Dataset Performance

Dataset Performance Detail Trend workspace		
View	**View type**	**Attribute group**
Dataset Performance Application Trend	TABLE	Dataset Performance
Response Time Analysis	PLOT	Dataset Performance

Dataset Performance Event workspace		
View	**View type**	**Attribute group**
Dataset Performance Summary Report	TABLE	Dataset Performance
Dataset Response Time	BAR	Dataset Performance

Highest Volume Response Time workspace		
View	**View type**	**Attribute group**
Highest Volume Response Time Report	TABLE	DASD Volume Performance
Volume Response Time	BAR	DASD Volume Performance

Highest Volume MPL workspace		
View	**View type**	**Attribute group**
Highest Volume MPL Report	TABLE	DASD Volume Performance
Volume MPL	BAR	DASD Volume Performance
Volume Response Time	BAR	DASD Volume Performance

Highest Volume I/O Rate workspace		
View	**View type**	**Attribute group**
Highest Volume I/O Report	TABLE	DASD Volume Performance
Provide summary performance information for the datasets on a specified volume.	BAR	DASD Volume Performance
Volume Response Time	BAR	DASD Volume Performance

Highest Volume Busy Percent workspace		
View	**View type**	**Attribute group**
Highest Volume Busy Percent Report	TABLE	DASD Volume Performance
Volume Busy Percent	BAR	DASD Volume Performance
Volume Response Time	BAR	DASD Volume Performance

Lowest Volume Connect Percent workspace		
View	**View type**	**Attribute group**
Lowest Volume Connect Time Percent Report	TABLE	DASD Volume Performance
Volume Connect Time Percent	BAR	DASD Volume Performance
Volume Response Time	BAR	DASD Volume Performance

Highest Volume Group Response Time workspace		
View	**View type**	**Attribute group**
Highest Volume Performance Report	TABLE	DASD Volume Performance
Volume Response Time	BAR	DASD Volume Performance

Highest Volume Group I/O Rate workspace		
View	**View type**	**Attribute group**
Highest Volume I/O Report	TABLE	DASD Volume Performance
Volume I/O Rate	BAR	DASD Volume Performance
Volume Response Time	BAR	DASD Volume Performance

Highest Volume Group MPL workspace		
View	**View type**	**Attribute group**
Highest Volume MPL Report	TABLE	DASD Volume Performance
Volume MPL	BAR	DASD Volume Performance
Volume Response Time	BAR	DASD Volume Performance

Highest Volume Group Busy Percent workspace		
View	**View type**	**Attribute group**
Highest Volume Busy Percent Report	TABLE	DASD Volume Performance
Volume Busy Percent	BAR	DASD Volume Performance
Volume Response Time	BAR	DASD Volume Performance

Lowest Volume Group Connect Percent workspace		
View	**View type**	**Attribute group**
Volume Connect Time Percent	BAR	DASD Volume Performance
Lowest Volume Connect Time Percent Report	TABLE	DASD Volume Performance
Volume Response Time	BAR	DASD Volume Performance

Monitoring SMS storage groups space

This topic describes the workspaces that monitor the status of SMS Storage Groups Space. You access these workspaces through the SMS Storage Groups Space node in the Tivoli® Enterprise Portal.

See the "Monitoring the performance of user DASD groups" on page 70 section for information on how to add, edit, and delete a user DASD volume group.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

In addition, with OMEGAMON® for Storage on z/OS® Version 4.2.0 Interim Feature 3 or later, you can issue volume-level Storage Toolkit commands from the SMS Storage Group report level. To access these commands, open the SMS Storage Groups Space workspace from the Navigator, right-click a row on the table view and select the **Volume Actions...** menu item, then select the action you want to perform.

Navigating the workspaces of the SMS Storage Groups Space node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the SMS Storage Groups Space node to navigate to the workspaces that you see in "SMS Storage Groups Space node: bulleted list showing the hierarchy of linked workspaces" on page 67. You can access the workspaces in this table using the standard path and method for navigating workspaces.

SMS Storage Groups Space node: bulleted list showing the hierarchy of linked workspaces

SMS Storage Groups Space

- SMS Storage Group Volume Space
  - Dataset Space Summary
    - Dataset Details
  - Device Space Details
  - Volume Space Trend
- SMS Storage Group Space Trend
  - SMS Storage Group Volume Space History

These workspaces list the volumes with the least favorable readings for the attribute within the SMS Storage Group:

- Lowest Volume Free Space
- Lowest Volume Free Space Percent
- Highest Volume Fragmentation Index
- Lowest Volume Group Free Space
- Lowest Volume Group Free Space Percent
- Highest Volume Group Fragmentation Index
- Storage Group Configurations

Access the following linked workspaces by selecting the SMS Storage Groups Space node in the Navigator.

Right-click the node. Select the workspace that you want to access in the pop-up menu.

These workspaces list the volumes with the least favorable readings for the attribute within the z/OS® image.

- Lowest Volume Free Space
- Lowest Volume Free Space Percent
- Highest Volume Fragmentation Index

Each of the workspaces in the preceding bulleted list link to the following set of related workspaces:

- Dataset Space Summary
  - Dataset Details
- Device Space Details
- Volume Space Trend

SMS Storage Groups Space workspace

This topic describes the main components of the principal workspace for this node.

The direct situation analysis feature is available for this workspace. The feature can help you analyze and understand a problem that is displayed in this workspace. See “Directly monitoring situations” on page 133.

Views

This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

SMS storage groups space table view

The table view contains one row of data for each SMS storage group. The view contains the following information:

- The group name
- The storage group status
- The type and whether or not it is enabled
- The number of volumes in the group

- The status of the VTOC index
- The lowest freespace percent of all volumes in the group
- The highest fragmentation index of all volumes
- The largest free extent
- The total free space and the percent of free space for the group
- The total amount of space for the entire storage group

Storage groups with free space % < 10 bar graph

This graph contains one bar for each storage group with a combined free space percentage that falls below 10%.

Workspaces that are linked to this workspace

See "Navigating the workspaces of the SMS Storage Groups Space node" on page 66.
Learn how to navigate to workspaces here: "Workspaces" on page 14.

See "Workspaces" on page 14 for a list of all nodes provided in IBM® Z OMEGAMON® AI for Storage®.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the SMS Storage Groups Space node.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

See "SMS Storage Groups Space workspace" on page 67 for more information on the primary workspace for this node.

SMS Storage Groups Space workspace		
View	**View type**	**Attribute group**
SMS Storage Groups Space Summary	TABLE	Volume Group Summary
Storage Groups with Free Space % < 10	BAR	Volume Group Summary

SMS Storage Group Volume Space workspace		
View	**View type**	**Attribute group**
Volume Space Report	TABLE	DASD Volume Space
Volumes with Free Space % < 10	BAR	DASD Volume Space
Volumes with Fragmentation Index > 650	BAR	DASD Volume Space

Dataset Space Summary workspace		
View	**View type**	**Attribute group**
Dataset Space Summary Report	TABLE	Dataset Space
Volume Space Allocation	PIE	DASD Volume Space

Device Space Details workspace		
View	**View type**	**Attribute group**

Volume Free Space Details	TABLE	DASD Volume Space
Track-Managed Free Space Details	TABLE	DASD Volume Space
Volume Allocated Space Details	TABLE	DASD Volume Space
Volume Status	TABLE	Volume Status
Volume Details	TABLE	DASD Volume Performance

SMS Storage Group Space Trend workspace		
View	**View type**	**Attribute group**
SMS Storage Group Space Trend	TABLE	Volume Group Summary
MB of Free Space	BAR	Volume Group Summary

Volume Group Space History workspace		
View	**View type**	**Attribute group**
Volume Space Report	TABLE	DASD Volume Space
Volumes with Free Space % < 10	BAR	DASD Volume Space
Volumes with Fragmentation Index > 650	BAR	DASD Volume Space

Storage Group Percent Free workspace		
View	**View type**	**Attribute group**
Volume Free Space Percent	BAR	DASD Volume Space
Volume Space Report	TABLE	DASD Volume Space
Volume Fragmentation Index	BAR	DASD Volume Space

Storage Group Percent Free Trend workspace		
View	**View type**	**Attribute group**
Storage Group Free Space Percent	PLOT	Volume Group Summary
SMS Storage Group Space Trend	TABLE	Volume Group Summary

Dataset Space Event workspace		
View	**View type**	**Attribute group**
Volume Space Allocation	PIE	DASD Volume Space
Dataset Space Summary Report	TABLE	Dataset Space

Volume Space Trend workspace		
View	**View type**	**Attribute group**
Volume Space Trend	TABLE	DASD Volume Space
Percent Free Space	PLOT	DASD Volume Space

Lowest Volume Freespace workspace		
View	**View type**	**Attribute group**
Volume Free Space Report	TABLE	DASD Volume Space
Free Space	BAR	DASD Volume Space
Free Space Percent	BAR	DASD Volume Space

Lowest Volume Free Space Percent workspace		
View	**View type**	**Attribute group**
Volume Free Space Percent Report	TABLE	DASD Volume Space
Free Space Percent	BAR	DASD Volume Space
Free Space	BAR	DASD Volume Space

Highest Volume Fragmentation Index workspace		
View	**View type**	**Attribute group**
Volume Fragmentation Index Report	TABLE	DASD Volume Space
Fragmentation Index	BAR	DASD Volume Space
Free Space	BAR	DASD Volume Space

Lowest Volume Group Freespace workspace		
View	**View type**	**Attribute group**
Volume Free Space Report	TABLE	DASD Volume Space
Free Space	BAR	DASD Volume Space
Free Space Percent	BAR	DASD Volume Space

Lowest Volume Group Free Space Percent workspace		
View	**View type**	**Attribute group**
Volume Free Space Percent Report	TABLE	DASD Volume Space
Free Space Percent	BAR	DASD Volume Space
Free Space	BAR	DASD Volume Space

Highest Volume Group Fragmentation Index workspace		
View	**View type**	**Attribute group**
Volume Fragmentation Index Report	TABLE	DASD Volume Space
Fragmentation Index	BAR	DASD Volume Space
Free Space	BAR	DASD Volume Space

Storage Groups Configurations Summary workspace		
View	**View type**	**Attribute group**
Storage Group Configs Report	TABLE	Storage Group Configuration
Volume Models Report	TABLE	Volume Model
Volume Pools Report	TABLE	Volume Pool

Monitoring the performance of user DASD groups

This topic describes the workspaces that report the performance of user DASD groups.

A user DASD group is a user-defined logical grouping of DASD volumes. Volumes within a user DASD group are typically related according to specific criteria. You can define a user DASD group based on volume name, device address, SMS storage group, or one or more DASD volume attributes (for example, Fragmentation Index and Response Time). After you define a user DASD group, you can use the User DASD Group Performance and User DASD Group Space workspaces to view aggregated performance and space metrics for the group as a whole.

From these two workspaces, you can drill down to other workspaces that display performance and space metrics for each volume in the group.

You dynamically allocate user DASD volume groups. You access the user DASD workspaces through the User DASD Groups Performance node in the Tivoli® Enterprise Portal.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

In addition, with OMEGAMON® for Storage on z/OS®, you can issue volume-level Storage Toolkit commands from the User DASD Group report level. To access these commands, open the User DASD Groups Performance workspace from the Navigator, right-click a row on the table view and select the **Volume Actions...** menu item, then select the action you want to perform.

Navigating the workspaces of the User DASD Groups Performance node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the User DASD Groups Performance node to navigate to the workspaces that you see in “User DASD Groups Performance node: list showing the hierarchy of linked workspaces” on page 72. You can access the workspaces in this table using the standard path and method for navigating workspaces.

User DASD Groups Performance node: list showing the hierarchy of linked workspaces

User DASD Groups Performance

- User DASD Group Volume Performance
  - Dataset Performance Summary [1]
  - Physical Volume Group Performance
    - Dataset Performance Summary [1]
  - Device Performance Details
  - Cross System Volume Performance
  - Volume Performance Trend
    - Dataset Historic Performance Summary
      - Dataset Performance Historical Detail
- User DASD Group Volume Cache
  - Dataset Performance Summary [1]
  - Physical Volume Group Cache
    - Dataset Performance Summary [1]
  - Device Performance Details
  - Cross System Volume Performance
  - Volume Cache Trend
- Cross System Group Summary
  - Cross System Volume Performance
- User DASD Performance Trend
  - User DASD Group Historic Volume Performance
    - Dataset Historic Performance Summary
      - Dataset Performance Historical Detail
  - User DASD Group Historic Volume Cache Performance
    - Dataset Historic Performance Summary
      - Dataset Performance Historical Detail

The following workspaces display the volumes that have the least favorable readings for the attribute within the user DASD group selected for navigation:

- Highest User DASD Group Volumes Response Time (Top 10) [2]
- Highest User DASD Group Volumes MPL (Top 10) [2]
- Highest User DASD Group Volumes I/O Rate (Top 10) [2]
- Highest User DASD Group Volumes Busy Percent (Top 10) [2]
- Highest User DASD Group Volumes Connect Percent (Top 10) [2]

Footnotes for “User DASD Groups Performance node: list showing the hierarchy of linked workspaces” on page 72:

1. This group of workspaces is organized in the following nested pattern:
   - Dataset Performance Summary
     - Dataset Performance Detail
       - Dataset Details
       - Dataset Performance Detail Trend
     - Dataset Details
     - Dataset Performance Summary Trend
       - Dataset Performance Historical Detail

2. This workspace links to the following set of related workspaces:
    - Dataset Performance Summary (See the preceding note ([1]) for details.)
    - Physical Volume Group Performance
        - Dataset Performance Summary (See the preceding note ([1]) for details.)
        - Volume Performance Trend
            - Dataset Performance Summary History
    - Device Performance Details
    - Cross System Volume Performance
    - Volume Performance Trend
        - Dataset Historic Performance Summary
            - Dataset Performance Historical Detail

User DASD Groups Performance workspace

This topic describes the main components of the principal workspace for this node.

The following sections describe various components of this workspace.

Views

This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

User DASD groups performance table view

The table view contains one row of data for each defined user DASD group. Each row includes the following list entries:

- The name of the group
- The number of SMS and non-SMS volumes
- The total number of volumes in the group
- The current highest response time for any volume in the group
- The current high busy percent for any volume in the group
- The device multi-processing level
- The low read and write hit percentages among all volumes in the group
- The highest DFW retry percentage in the group

High volume response time > 25 bar graph

This graph contains one bar for each user DASD group with a highest volume response time that exceeds 25 milliseconds. The height of the bar indicates the actual response time for that volume.

Workspaces that are linked to this workspace

See “Navigating the workspaces of the User DASD Groups Performance node” on page 71.
Learn how to navigate among workspaces here: “Workspaces” on page 14.

See “Workspaces” on page 14 for a list of all nodes provided in IBM® Z OMEGAMON® AI for Storage®.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the User DASD Groups Performance node.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.

- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

See "User DASD Groups Performance workspace" on page 73 for more information on the primary workspace for this node.

User DASD Groups Performance workspace		
View	**View type**	**Attribute group**
User DASD Groups Performance Report	TABLE	Volume Group Summary
High Volume Response Time > 25	BAR	Volume Group Summary

User DASD Group Volume Performance workspace		
View	**View type**	**Attribute group**
Volume Performance Report	TABLE	DASD Volume Performance
Volume Response Time > 25	BAR	DASD Volume Performance

User DASD Group Volume Cache workspace		
View	**View type**	**Attribute group**
Volume Cache Report	TABLE	Cache Devices
Read Hit Percent < 55	BAR	Cache Devices
Write Hit Percent < 30	BAR	Cache Devices

User DASD Group Performance Trend workspace		
View	**View type**	**Attribute group**
User DASD Groups Performance Trend Report	TABLE	Volume Group Summary
High Volume Response Time	BAR	Volume Group Summary

Adding a user DASD group

Use the **Add User DASD Group** dialog box to add a user DASD volume group.

A dynamic user DASD volume group consists of a group of volumes constrained by at least one volume, device address, device range, SMS storage group, or volume attribute.

A volume is included in the group if it meets the criteria of at least one filter that is entered on the **Volumes** tab (if any are entered) and it meets the criteria of all filters that are entered on the **Attributes** tab.

Group tab

Use this tab to specify your user DASD volume group.

Name

Enter a name for the user DASD volume group that you are creating. Names may consist of any combination of characters and must not exceed 30 characters in length. Names are case-sensitive.

Description

Enter an optional description for the user DASD volume group. Descriptions may consist of any combination of characters and must not exceed 50 characters in length.

Persist

Select this checkbox to indicate that you want the user DASD volume group definition to persist (or to be saved) across recycles of the Tivoli® Enterprise Monitoring Server (monitoring server or TEMS). User DASD group definitions are discarded during a recycle of the monitoring server if the **Persist** option is not selected.

Paste from clipboard

Click this button and paste copied group definition in a special format to fill the group's fields automatically.

Volumes tab

Use this tab to specify the DASD devices of your user DASD volume group. Define the volume, device address, device range, and SMS storage group names in the following fields:

Volser or Pattern

Enter one or more volume names or volume masks. You can specify a partial name pattern and place the wildcard character (*) at the end. For example, a mask of PRI* includes all volumes with names that begin with "PRI".

First Device/Last Device

Enter one or more device addresses or device address ranges in the table to include volumes with those addresses in the user DASD volume group. Valid device addresses consist of three to four hexadecimal characters. The last device must be greater than the first device if a device address range is specified.

SMS Storage Group

Enter one or more SMS storage groups.

Attributes tab

Use this tab to specify the attribute constraints that are used in conjunction with the DASD device constraints for your user DASD volume group.
Use the drop-down list in the **Attribute** column to define the attribute name.

Use the drop-down list in the **Operator** column to select an operator.

Enter an appropriate value in the **Value** column or use the drop-down list to select an enumerated value.

Editing a user DASD group

Use this dialog box to edit a user DASD volume group.

For descriptions of the fields in the **Group**, **Volume**, and **Attributes** tabs see “Adding a user DASD group” on page 74.

Deleting a user DASD group

In the User DASD Groups Performance workspace, you can delete a user DASD volume group.

To delete a dynamically defined user DASD volume group, complete the following steps:

1. Navigate to the User DASD Groups Performance workspace.
2. Right-click the table row for the group that you want to delete.
3. From the menu that appears, select **Delete Group**.

The selected user DASD volume group is deleted.

Monitoring the space of user DASD groups

This topic describes the workspaces that report the status of user DASD groups space. This topic also describes the dialog boxes that you use to add, edit, and delete a user DASD volume group.

A user DASD group is a user-defined logical grouping of DASD volumes. Volumes within a user DASD group are typically related according to some criteria. You can define a user DASD group based on volume name, device address, SMS storage group, or one or more DASD volume attributes (for example, Fragmentation Index, Response Time, and other attributes).

After you define a user DASD group, you can use the User DASD Group Performance and User DASD Group Space workspaces to view aggregated performance and space metrics for the group as a whole. From these two workspaces, you can drill down to other workspaces that display performance and space metrics for each volume in the group.

You dynamically allocate user DASD volume groups. You access these workspaces through the User DASD Groups Space node in the Tivoli® Enterprise Portal.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

Additionally, with OMEGAMON® for Storage on z/OS®, you can issue volume-level Storage Toolkit commands from the User DASD Group report level.

To access these commands, open the User DASD Groups Space workspace from the Navigator, right-click a row on the table view and select the **Volume Actions...** menu item, then select the action you want to perform.

Navigating the workspaces of the User DASD Groups Space node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the User DASD Groups Space node to navigate to the workspaces that you see in "User DASD Groups Space node: bulleted list showing the hierarchy of linked workspaces" on page 76. You can access the workspaces in this table using the standard path and method for navigating workspaces.

User DASD Groups Space node: bulleted list showing the hierarchy of linked workspaces

User DASD Groups Space

- User DASD Group Volume Space
  - Dataset Space Summary
    - Dataset Details
  - Device Space Details
- User DASD Group Space Trend
  - User DASD Group Volume Historical Space

The following workspaces display the volumes that have the least favorable readings for the attribute within the User DASD Group:

- Lowest User DASD Group Volumes Free Space (Top 10)
- Lowest User DASD Group Volumes Free Space Percent (Top 10)
- Highest User DASD Group Volumes Fragmentation Index (Top 10)

User DASD Groups Space workspace

This topic describes the main components of the principal workspace for this node.

The following sections describe various components of this workspace.

Views

This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

User DASD groups space table view

The table view contains one row of data for each user DASD group. The view contains the following information:

- The group name
- The number of SMS and non-SMS volumes and the total number of volumes in the group
- The status of the VTOC index

- The lowest freespace percent of all volumes in the group
- The highest fragmentation index of all volumes
- The largest free extent
- The total free space
- The percentage of free space for the group and the total amount of space for the entire storage group

User DASD groups with free space % < 10 bar graph

This graph contains one bar for each user DASD group with a combined free space percentage that falls below 10%.

Workspaces that are linked to this workspace

See "Navigating the workspaces of the User DASD Groups Space node" on page 76.
See "Workspaces" on page 14 to learn how to navigate among workspaces and check the list of all nodes provided in IBM® Z OMEGAMON® AI for Storage®.

Issuing Group level Toolkit commands

This section describes a sample scenario for issuing a Group level Toolkit command to schedule weekly defragmentation of volumes based on specific criteria.

With Interim Feature 3 or later, you can issue volume-level Storage Toolkit commands from the User DASD Group report level. For example, consider a user who wants to defragment all volumes with a fragmentation index greater than 700 every Monday morning.

Complete the following steps to schedule weekly defragmentation of specific volumes:

1. Open the User DASD Groups Space workspace from the navigation tree.
2. Add a user DASD group with the attributes **Fragmentation Index**, the operator **> Greater than**, and the value **700** and click **OK**.

   Note: Specify these attributes on the **Attributes** tab of the **Add User DASD Group** dialog box.

3. Right-click the new row on the table view and select the **Volume Actions...Defrag** menu item.
4. On the **Schedule** tab, schedule the command to run at 1:00 am, every Monday.
5. Click **OK**. Volumes with a fragmentation index greater than 700 will be defragmented every Monday at 1:00 am.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the User DASD Groups Space node.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

See "User DASD Groups Space workspace" on page 76 for more information on the primary workspace for this node.

User DASD Groups Space workspace		
View	**View type**	**Attribute group**
User DASD Groups Space Report	TABLE	Volume Group Summary

User DASD Groups with Free Space % < 10	BAR	Volume Group Summary

User DASD Group Volume Space workspace		
View	**View type**	**Attribute group**
Volume Space Report	TABLE	DASD Volume Space
Volumes with Free Space % < 10	BAR	DASD Volume Space
Volumes with Fragmentation Index > 650	BAR	DASD Volume Space

User DASD Group Space Trend workspace		
View	**View type**	**Attribute group**
User DASD Groups Space Trend Report	TABLE	Volume Group Summary
MBs of Free Space	BAR	Volume Group Summary

DFSMShsm™ Status workspace		
View	**View type**	**Attribute group**
HSM Functions Status Report	TABLE	HSM Function Summary
Waiting Requests	BAR	HSM Function Summary
HSM Control Dataset Report	TABLE	HSM CDS
HSM Status Report	TABLE	HSM Status

Adding a user DASD group

Use **Add User DASD Groups Space** dialog box to add a user DASD volume group.

A volume is included in the group if it meets the criteria of at least one filter that is entered on the **Volumes** tab (if any are entered) and it meets the criteria of all filters that are entered on the **Attributes** tab.

Group tab

Use this tab to specify your user DASD volume group.

Name

Enter a name for the user DASD volume group that you are creating. Names may consist of any combination of characters and must not exceed 30 characters in length. Names are case-sensitive.

Description

Enter an optional description for the user DASD volume group. Descriptions may consist of any combination of characters and must not exceed 50 characters in length.

Persist

Select this checkbox to indicate that you want the user DASD volume group definition to persist (or to be saved) across recycles of the Tivoli® Enterprise Monitoring Server (monitoring server or TEMS). User DASD group definitions are discarded during a recycle of the monitoring server if the **Persist** option is not selected.

Volumes tab

Use this tab to specify the DASD devices of your user DASD volume group. Define the volume, device address, device range, and SMS storage group names in the following fields:

Volser or Pattern

Enter one or more volume names or volume masks. You can specify a partial name pattern and place the wildcard character (*) at the end. For example, a mask of PRI* includes all volumes with names that begin with "PRI".

First Device/Last Device

Enter one or more device addresses or device address ranges in the table to include volumes with those addresses in the user DASD volume group. Valid device addresses consist of three to four hexadecimal characters. The last device must be greater than the first device if a device address range is specified.

SMS Storage Group

Enter one or more SMS storage groups.

Attributes tab

Use this tab to specify the attribute constraints that are used in conjunction with the DASD device constraints for your user DASD volume group.
Use the drop-down list in the **Attribute** column to define the attribute name.

Use the drop-down list in the **Operator** column to select an operator.

Enter an appropriate value in the **Value** column or use the drop-down list to select an enumerated value.

Editing a user DASD group

Use this dialog box to edit a user DASD volume group.

For descriptions of the fields in the **Group**, **Volume**, and **Attributes** tabs, see "Adding a user DASD group" on page 78.

Deleting a user DASD group

In the User DASD Groups Space workspace, you can delete a user DASD volume group.

To delete a dynamically defined user DASD volume group, complete the following steps:

1. Navigate to the User DASD Groups Space workspace.
2. Right-click the table row for the group that you want to delete.
3. From the menu that appears, select **Delete Group**.

The selected user DASD volume group is deleted.

Monitoring user storage groups

This topic describes the User Groups for Storage Groups workspace that allows you to create custom groups for storage groups and review the summarized data in one place.

To access this workspace in the OMEGAMON® enhanced 3270 user interface (enhanced 3270UI), do the following:

1. Select the LPAR.
2. Select **DASD Storage and User Groups**.
3. Select **User Storage Groups**.

> **Note:** The User Groups for Storage Groups workspace is not available in the Tivoli® Enterprise Portal.

See "Workspaces" on page 14 for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

Linked workspaces

This section summarizes the workspaces available through User Groups for Storage Groups. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

User Groups for Storage Groups

- **Storage Groups for [Selected Storage Group Name]**
  - Storage Groups Space Metrics
    - **SMS Storage Group Volume Space**
      - Volume Space Report
        - **Dataset Space Summary**
          - Dataset Space Summary Report
            - **Dataset Details**
              - Dataset Space Attributes
              - Dataset Users
              - Dataset SMS Constructs
              - Dataset Performance Summary
  - Storage Groups Performance Metrics
    - **SMS Storage Group Volume Performance**
      - Volume Performance Report
      - Volume Cache Report
        - **Dataset Performance Summary**
          - Dataset Performance Summary Report
          - Datasets with MSR > Storage Class Objective
            - **Dataset Performance Detail**
              - Dataset Performance Detail Report
                - **Dataset Details**
                  - Dataset Space Attributes
                  - Dataset Users
                  - Dataset SMS Constructs
                  - Dataset Performance Summary

User Groups for Storage Groups workspace

The User Groups for Storage Groups workspace includes the table view for the User Storage Groups attribute group. For descriptions of the attributes used as column headings in the workspace table view, see S3 User Storage Groups.

The table view contains one row of data for each defined user storage group. Each row includes the following list entries:

- The name of the group
- The description of the group
- If the user group for SMS storage groups definition persists or not
- The Storage group count
- The total number of volumes in the group
- The size of the largest free extent
- The lowest free space value for a volume in the group
- The total amount of free space in the group
- The total amount of inactive space in the group
- The total amount of unused space in the group

- The total amount of track-managed free, inactive, and unused space in the group

Adding a user storage group

Use the **Add Group** dialog box to add a storage group to user storage groups.

A storage group is included in User Storage Groups if it meets the criteria of at least one filter that is entered (if any are entered) and if it meets the criteria of all filters that are entered in the **Attributes** tab.

To access the **Add Group** dialog box in User Storage Groups, type a before any name in the **Group Name** column, and press Enter.

Group tab

Use this tab to specify your user group for storage groups.

Name

Enter a name for the user group for storage groups that you are creating. Names may consist of any combination of characters and must not exceed 30 characters in length. Names are case-sensitive. The **Name** is required.

Description

Enter an optional description for the user group for storage groups. Descriptions may consist of any combination of characters and must not exceed 50 characters in length.

Persist

Select the **Persist** checkbox if you want the user group for storage groups definition to persist (or be saved) across recycles of the Tivoli Enterprise Monitoring Server (TEMS). User group definitions are discarded during a recycle of the monitoring server if the **Persist** field is not selected.

Properties tab

Use this tab to specify properties for your user group for storage groups.
Storage groups that match any of the properties that you specify will be included in your group, unless further constrained by the criteria you specify in the **Attributes** tab.

To specify properties, do the following:

1. In the **Attribute** column, type a before **n/a**, and press Enter.
2. Select one attribute from the **Attribute Name** list, and press Enter.

   > **Note:** You can select only one attribute at a time.

3. From the **Operator** drop-down list, select an operator.
4. Under **Value**, enter an appropriate value in the **Custom** field or select a value from the **Predefined** drop-down list.
5. Click **OK**.

Attributes tab

Use this tab to specify attribute constraints for your user group for storage groups.
Storage groups that match any of the properties and all of the attributes that you specify will be included in your group.

To specify attributes, do the following:

1. In the **Attribute** column, type a before **n/a**, and press Enter.
2. Select one attribute from the **Attribute Name** list, and press Enter.

   > **Note:** You can select only one attribute at a time.

3. From the **Operator** drop-down list, select an operator.

4. Under **Value**, enter an appropriate value in the **Custom** field or select a value from the **Predefined** drop-down list.
5. Click **OK**.

> **Important:** Property and attribute values cannot contain commas (,) or semi-colons (;).

Editing a user storage group

Use the **Edit Group** dialog box to edit a user storage group.

To access the **Edit Group** dialog box in User Storage Groups, type e before the name of the storage group that you want to edit, and press Enter.

You can edit all values except for the name of the user storage group.

For the descriptions of the **Group**, **Properties**, **Attributes** tabs, see Adding a user group for storage groups.

Deleting a user storage group

Use the **Delete Group** dialog box to delete a user storage group.

To access the **Delete Group** dialog box in User Storage Groups, type d before the name of the storage group that you want to delete, and press Enter.

After the **Delete Confirmation** pop-up window appears, press Enter to confirm the action.

Monitoring DFSMShsm™ status

This topic describes the workspaces that monitor the status of DFSMShsm™ activity. You access these workspaces through the DFSMShsm™ Status node in the Tivoli® Enterprise Portal.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

Navigating the workspaces of the DFSMShsm™ Status node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the DFSMShsm™ Status node to navigate to the workspaces that you see in “DFSMShsm™ Status node: bulleted list showing the hierarchy of linked workspaces” on page 83. You can access the workspaces in this table using the standard path and method for navigating workspaces.

DFSMShsm™ Status node: bulleted list showing the hierarchy of linked workspaces
DFSMShsm™ Status • DFSMShsm™ Request Details ◦ DFSMShsm™ Function Details • DFSMShsm™ Storage Usage • DFSMShsm™ Host Details ◦ CRQplex Details ◦ CRQplex Requests • HSMplex CRQplex Details ◦ CRQplex Requests • CRQplex Details ◦ CRQplex Requests

DFSMShsm™ Status workspace

This topic describes the main components of the principal workspace for this node.

When you click the DFSMShsm™ Status node in the Navigator of the Tivoli® Enterprise Portal, you see the DFSMShsm™ Status workspace. This workspace monitors DFSMShsm™ activity for the entire HSM subsystem, including utilization of volumes and CDS space. The main components of the workspace are as follows:

The direct situation analysis feature is available for this workspace. The feature can help you analyze and understand a problem that is displayed in this workspace. See "Directly monitoring situations" on page 133.

Views

This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

DFSMShsm™ Status workspace		
View	**Type**	**Attribute Group**
DFSMShsm™ Functions Summary	TABLE	HSM Function Summary
DFSMShsm™ Control Dataset Report	TABLE	HSM CDS
DFSMShsm™ Control Dataset Report (type File)	TABLE	HSM CDS
DFSMShsm™ Waiting Requests	BAR	HSM Host Request Summary
DFSMShsm™ Function Status	TABLE	HSM Host Function Status
DFSMShsm™ Status Report	TABLE	HSM Status

DFSMShsm™ Function Status Report

The tabular report in the top-center portion of this workspace contains one row for each DFSMShsm™ function. Each row displays:

- The status of that function for the reporting LPAR
- The number of dataset requests and volume requests for that function and how many are active and how many waiting
- The name of the HSMplex and the base name of the CRQplex that the HSM Host is a member of, or n/a, if the host is not a member of a CRQplex

DFSMShsm™ Control Dataset Report

The tabular report covering the entire width of the workspace in the vertical center contains one row for each of the DFSMShsm™ control datasets. Each row displays:

- The DD name of the dataset

- The number of extents for the data and index components
- The percent of Available Space for the data and index components of the dataset
- The percent of Free Space for the data and index components of the dataset
- The total size of the data and the index components of the dataset

DFSMShsm™ Status Report

The tabular report covering the entire width of the workspace at the bottom of the screen contains a single row of data describing DFSMShsm™ status in the reporting LPAR. This row indicates:

- Whether DFSMShsm™ is active
- The address space ID
- The version of DFHSM
- DFSMShsm™ start time
- The interval start time
- The number of DFSMShsm™ requests during this collection interval
- The CPU usage during this interval
- The number of requests that abended during this interval
- The number of volume migrations, backups and the number of migration and backup volumes recycled during this interval
- The name of the HSMplex and an indication of whether at least one HSM host on the z/OS® image is in a CRQplex

Waiting requests bar graph

This graph contains one bar for each DFSMShsm™ function which has at least one waiting request. The height of the bar indicates the total number of waiting requests for this function.

Workspaces that are linked to this workspace

See "Navigating the workspaces of the DFSMShsm Status node" on page 82.
Learn how to navigate among workspaces here: "Workspaces" on page 14.

See "Workspaces" on page 14 for a list of all nodes provided in IBM® Z OMEGAMON® AI for Storage®.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the DFSMShsm™ Status node.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

See "DFSMShsm Status workspace" on page 83 for more information on the primary workspace for this node.

DFSMShsm™ Status workspace		
View	**View type**	**Attribute group**
DFSMShsm™ Functions Summary	TABLE	HSM Function Summary
DFSMShsm™ Control Dataset Report	TABLE	HSM CDS
DFSMShsm™ Waiting Requests	BAR	HSM Host Request Summary
DFSMShsm™ Function Status	TABLE	HSM Host Function Status
DFSMShsm™ Status Report	TABLE	HSM Status

DFSMShsm™ Request Details workspace		
View	**View type**	**Attribute group**
DFSMShsm™ Function Statistics	TABLE	HSM Function Statistics
DFSMShsm™ Requests	TABLE	HSM Requests

DFSMShsm™ Storage Usage workspace		
View	**View type**	**Attribute group**
DFSMShsm™ Host Storage Utilization	TABLE	HSM Private Storage
DFSMShsm™ Host Frame Usage	BAR	HSM Private Storage
DFSMShsm™ Common Storage	TABLE	HSM Common Storage
DFSMShsm™ Common Storage Utilization	BAR	HSM Common Storage

DFSMShsm™ Host Details workspace		
View	**View type**	**Attribute group**
DFSMShsm™ Host Details Report	TABLE	HSM Host Status
DFSMShsm™ Requests	BAR	HSM Host Status
DFSMShsm™ Failed Requests	BAR	HSM Host Status

DFSMShsm™ Function Details workspace		
View	**View type**	**Attribute group**
DFSMShsm™ Host Function Details	TABLE	HSM Function Details
DFSMShsm™ Function Response Time	BAR	HSM Function Details

HSMplex CRQplex Details workspace		
View	**View type**	**Attribute group**
HSMplex CRQplex Host Details Report	TABLE	HSM Cross System CRQ Hosts
HSMplex CRQplex Details Report	TABLE	HSM Cross System CRQplex
CRQplex Usage	BAR	HSM Cross System CRQplex
CRQplex Requests	BAR	HSM Cross System CRQplex

CRQplex Details workspace		
View	**View type**	**Attribute group**
CRQplex Host Details Report	TABLE	HSM Cross System CRQ Hosts
CRQplex Details Report	TABLE	HSM Cross System CRQplex
CRQplex Usage	BAR	HSM Cross System CRQplex
CRQplex Requests	BAR	HSM Cross System CRQplex

CRQplex Requests workspace		
View	**View type**	**Attribute group**

CRQplex Details Report	TABLE	HSM Cross System CRQplex
CRQplex Usage	TABLE	HSM Cross System CRQplex
CRQplex Requests	BAR	HSM Cross System CRQplex
CRQplex Requests Report	BAR	HSM CRQ Requests

Monitoring Tape Management status

This topic introduces the workspaces that monitor the status of DFSMSrmm™ and CA1 resources. You access these workspaces through Tape Management Status node in the Tivoli® Enterprise Portal.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

Navigating the workspaces of the Tape Management Status node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources. See “Workspaces” on page 14 to learn more about navigating from workspace to workspace.

You access the primary workspaces of this node as follows:

1. Click the node in the Tivoli® Enterprise Portal. The node is activated and the first workspace is displayed in the portal.
2. Right-click the node to display the pop-up menu.
3. Select the **Workspace** submenu to see a list of the available linked workspaces. These linked workspaces, in turn, link to many other related workspaces.

You can access the workspaces using the standard path and method for navigating workspaces.

Monitoring DFSMSrmm™ status

This topic describes the workspaces that monitor the status of DFSMSrmm™ resources.

You access these workspaces through the Tape Management Status node in the Tivoli® Enterprise Portal.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

Navigating the workspaces of the DFSMSrmm™ Status node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources. See “Workspaces” on page 14 to learn more about navigating from workspace to workspace.

You access the primary workspaces of this node as follows:

1. Click the node in the Tivoli® Enterprise Portal. The node is activated and the first workspace is displayed in the portal.

2. Right-click the node to display the pop-up menu.
3. Select the **Workspace** submenu to see a list of the available linked workspaces. These linked workspaces, in turn, link to many other, related workspaces as shown in “DFSMSrmm™ Status node: bulleted list showing the hierarchy of linked workspaces” on page 88.

Workspaces that are available in the "Workspaces" pop-up menu of the Navigator node:

DFSMSrmm™ Location Summary
DFSMSrmm™ Owner Summary
DFSMSrmm™ Account Summary
DFSMSrmm™ Job Summary
DFSMSrmm™ Program Summary
DFSMSrmm™ VRS Summary
DFSMSrmm™ Error Summary
DFSMSrmm™ Pending Actions
DFSMSrmm™ Configuration

You start in the “DFSMSrmm System Summary workspace” on page 90 to access all the other workspaces that are listed in “DFSMSrmm™ Status node: bulleted list showing the hierarchy of linked workspaces” on page 88. You can access the workspaces in this table using the standard path and method for navigating workspaces.

Note: The following italic phrase is repeated in “DFSMSrmm™ Status node: bulleted list showing the hierarchy of linked workspaces” on page 88: "*DFSMSrmm™ Dataset Details workspaces are provided here.*" A set of standard workspaces is available where you see this phrase in the workspace hierarchy. The phrase includes a link to a description of the set of workspaces, which detail the volume, owner, VRS, and SMS construct of the relevant dataset.

DFSMSrmm™ Status node: bulleted list showing the hierarchy of linked workspaces

DFSMSrmm™ Status

- DFSMSrmm™ Location Summary
  - DFSMSrmm™ Location Details
  - DFSMSrmm™ Location Summary Volumes
    - DFSMSrmm™ Volume Details
      - DFSMSrmm™ Realtime Datasets
        - *DFSMSrmm™ Dataset Details workspaces are provided here.*
    - DFSMSrmm™ Realtime Datasets
      - *DFSMSrmm™ Dataset Details workspaces are provided here.*
  - DFSMSrmm™ Location Realtime Volumes
    - DFSMSrmm™ Volume Details
      - DFSMSrmm™ Realtime Datasets
        - *DFSMSrmm™ Dataset Details workspaces are provided here.*
    - DFSMSrmm™ Owner Details
    - DFSMSrmm™ Realtime Datasets
      - *DFSMSrmm™ Dataset Details workspaces are provided here.*
- DFSMSrmm™ Owner Summary
  - DFSMSrmm™ Owner Details
  - DFSMSrmm™ Owner Summary Volumes
    - DFSMSrmm™ Volume Details
      - DFSMSrmm™ Realtime Datasets
        - *DFSMSrmm™ Dataset Details workspaces are provided here.*
    - DFSMSrmm™ Realtime Datasets
      - *DFSMSrmm™ Dataset Details workspaces are provided here.*
  - DFSMSrmm™ Owner Realtime Volumes
    - DFSMSrmm™ Volume Details
      - DFSMSrmm™ Realtime Datasets
        - *DFSMSrmm™ Dataset Details workspaces are provided here.*
      - DFSMSrmm™ Owner Details
    - DFSMSrmm™ Realtime Datasets
      - *DFSMSrmm™ Dataset Details workspaces are provided here.*

- DFSMSrmm™ Account Summary
  - DFSMSrmm™ Account Summary Volumes
    - DFSMSrmm™ Volume Details
      - DFSMSrmm™ Realtime Datasets
        - *DFSMSrmm™ Dataset Details workspaces are provided here.*
    - DFSMSrmm™ Realtime Datasets
      - *DFSMSrmm™ Dataset Details workspaces are provided here.*
- DFSMSrmm™ Job Summary
  - DFSMSrmm™ Job Summary Datasets
    - *DFSMSrmm™ Dataset Details workspaces are provided here.*
    - DFSMSrmm™ Volume Details
      - DFSMSrmm™ Realtime Datasets
        - *DFSMSrmm™ Dataset Details workspaces are provided here.*
      - DFSMSrmm™ Owner Details

- DFSMSrmm™ Program Summary
  - DFSMSrmm™ Program Summary Datasets
    - *DFSMSrmm™ Dataset Details workspaces are provided here.*
    - DFSMSrmm™ Volume Details
      - DFSMSrmm™ Realtime Datasets
        - *DFSMSrmm™ Dataset Details workspaces are provided here.*
      - DFSMSrmm™ Owner Details
- DFSMSrmm™ VRS Summary
  - DFSMSrmm™ VRS Summary Datasets
    - *DFSMSrmm™ Dataset Details workspaces are provided here.*
    - DFSMSrmm™ Volume Details
      - DFSMSrmm™ Realtime Datasets
        - *DFSMSrmm™ Dataset Details workspaces are provided here.*
      - DFSMSrmm™ Owner Details
  - DFSMSrmm™ VRS Details

- DFSMSrmm™ Error Summary
  - DFSMSrmm™ Volume Details
    - DFSMSrmm™ Realtime Datasets
      - *DFSMSrmm™ Dataset Details workspaces are provided here.*
    - DFSMSrmm™ Owner Details
- DFSMSrmm™ Pending Actions
  - DFSMSrmm™ Pending Action Volumes
    - DFSMSrmm™ Realtime Datasets
      - *DFSMSrmm™ Dataset Details workspaces are provided here.*
    - DFSMSrmm™ Volume Details
      - DFSMSrmm™ Realtime Datasets
        - *DFSMSrmm™ Dataset Details workspaces are provided here.*
      - DFSMSrmm™ Owner Details
  - DFSMSrmm™ Pending Move Volumes
    - DFSMSrmm™ Realtime Datasets
      - *DFSMSrmm™ Dataset Details workspaces are provided here.*
    - DFSMSrmm™ Volume Details
      - DFSMSrmm™ Realtime Datasets
        - *DFSMSrmm™ Dataset Details workspaces are provided here.*
      - DFSMSrmm™ Owner Details

- DFSMSrmm™ Volume Search
  - DFSMSrmm™ Volume Details
    - DFSMSrmm™ Realtime Datasets
      - *DFSMSrmm™ Dataset Details workspaces are provided here.*
  - DFSMSrmm™ Realtime Datasets
    - *DFSMSrmm™ Dataset Details workspaces are provided here.*
- DFSMSrmm™ Dataset Search
  - *DFSMSrmm™ Dataset Details workspaces are provided here.*
  - DFSMSrmm™ Volume Details
    - DFSMSrmm™ Realtime Datasets
      - *DFSMSrmm™ Dataset Details workspaces are provided here.*
- DFSMSrmm™ Configuration
  - Dataset Details

DFSMSrmm™ System Summary workspace

This topic describes the main components of the principal workspace for this node.

When you click the DFSMSrmm™ Status node in the Navigator of the Tivoli® Enterprise Portal, you see the DFSMSrmm™ System Summary workspace. The main components of the workspace are listed here.

Note: DFSMSrmm Data collection is turned off by default, so this workspace is initially empty. Additionally, if your installation does not use DFSMSrmm, this workspace is always empty.

Views

This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

DFSMSrmm Control Dataset Space bar graph

This graph contains bars whose height represents the percentage utilization of the Control Dataset data component and the Journal, as reported by DFSMSrmm.

DFSMSrmm Pending Action Summary Volume Counts bar graph

This graph contains bars representing the total number of volumes in each of the possible DFSMSrmm Pending Action states: Scratch, Replace, Return, Init, Erase and Notify.

DFSMSrmm Volume Type Distribution Summary pie graph

This graph shows the number of non-scratch volumes of each type: Master, User, Init and Entry.

DFSMSrmm Volume Utilization Summary bar graph

This graph contains shows the distribution of volumes based on their percentage utilization. The height of each bar represents the number of volumes in that utilization range (for example, utilization between 21% and 30%).

DFSMSrmm System Summary Total table view

This table contains a single row that summarizes the entire DFSMSrmm environment and contains the following attributes. Notice that many of these same attributes are shown on the linked-to Summary workspaces described below. In those cases the attribute values are relative to the specific object that you are summarizing (for example, a location, an owner, a creating job, and so on). The attribute values are not relative to the entire DFSMSrmm environment:

- Collection Status indicates whether a DFSMSrmm collection cycle has been run; if it was run to completion or was stopped, and so on.
- Summary Collection Start Time - the time when the last collection cycle started. All the summary data displayed on this and other summary workspaces is based on the contents of the DFSMSrmm CDS between the indicated Collection Start Time and Collection End Time.
- Summary Collection End Time
- Total Volumes - the total number of volumes found in the CDS.
- Total Datasets - the total number of datasets contained on all volumes.
- Physical Volumes - the number of volumes that are neither Logical nor Stacked.
- Scratch Volumes
- Master Volumes
- User Volumes
- Init Volumes
- Entry Volumes
- New Scratch Volumes - the number of volumes returned to scratch status during or since the last DFSMSrmm expiration processing.
- Total Errors - the total number of temporary and permanent read and write errors on all volumes.
- Total Permanent Errors
- Total Temporary Errors

DFSMSrmm System Usage and Activity Summary table view

This table contains a single row that summarizes the entire DFSMSrmm environment, and contains the following attributes. Some of these attributes can be calculated only when you have selected the **Summarize RMM Dataset Attributes** option in the "RMM Collection Control dialog box" on page 223. Notice that many of these same attributes are shown on the linked-to Summary workspaces described below. In those cases the attribute values are relative to the specific object that you are summarizing (for example, a location, an owner, a creating job, and so on). The attribute values are not relative to the entire DFSMSrmm environment:

- Total Capacity - the capacity of all volumes in GB.
- Total Usage - the total usage on all volumes in GB.

- Average Volume Utilization - the average utilization across all volumes, in percent.
- Average Compression Ratio - the calculated percentage compression achieved between the data written by the application and the data recorded on the tape media.
- Average Volume Use Count - the average number of times each volume was read or written.
- Volumes R/W Previous Day - the number of volumes that were read or written on the previous Julian day.
- % Volumes R/W Previous Day - the percentage of the total volumes that were read or written on the previous Julian day.
- Average Datasets per Volumes - the average number of datasets contained on each tape volume.
- Average Dataset Size - the average size of a dataset, in MB to 1 decimal place.
- Datasets R/W Previous Day - the number of datasets that were read or written on the previous Julian day.
- Datasets Created Previous Day - the number of datasets that were created on the previous Julian day.
- Data Created Previous Day - the amount of data that was created on the previous Julian day in all newly created datasets, expressed as MB to 1 decimal place.

Workspaces that are linked to this workspace

See "Navigating the workspaces of the DFSMSrmm Status node" on page 86.
See "Workspaces" on page 14 to check a list of all nodes provided in IBM® Z OMEGAMON® AI for Storage® and learn how to navigate among workspaces.

See "DFSMSrmm Status feature" on page 219 for additional information.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the DFSMSrmm™ Status node.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

This workspace monitors DFSMSrmm™ activity for the entire RMM subsystem, including utilization of volumes and CDS space.

DFSMSrmm™ Status workspace		
View	**View type**	**Attribute group**
DFSMSrmm™ Control Dataset Space	BAR	RMM Control Dataset
DFSMSrmm™ Pending Action Summary Volume Counts	BAR	RMM Summary
DFSMSrmm™ Volume Utilization Summary	BAR	RMM Summary
DFSMSrmm™ Volume Type Distribution Summary	PIE	RMM Summary
DFSMSrmm™ System Summary Totals	TABLE	RMM Summary
DFSMSrmm™ System Usage & Activity Summary	TABLE	RMM Summary

This workspace contains views displaying summary attributes for all Locations that were found during the last Summary Collection cycle:

DFSMSrmm™ Location Summary workspace		
View	**View type**	**Attribute group**
DFSMSrmm™ Location Summary Totals	TABLE	RMM Summary
DFSMSrmm™ Location Usage & Activity Summary	TABLE	RMM Summary
DFSMSrmm™ Top 10 Locations by Total Volumes	BAR	RMM Summary

This workspace contains both static and dynamic values from the control dataset (CDS), which is a dataset that contains configuration, operational, and communication information for the DFSMSrmm™ functions.

DFSMSrmm™ Configuration workspace		
View	**View type**	**Attribute group**
DFSMSrmm™ Control Dataset	TABLE	RMM Control Dataset
DFSMSrmm™ Control Dataset Space	BAR	RMM Control Dataset
DFSMSrmm™ Days Since Last Utility Processing	BAR	RMM Config
DFSMSrmm™ Configuration Options	TABLE	RMM Config
DFSMSrmm™ Configuration State	TABLE	RMM Config

This workspace displays the top volumes with non-zero temporary and permanent error counts:

DFSMSrmm™ Error Summary workspace		
View	**View type**	**Attribute group**
DFSMSrmm™ Top 100 Temporary Error Volumes	TABLE	RMM Summary
DFSMSrmm™ Top 100 Permanent Error Volumes	TABLE	RMM Summary
DFSMSrmm™ Top 10 Permanent Error Volumes	BAR	RMM Summary
DFSMSrmm™ Top 10 Temporary Error Volumes	BAR	RMM Summary

DFSMSrmm™ Realtime Datasets workspace		
View	**View type**	**Attribute group**
DFSMSrmm™ Datasets	TABLE	RMM Dataset
DFSMSrmm™ Dataset Block Count	BAR	RMM Dataset

This workspace displays attributes for all actions. The metrics that are displayed for each action status correspond to the last collection cycle:

DFSMSrmm™ Pending Actions workspace		
View	**View type**	**Attribute group**
DFSMSrmm™ Action Status	TABLE	RMM Action
DFSMSrmm™ Move Status	TABLE	RMM Action
DFSMSrmm™ Action Summary Volume Counts	BAR	RMM Summary

This workspace contains views displaying summary attributes for all Owners that were found during the last Summary Collection cycle:

DFSMSrmm™ Owner Summary workspace		
View	**View type**	**Attribute group**
DFSMSrmm™ Top 10 Owners by Total Volumes	BAR	RMM Summary
DFSMSrmm™ Owner Usage & Activity Summary	TABLE	RMM Summary
DFSMSrmm™ Owner Summary Totals	TABLE	RMM Summary

This workspace contains views displaying summary attributes for all accounts that were found during the last Summary Collection cycle:

DFSMSrmm™ Account Summary workspace		
View	**View type**	**Attribute group**
DFSMSrmm™ Account Summary Activity	TABLE	RMM Summary
DFSMSrmm™ Account Summary Totals	TABLE	RMM Summary
DFSMSrmm™ Top 10 Summary Accounts by Total Volumes	BAR	RMM Summary

This workspace displays all attributes for all DFSMSrmm™ Locations.

DFSMSrmm™ Location Details workspace		
View	**View type**	**Attribute group**
DFSMSrmm™ Location Details	TABLE	RMM Location
DFSMSrmm™ Location Summary Volumes	BAR	RMM Summary

This workspace contains views displaying summary attributes for all Job objects that were found during the last Summary Collection cycle:

DFSMSrmm™ Job Summary workspace		
View	**View type**	**Attribute group**
DFSMSrmm™ Job Summary	TABLE	RMM Summary
DFSMSrmm™ Top Jobs by Total Datasets	BAR	RMM Summary

This workspace contains views displaying summary attributes for all Program objects that were found during the last Summary Collection cycle:

DFSMSrmm™ Program Summary workspace		
View	**View type**	**Attribute group**
DFSMSrmm™ Program Summary	TABLE	RMM Summary
DFSMSrmm™ Top Programs by Total Datasets	BAR	RMM Summary

This workspace contains views displaying summary attributes for all VRS objects that were found during the last Summary Collection cycle:

DFSMSrmm™ VRS Summary workspace		
View	**View type**	**Attribute group**
DFSMSrmm™ Top 10 VRS by Total Datasets	BAR	RMM Summary

DFSMSrmm™ VRS Summary	TABLE	RMM Summary

This workspace displays the volumes that are in the selected action state.

DFSMSrmm™ Pending Action Volumes workspace		
View	**View type**	**Attribute group**
DFSMSrmm™ Pending Action Volumes	TABLE	RMM Volume
DFSMSrmm™ Pending Action Summary Volume Counts	BAR	RMM Summary

This workspace displays a subset of attributes for a list of volumes that were obtained during the last summary collection cycle. These volumes are associated with a specific object that was summarized in another workspace.

DFSMSrmm™ Summary Volumes workspace		
View	**View type**	**Attribute group**
DFSMSrmm™ Summary Volumes	TABLE	RMM Volume
DFSMSrmm™ Summary Volume Total	BAR	RMM Summary

This workspace displays selected attributes for a list of datasets.

DFSMSrmm™ Dataset Search workspace		
View	**View type**	**Attribute group**
DFSMSrmm™ Dataset Block Count	BAR	RMM Dataset
DFSMSrmm™ Datasets	TABLE	RMM Dataset

DFSMSrmm™ Volume Search workspace		
View	**View type**	**Attribute group**
DFSMSrmm™ Permanent Volume Errors > 0	BAR	RMM Volume
DFSMSrmm™ Volumes	TABLE	RMM Volume

DFSMSrmm™ Summary Datasets workspace		
View	**View type**	**Attribute group**
DFSMSrmm™ Summary Datasets	TABLE	RMM Volume
DFSMSrmm™ Summary Dataset Total	BAR	RMM Summary

This workspace displays the volumes that are in the move state.

DFSMSrmm™ Pending Move Volumes workspace		
View	**View type**	**Attribute group**
DFSMSrmm™ Pending Move Volumes	TABLE	RMM Volume
DFSMSrmm™ Pending Action Summary Volume Counts	BAR	RMM Summary

This workspace displays a filler set of attributes for one or more volumes obtained in real-time from the DFSMSrmm™ facility. The list of volumes that this workspace displays is associated with a selected Owner or Location, with a selected Pending Action, or is specified through the pop-up menu the DFSMSrmm™ Volume Search link.

DFSMSrmm™ Realtime Volumes workspace		
View	**View type**	**Attribute group**

DFSMSrmm™ Permanent Volume Errors > 0	BAR	RMM Volume
DFSMSrmm™ Realtime Volumes	TABLE	RMM Volume

DFSMSrmm™ Dataset Details workspaces

The workspaces in this set of workspaces provide information about a dataset and its volume, owner, VRS, and SMS construct. The set workspaces is provided in many places in the hierarchy of workspaces, and represents the deepest level of detail that is available.

This workspace displays all attributes for a single DFSMSrmm™ dataset.

DFSMSrmm™ Dataset Details workspace		
View	**View type**	**Attribute group**
DFSMSrmm™ Dataset Information	TABLE	RMM Dataset
DFSMSrmm™ Dataset Usage	BAR	RMM Dataset
DFSMSrmm™ Datasets Statistics	TABLE	RMM Dataset
DFSMSrmm™ Dataset Dates	TABLE	RMM Dataset
DFSMSrmm™ Dataset VRS	TABLE	RMM Dataset

This workspace displays all attributes for a single DFSMSrmm™ volume.

DFSMSrmm™ Volume Details workspace		
View	**View type**	**Attribute group**
DFSMSrmm™ Volume Information	TABLE	RMM Volume
DFSMSrmm™ Volume Date Information	TABLE	RMM Volume
DFSMSrmm™ Volume Statistics	TABLE	RMM Volume
DFSMSrmm™ Volume Storage and Movement	TABLE	RMM Volume
Volume Errors	BAR	RMM Volume

This workspace displays all attributes for a single DFSMSrmm™ Owner.

DFSMSrmm™ Owner Details workspace		
View	**View type**	**Attribute group**
DFSMSrmm™ Owner Details	TABLE	RMM Owner
DFSMSrmm™ Owner Volumes	BAR	RMM Owner

This workspace displays all attributes for a single DFSMSrmm™ Vital Record Specification.

DFSMSrmm™ VRS Details workspace		
View	**View type**	**Attribute group**
DFSMSrmm™ VRS Details	TABLE	RMM VRS
DFSMSrmm™ VRS Priority	BAR	RMM VRS
DFSMSrmm™ VRS Count	BAR	RMM VRS

Monitoring CA 1 status

This topic describes the workspaces that monitor the status of CA1 resources.

You access these workspaces through Tape Management Status node in the Tivoli® Enterprise Portal.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

Navigating the workspaces of the CA 1 Status node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources. See "Workspaces" on page 14 to learn more about navigating from workspace to workspace.

You access the primary workspaces of this node as follows:

1. Click the node in the Tivoli® Enterprise Portal. The node is activated and the first workspace is displayed in the portal.
2. Right-click the node to display the pop-up menu.
3. Select the **Workspace** submenu to see a list of the available linked workspaces. These linked workspaces, in turn, link to many other, related workspaces as shown in "CA 1 Status node: bulleted list showing the hierarchy of linked workspaces" on page 97.

You start in the "CA 1 Status workspace" on page 97 to access all the other workspaces that are listed in "CA 1 Status node: bulleted list showing the hierarchy of linked workspaces" on page 97. You can access the workspaces in this table using the standard path and method for navigating workspaces.

CA 1 Status node: bulleted list showing the hierarchy of linked workspaces

- CA 1 Summary
- CA 1 Job Summary
  - CA 1 Summary Datasets
    - CA 1 Dataset Details
    - CA 1 Volume Details
      - CA 1 Volume Datasets
      - CA 1 Volume Search
      - CA 1 Dataset Search
- CA 1 Program Summary
  - CA 1 Summary Datasets
    - CA 1 Dataset Details
    - CA 1 Volume Details
      - CA 1 Volume Datasets
      - CA 1 Volume Search
      - CA 1 Dataset Search
- CA 1 Error Summary
  - CA 1 Volume Details
    - CA 1 Volume Datasets
    - CA 1 Volume Search
    - CA 1 Dataset Search

CA 1 Status workspace

This topic describes the main components of the principal workspace for this node.

When you click the CA 1 Status node in the Navigator of the Tivoli® Enterprise Portal, you see the CA 1 Status workspace. The main components of the workspace are listed here.

Note: CA 1 Data collection is turned off by default, so this workspace is initially empty. In addition, if your installation does not use CA 1, this workspace is always empty.

Views

This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

CA 1 Scratch Volumes pie graph

This graph shows the number of total volumes and total scratch volumes.

CA 1 DSNB Allocation pie graph

This graph shows the number of total DSNB allocations and the active DSNBs.

CA 1 Volume Utilization bar graph

This graph contains shows the volume utilization. The height of each bar represents the percentage utilized range (for example, utilization between 21% and 30%).

CA 1 System Summary table view

This table contains a single row that summarizes performance information for the CA1 tape management system, and contains the following attributes. Notice that many of these same attributes are shown on the linked-to Summary workspaces described below. In those cases the attribute values are relative to the specific object that you are summarizing (for example, a location, an owner, a creating job, and so on). The attribute values are not relative to the entire CA 1 environment:

- Average Compression Ratio
- Average Datasets Per Volume - The average number of datasets per volume.
- Average Dataset Size - The average dataset size, in MB to 1 decimal place.
- Average Dataset Usage Percent
- Average Volume Open Count - The average number of times the tape was opened since the birth date.
- Average Volume Usage Percent
- Collection End Time - The end time of data collection.
- Collection Start Time - The start time of data collection.
- Collection Status - The status of the data collector.
- Managed System - The SMS name of this system.
- Perm Read Errors Since Clean - The number of permanent read errors since clean.
- Perm Read Errors Since Init - The number of permanent read errors since initialization.
- Perm Write Errors Since Clean - The number of permanent write errors since clean.
- Perm Write Errors Since Init - The number of permanent write errors since initialization.
- Temp Read Errors Since Clean - The number of temporary read errors since clean.
- Temp Read Errors Since Init - The number of temporary read errors since initialization.
- Temp Write Errors Since Clean - The number of temporary write errors since clean.
- Temp Write Errors Since Init - The number of temporary write errors since initialization.
- Total Datasets - The total number of datasets.
- Total Errors - The total number of errors.
- Total New Scratch Volumes - The total number of new (uninitialized) scratch volumes.
- Total Perm Errors - The total number of permanent errors (since initialization or clean).
- Total Perm Write Errors - The total number of permanent write errors (since initialization or clean).

- Total Perm Read Errors - The total number of permanent read errors (since initialization or clean).
- Total RDS Override Volumes - The total number of volumes eligible for RDS Override.
- Total Scratch Volumes - The total number of scratch volumes.
- Total Temp Errors - The total number of temporary errors (since initialization or clean).
- Total Temp Read Errors - The total number of temporary read errors (since initialization or clean).
- Total Temp Write Errors - The total number of temporary write errors (since initialization or clean).
- Total Usage - The amount of uncompressed data written to the tape, expressed as GB.
- Total Vol Perm Errors - The total number of volumes with permanent errors.
- Total Vol Perm Read Errors - The total number of volumes with permanent read errors (since initialization or clean).
- Total Vol Perm Write Errors - The total number of volumes with permanent write errors (since initialization or clean).
- Total Vol Temp Errors - The total number of volumes with temporary errors (since initialization or clean).
- Total Vol Temp Read Errors - The total number of volumes with temporary read errors (since initialization or clean).
- Total Vol Temp Write Errors - The total number of volumes with temporary write errors (since initialization or clean).
- Total Volume Errors - The total number of volumes with errors.
- Total Volumes - The total number of active volumes.
- Total Volumes Closed by ABEND - The total number of volumes closed by ABEND.
- Total Volumes Closed Normally - The total number of volumes closed normally.
- Total Volumes in Delete Status - The total number of volumes in delete status.
- Total Volumes To Be Cleaned - The total number of volumes listed to be cleaned.
- Vol Perm Read Errors Since Clean - The number of volumes with permanent read errors since clean.
- Vol Perm Read Errs Since Init - The number of volumes with permanent read errors since initialization.
- Vol Perm Write Errors Since Clean - The number of volumes with permanent write errors since clean.
- Vol Perm Write Errors Since Init - The number of volumes with permanent write errors since initialization.
- Vol Temp Read Errs Since Clean - The number of volumes with temporary read errors since clean.
- Vol Temp Read Errors Since Init - The number of volumes with temporary read errors since initialization.
- Vol Temp Write Errors Since Clean - The number of volumes with temporary write errors since clean.
- Vol Temp Write Errs Since Init - The number of volumes with temporary write errors since initialization.
- Volume 0% Utilized - The number of volumes not being utilized.
- Volume 1-10% Utilized - The number of volumes being utilized 1-10%.
- Volume 11-20% Utilized - The number of volumes being utilized 11-20%.
- Volume 21-30% Utilized - The number of volumes being utilized 21-30%.
- Volume 31-40% Utilized - The number of volumes being utilized 31-40%.

- Volume 41-50% Utilized - The number of volumes being utilized 41-50%.
- Volume 51-60% Utilized - The number of volumes being utilized 51-60%.
- Volume 61-70% Utilized - The number of volumes being utilized 61-70%.
- Volume 71-80% Utilized - The number of volumes being utilized 71-80%.
- Volume 81-90% Utilized - The number of volumes being utilized 81-90%.
- Volume 91-100% Utilized - The number of volumes being utilized 91-100%.

CA 1 Configuration table view

This table contains a single row that provides configuration and control information for the CA1 tape management system. Notice that many of these same attributes are shown on the linked-to Summary workspaces described below. In those cases the attribute values are relative to the specific object that you are summarizing (for example, a location, an owner, a creating job, and so on). The attribute values are not relative to the entire CA 1 environment:

- First Audit Record Since Last Backup
- First DSNB Offset - The relative address of the first DSNB record.
- First Volume Definition Offset - The relative address of the first volume definition record.
- Last TMC Backup - The time of the last TMC backup.
- Managed System - The SMS name of this system.
- New TMC VOLSER - The VOLSER of the new backup tape.
- Next Audit Offset - The relative address of the next available audit record.
- Next Free DSNB Offset - The relative address of the next free DSNB record.
- Number of Active DSNB Records - The total number of DSNB records currently in use.
- Number of Audit Records Generated
- Number of Formatted DSNB Records - The total number of formatted (ready to use) DSNB records.
- Number of Volume Definition Records
- TMC Backup DSN - The DSN of the last backup tape.
- TMC Backup Indicator
- TMSEXTEND Indicator - TMS EXTEND audit flag.
- VOLSER - The VOLSER of the last backup tape.

Workspaces that are linked to this workspace

See “Navigating the workspaces of the CA 1 Status node” on page 97.
See “Workspaces” on page 14 to check a list of all nodes provided in IBM® Z OMEGAMON® AI for Storage® and learn how to navigate among workspaces.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the CA 1 Summary node.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

CA 1 Summary workspace		
View	**View type**	**Attribute group**

CA 1 Scratch Volumes	PIE	CA1 Summary
CA 1 DSNB Allocation	PIE	CA1 Configuration
CA 1 Volume Utilization Summary	BAR	CA1 Summary
CA 1 System Summary	TABLE	CA1 Summary
CA 1 Configuration	TABLE	CA1 Configuration

This workspace contains views displaying summary attributes for all the Jobs that were found during the last Summary Collection cycle:

CA 1 Job Summary workspace		
View	**View type**	**Attribute group**
CA 1 Top Jobs by Total Datasets	BAR	CA1 Job Summary
CA 1 Job Summary	TABLE	CA1 Job Summary

This workspace contains views displaying summary attributes for all the Programs that were found during the last Summary Collection cycle:

CA 1 Program Summary workspace		
View	**View type**	**Attribute group**
CA 1 Top Programs by Total Datasets	BAR	CA1 Program Summary
CA 1 Program Summary	TABLE	CA1 Program Summary

This workspace contains views displaying summary attributes for all the Errors that were found during the last Summary Collection cycle:

CA 1 Error Summary workspace		
View	**View type**	**Attribute group**
CA 1 Top 10 Permanent Error Volumes	BAR	CA1 Permanent Error Summary
CA 1 Top 10 Temporary Error Volumes	BAR	CA1 Temporary Error Summary
CA 1 Top 100 Permanent Error Volumes	TABLE	CA1 Permanent Error Summary
CA 1 Top 10 Temporary Error Volumes	TABLE	CA1 Temporary Error Summary

This workspace contains views displaying summary attributes for all the Datasets that were found during the last Summary Collection cycle:

CA 1 Summary Dataset workspace		
View	**View type**	**Attribute group**
CA 1 Summary Datasets	TABLE	CA1 Dataset

This workspace contains views displaying information about a dataset:

CA 1 Dataset Details workspace		
View	**View type**	**Attribute group**
CA 1 Dataset Information	TABLE	CA1 Dataset
CA 1 Dataset Creation Information	TABLE	CA1 Dataset

This workspace contains views displaying volume information:

CA 1 Volume Details workspace		
View	**View type**	**Attribute group**
CA 1 Volume Errors	BAR	CA1 Volume

CA 1 Volume Information	TABLE	CA1 Volume
CA 1 Volume Interaction Information	TABLE	CA1 Volume
CA 1 Volume Retention Information	TABLE	CA1 Volume

This workspace contains views displaying information about Datasets:

CA 1 Volume Datasets workspace		
View	**View type**	**Attribute group**
CA 1 Dataset Block Count	BAR	CA1 Dataset
CA 1 Datasets	TABLE	CA1 Dataset

This workspace contains views displaying information about VOLSERs:

CA 1 Volume Search workspace		
View	**View type**	**Attribute group**
CA 1 Permanent Volume Errors > 0	BAR	CA1 Volume
CA 1 Volumes	TABLE	CA1 Volume

This workspace contains views displaying information about Datasets:

CA 1 Dataset Search workspace		
View	**View type**	**Attribute group**
CA 1 Dataset Block Count	BAR	CA1 Dataset
CA 1 Datasets	TABLE	CA1 Dataset

Monitoring Record Level Sharing

This topic describes the workspaces that monitor the status of Record Level Sharing. You access these workspaces through the Record Level Sharing Summary node in the Tivoli® Enterprise Portal.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

In addition, with OMEGAMON® for Storage on z/OS® Version 4.2.0 Interim Feature 3 or later, you can issue dataset-level Storage Toolkit commands from the dataset group level. To access these commands, open the RLS Summary workspace from the Navigator, right-click a row on the table view and select the **Dataset Actions...** menu item, then select the action you want to perform.

Navigating the workspaces of the Record Level Sharing Summary node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the Record Level Sharing Summary node to navigate to the workspaces that you see in "Record Level Sharing Summary node: bulleted list showing the hierarchy of linked workspaces" on page 103. You can view details regarding the following types of dataset groups: Non-VSAM, VSAM, and GDG (generation data group). You can access the workspaces in this table using the standard path and method for navigating workspaces.

Record Level Sharing Summary node: bulleted list showing the hierarchy of linked workspaces

Record Level Sharing Summary
• RLS Buffer LRU Detail
  ◦ RLS Buffer LRU System
  ◦ RLS Buffer Pool Detail
• RLS Cache Structure Datasets
  ◦ RLS Storage Class Dataset Performance
  ◦ RLS Dataset Performance Detail
• RLS Storage Class Details
  ◦ RLS Storage Class Datasets
  ◦ RLS Storage Class Dataset Performance
  ◦ RLS Dataset Performance Detail
• RLS Lock Structure Detail

Record Level Sharing Summary workspace

This topic describes the main components of the principal workspace for this node.

The following sections describe various components of this workspace.

Views

This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

RLS Lock Structure CF Details table view

The tabular portion covering the entire width of the workspace on the screen displays one row of data for each lock structure. The rows contain:

- The duplex elapsed time in seconds.
- The percent of lock structure entries in use in tenths of a percent.
- The lock structure mode. Valid values are duplex, duplexing, and simplex.
- The name of the lock structure.
- The lock status.
- An internal node name used to identify the system.
- The primary lock structure version.
- The secondary lock structure version.
- The size of the lock structure in kb.
- The total number of entries in the lock structure.
- The number of entries in use in the lock structure.

RLS Overview table view

The tabular portion covering the entire width of the workspace on the screen displays one row of data for each lock structure. The rows contain:

- The average number of requests waiting for locks in thousandths.
- The percentage of lock requests that were globally managed in thousandths of a percent.
- The percentage of lock requests that were falsely globally managed in thousandths of a percent.
- The time interval for which data was captured.
- The name of the lock structure.

- The number of lock requests per second in thousandths.
- The internal node name used to identify the system.
- The SMF ID.
- The statistics level, indicating if the data is for the sysplex or an individual system.
- The sysplex or system name.
- The number of systems included in the sysplex statistics.

Lock Structure Summary table view

The tabular portion covering the entire width of the workspace on the screen displays one row of data for each lock structure. The rows contain:

- The total number of ESDS Add to End lock requests in false contention.
- The number of ESDS Add to End lock requests per minute in hundredths.
- The total number of ESDS Add to End lock requests in true contention.
- The percent of ESDS Add to End lock requests in false contention in tenths of a percent.
- The number of ESDS Add to End lock requests in false contention per minute in hundredths.
- The percent of ESDS Add to End lock requests in true contention in tenths of a percent.
- The number of ESDS Add to End lock requests in true contention per minute in hundredths.
- The average response time for the dataset in .01 milliseconds.
- The time span over which data was collected in seconds.
- The time when data was collected.
- The number of Data Insert Work Area (DIWA) lock requests that had false contention per minute in hundredths.
- The number of Data Insert Work Area (DIWA) lock requests that had false contention.
- The percent of Data Insert Work Area (DIWA) lock requests that had false contention in tenths of a percent.
- The number of Data Insert Work Area (DIWA) lock requests that had true contention.
- The number of Data Insert Work Area (DIWA) lock requests.
- The number of Data Insert Work Area (DIWA) lock requests per minute in hundredths.
- The number of Data Insert Work Area (DIWA) lock requests that had true contention per minute in hundredths.
- The percent of Data Insert Work Area (DIWA) lock requests that had true contention in tenths of a percent.
- The total number of component 2 (ESDS Add to End) lock requests.
- The VSAM RLS exception path lock requests per minute in hundredths.
- The total number of lock requests. VSAM RLS exception path.
- The total number of false contention VSAM RLS exception path lock requests.
- The number of false contention VSAM RLS exception path lock requests per minute in hundredths.
- The percent of exception path lock requests that were in false contention in tenths of a percent.
- The total number of true contention VSAM RLS exception path lock requests.
- The number of true contention VSAM RLS exception path lock requests per minute in hundredths.
- The percent of exception path lock requests that were in true contention in tenths of a percent.
- The DFSMS lock table name.
- The total number of false contention VSAM RLS mainline lock requests.

- The number of false contention VSAM RLS mainline lock requests per minute in hundredths.
- The percent of mainline lock requests that were in false contention in tenths of a percent.
- The total number of true contention VSAM RLS mainline lock requests.
- The number of true contention VSAM RLS mainline lock requests per minute in hundredths.
- The percent of mainline lock requests that were in true contention in tenths of a percent.
- The VSAM RLS mainline path lock requests per minute in hundredths.
- The total number of lock requests. VSAM RLS mainline path.
- The internal node name used to identify the system.
- The percent of record lock requests that were in false contention in tenths of a percent.
- The number of record lock requests in false contention per minute in hundredths.
- The number of record lock requests (obtain, alter, promote) in false contention.
- The number of record lock requests in true contention per minute in hundredths.
- The number of record lock requests (obtain, alter, promote) in true contention.
- The percent of record lock requests that were in true contention in tenths of a percent.
- The number of record lock requests per minute in hundredths.
- The number of record lock requests (obtain, alter, promote).
- The SMF ID.
- The statistics level, indicating if the data is for the sysplex or an individual system
- The sysplex or system name.
- The number of upgrade locks in false contention per minute in hundredths.
- The number of upgrade locks in true contention per minute in hundredths.
- The number of upgrade locks (component 1 class 2).
- The percentage of upgrade lock requests that had false contention in tenths of a percent.
- The number of upgrade locks in false contention.
- The number of upgrade locks per minute in hundredths.
- The percentage of upgrade lock requests that had true contention in tenths of a percent.
- The number of upgrade locks in true contention.

Buffer LRU Summary table view

The tabular portion covering the entire width of the workspace on the screen provides the status of the activity for all buffers. The rows contain:

- The average CF Cache read hit percent over the collection interval.
- The average DASD hit percent over the collection interval.
- The average percentage of RE-DO operations over the collection interval.
- The average percentage of recursive RE-DO operations across the sysplex.
- The total number of buffer manager LRU intervals processed, where BMF was over the goal, accelerated the aging, but did not go into panic mode.
- The percentage of average intervals where BMF was over the goal, accelerated the aging but did not go into panic mode in tenths of a percent.
- The total number of Buffer Manager LRU intervals where BMF was over the goal, and normal algorithms were bypassed to reclaim buffers.
- The percent of LRU intervals during which BMF was in panic mode in tenths of a percent.

- The total write requests through BMF.
- The time span over which data was collected in seconds.
- The time when data was collected.
- The percentage of recursive RE-DO operations over the current LRU interval.
- The BMF read hit percent for the current LRU interval.
- The percentage of CF cache read hits over the current LRU interval.
- The percentage of DASD hits over the current LRU interval.
- The percentage of RE-DO operations over the current LRU interval.
- Indicates if the statistics are for buffers above or below the bar.
- The maximum BMF read hit percent for any LRU interval over the collection interval.
- The maximum percentage of recursive RE-DO operations for any LRU interval over the collection interval.
- The maximum CF Cache read hit percent for any LRU interval over the collection interval.
- The maximum DASD hit percent for any LRU interval over the collection interval.
- The maximum percentage of RE-DO operations for any LRU interval over the collection interval.
- The minimum BMF read hit percent for any LRU interval over the collection interval.
- The minimum percentage of recursive RE-DO operations for any LRU interval over the collection interval.
- The minimum CF Cache read hit percent for any LRU interval over the collection interval.
- The minimum DASD hit percent for any LRU interval over the collection interval.
- The minimum percentage of RE-DO operations for any LRU interval over the collection interval.
- The internal node name used to identify the system.
- The total number of recursive RE-DO operations.
- The percentage of RE-DO operations that went recursive in tenths of a percent.
- The SMF ID.
- The statistics level, indicating if the data is for the sysplex or an individual system.
- The sysplex or system name.
- The total number of Buffer Manager LRU intervals processed.
- The number of BMF read hits.
- The total number of times that BMF was called.
- The total number of CF cache read hits.
- The number of times requests were satisfied from DASD.
- The total number of redos.
- The percentage of castout contentions retried in tenths of a percent.
- The total number of CF read requests that encountered XCF cast-out lock contention.
- The total number of retries for cast-out lock.
- The total number of CF read requests that had to be retried more than once because they encountered contention on the cast-out lock.

Storage Class Summary table view

The tabular portion covering the entire width of the workspace on the screen provides performance statistics for each Storage Class with RLS activity in the sysplex. The rows contain:

- The average response time for the Storage class in .01 milliseconds.
- The average response time for all direct access requests in .01 milliseconds
- The average response time for all sequential access requests in .01 milliseconds.
- The number of BMF direct access false invalids.
- The number of BMF direct access false invalids per minute in hundredths.
- The number of BMF direct access requests per minute in hundredths.
- The number of BMF direct access requests.
- The percent of BMF requests that were false invalids in tenths of a percent.
- The number of BMF false invalids.
- The number of BMF false invalids per minute in hundredths.
- The number of BMF requests.
- The number of BMF requests per minute in hundredths.
- The number of BMF sequential access false invalids.
- The number of BMF sequential access false invalids per minute in hundredths.
- The number of BMF sequential access requests per minute in hundredths.
- The number of BMF sequential access requests.
- The number of control area splits per minute in hundredths for all datasets in the storage class.
- The percent of writes that resulted in control area splits for all datasets in the storage class.
- The cache set used by this dataset.
- The number of control interval splits per minute in hundredths for all datasets in the storage class.
- The percent of writes that resulted in a control interval split for all datasets in the storage class.
- The time span over which data was collected in seconds.
- The time when data was collected.
- The number of control area splits for all datasets in the storage class.
- The number of control area splits for all datasets in the storage class.
- The number of direct access reads with CR per minute in hundredths.
- The number of direct access reads with NRI per minute in hundredths.
- The number of direct access reads with Consistent Read.
- The number of direct access reads with No Read Integrity.
- The direct access requests per minute in hundredths.
- The total number of direct access requests.
- The number of direct access writes.
- The number of direct access writes per minute in hundredths.
- The percent of write requests that had Data Insert Work Area (DIWA) lock requests in tenths of a percent.
- The number of Data Insert Work Area (DIWA) lock requests that had true contention.
- The number of Data Insert Work Area (DIWA) lock requests.
- The number of Data Insert Work Area (DIWA) lock requests per minute in hundredths.
- The number of Data Insert Work Area (DIWA) lock requests that had true contention per minute in hundredths.

- The percent of Data Insert Work Area (DIWA) lock requests that had true contention in tenths of a percent.
- The number of ESDS Add to End lock requests.
- The percent of write requests that had ESDS Add to End lock requests in tenths of a percent.
- The number of ESDS Add to End lock requests per minute in hundredths.
- The number of ESDS Add to End lock requests in true contention per minute in hundredths.
- The number of ESDS Add to End lock requests in true contention.
- The number of lock requests in contention per minute in hundredths.
- The percent of lock requests in contention in .1 percent.
- The percent of lock requests in false contention in .1 percent.
- The number of lock requests in false contention per minute in hundredths.
- The number of lock requests in true contention per minute in hundredths.
- The number of lock requests.
- The number of lock requests in false contention.
- The number of lock requests per minute in hundredths.
- The number of lock requests in true contention.
- The lock structure name.
- The percent of lock requests in true contention in .1 percent.
- The internal node name used to identify the system.
- The number of sequential access reads with CR per minute in hundredths.
- The number of sequential access reads with NRI per minute in hundredths.
- The sequential access requests per minute in hundredths.
- The total number of sequential access requests.
- The number of sequential access reads with CR Integrity.
- The number of sequential access reads with No Read Integrity.
- The number of sequential access writes.
- The number of sequential access writes per minute in hundredths.
- The SMF ID if this is a system level record.
- The indicator if this record is for the entire sysplex or for an individual system.
- The Storage Class Name.
- The system or sysplex name from which the data is derived.
- The number of reads.
- The number of reads per minute in hundredths.
- The total number of requests.
- The number of requests per minute in hundredths.
- The number of write requests.
- The number of write requests per minute in hundredths.
- The number of upgrade lock requests per minute in hundredths.
- The number of upgrade lock requests in true contention per minute in hundredths.
- The number of upgrade lock requests in true contention.

- The number of upgrade lock requests.

Cache Structure Summary table view

The tabular portion covering the entire width of the workspace on the screen provides performance statistics for each Cache Structure with RLS activity in the sysplex. The rows contain:

- The DFSMS Cache Structure name.
- The Read-hit counter.
- The Read-miss directory-hit counter.
- The Read-miss assignment suppressed counter.
- The Read-miss target storage class full counter.
- The Directory entry reclaim counter.
- The XI for directory reclaim counter.
- The XI for write counter.
- The XI for name invalidation counter.
- The XI for complement invalidation counter.
- The Cast-out counter.
- The Target storage class full counter.
- The Directory entry counter.
- The Data table entry reclaim counter.

Workspaces that are linked to this workspace

See “Navigating the workspaces of the Record Level Sharing Summary node” on page 102.
Learn how to navigate to workspaces here: “Workspaces” on page 14.

See “Workspaces” on page 14 for a list of all nodes provided in IBM® Z OMEGAMON® AI for Storage®.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the Record Level Sharing Summary node.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

Record Level Sharing Summary workspace		
View	**View type**	**Attribute group**
RLS Lock Structure CF Details	TABLE	RLS Lock Structure CF Detail
RLS Overview	TABLE	RLS Performance Overview
Lock Structure Summary	TABLE	RLS Lock Structure
Buffer LRU Summary	TABLE	RLS Buffer LRU Summary
Storage Class Summary	TABLE	RLS Storage Class
Cache Structure Summary	TABLE	RLS Cache Structure Summary

RLS Storage Class Details workspace		
View	**View type**	**Attribute group**

Average Storage Class Response Time	BAR	RLS Storage Class
Storage Class DIWA Lock Requests Per Minute	BAR	RLS Storage Class
Storage Class Sysplex Summary	TABLE	RLS Storage Class
Storage Class System Summary	TABLE	RLS Storage Class

RLS Lock Structure Storage Class workspace		
View	**View type**	**Attribute group**
Average Storage Class Response Time	BAR	RLS Storage Class
Storage Class DIWA Lock Requests Per Minute	BAR	RLS Storage Class
Storage Class Sysplex Summary	TABLE	RLS Storage Class

RLS Storage Class Datasets workspace		
View	**View type**	**Attribute group**
Average Response Time	BAR	RLS Storage Class
DIWA Lock Requests Per Minute	BAR	RLS Storage Class
Storage Class Summary	TABLE	RLS Storage Class
Dataset Performance Summary	TABLE	RLS Storage Class

RLS Storage Class Dataset Performance workspace		
View	**View type**	**Attribute group**
System Level Data Component Average Response Time	BAR	RLS Dataset Details
System Level Index Component Average Response Time	BAR	RLS Dataset Details
Sysplex Level RLS Dataset Performance	TABLE	RLS Dataset Details
System Level RLS Dataset Performance	TABLE	RLS Dataset Details

RLS Lock Structure Detail workspace		
View	**View type**	**Attribute group**
DIWA Lock Requests per Minute	BAR	RLS Lock Structure
DIWA Lock True Contention Percent	BAR	RLS Lock Structure
Sysplex Level Lock Structure Summary	TABLE	RLS Lock Structure
System Level Lock Structure Summary	TABLE	RLS Lock Structure

RLS Buffer LRU Detail workspace		
View	**View type**	**Attribute group**
Sysplex Level Above the Bar Buffers	TABLE	RLS Buffer LRU Summary
Sysplex Level Below the Bar Buffers	TABLE	RLS Buffer LRU Summary

Sysplex Level Buffer Pool Summary	TABLE	RLS Buffer LRU Summary
Local Level Above the Bar Buffer Pool Summary	TABLE	RLS Buffer LRU Summary
Local Level Below the Bar Buffer Pool Summary	TABLE	RLS Buffer LRU Summary

RLS Buffer LRU System Detail workspace		
View	**View type**	**Attribute group**
Above the Bar Buffer Detail	BAR	RLS Buffer LRU Summary
Below the Bar Buffer Detail	BAR	RLS Buffer LRU Summary
Local Level Buffer Pool Summary	TABLE	RLS Buffer LRU Summary
Local Level Above the Bar Buffer Pool Detail	TABLE	RLS Buffer LRU Summary
Local Level Below the Bar Buffer Pool Detail	TABLE	RLS Buffer LRU Summary

RLS Buffer Pool Detail workspace		
View	**View type**	**Attribute group**
Above the Bar Buffer Detail	BAR	RLS Buffer Pool Detail
Below the Bar Buffer Detail	BAR	RLS Buffer Pool Detail
Sysplex Level Buffer Pool Summary	TABLE	RLS Buffer Pool Detail
Local Level Above the Bar Buffer Pool Detail	TABLE	RLS Buffer Pool Detail
Local Level Below the Bar Buffer Pool Detail	TABLE	RLS Buffer Pool Detail

RLS Lock Performance Detail workspace		
View	**View type**	**Attribute group**
RLS Lock Structure CF Details	TABLE	Lock Structure CF Detail
RLS Sysplex Level Lock Overview	TABLE	Lock Structure CF Detail
1 Min System Level Lock Performance	TABLE	RLS Performance Overview
1 Hour System Level Lock Performance	TABLE	RLS Performance Overview
8 Hour System Level Lock Performance	TABLE	RLS Performance Overview
1 Day System Level Lock Performance	TABLE	RLS Performance Overview

RLS Dataset Performance Detail workspace		
View	**View type**	**Attribute group**
Component Percent of Dataset MSR for All Jobs	PIE	Dataset Performance
Dataset Performance Detail Report	TABLE	Dataset Performance

RLS Cache Structure Datasets workspace		
View	**View type**	**Attribute group**
Dataset Performance Summary	TABLE	RLS Dataset Details

Monitoring Dataset Attributes System Summary data

This topic describes the workspaces that monitor the status of Dataset Attributes System Summary data. You access these workspaces through the Dataset Attributes System Summary node in the Tivoli® Enterprise Portal.

With IBM® Z OMEGAMON® AI for Storage®, you can also view multivolume datasets as a single entity, rather that as multiple entitites over multiple volumes. The Dataset Attributes System Summary workspace displays data for multivolume datasets, where applicable, so that you can monitor the largest datasets on your system, not just the largest on a single volume. In addition, the Multivolume Dataset Attribute Details workspace provides dataset attributes for a multivolume dataset, including dataset-level attributes and volume-specific attributes.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

You can also issue dataset-level Storage Toolkit commands from the dataset group level. To access these commands, open the Dataset Attributes Group Summary workspace from the Navigator, right-click a row on the table view and select the **Dataset Actions...** menu item, then select the action you want to perform.

Navigating the workspaces of the Dataset Attributes System Summary node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the Dataset Attributes System Summary node to navigate to the workspaces that you see in “Dataset Attributes System Summary node: bulleted list showing the hierarchy of workspace links” on page 113. You can access the workspaces in this table using the standard path and method for navigating workspaces.

Dataset Attributes System Summary node: bulleted list showing the hierarchy of workspace links

Dataset Attributes System Summary

- Dataset Space Summary
  - Dataset Attribute Details
    - Dataset Details
  - Multivolume Dataset Attribute Details
- DCB Summary
  - Largest Datasets in DSORG
    - Dataset Attribute Details
      - Dataset Details
    - Multivolume Dataset Attribute Details
  - Largest Inefficient Blocksize Datasets
    - Dataset Attribute Details
      - Dataset Details
    - Multivolume Dataset Attribute Details
  - Newborn Dataset Summary
    - Dataset Attribute Details
      - Dataset Details
    - Multivolume Dataset Attribute Details
  - Largest Never Referenced Datasets
    - Dataset Attribute Details
      - Dataset Details
    - Multivolume Dataset Attribute Details
  - Mature Dataset Summary – Largest Datasets Unreferenced > 366 Days
    - Dataset Attribute Details
      - Dataset Details
    - Multivolume Dataset Attribute Details
- Catalog Summary
  - Datasets in Catalog
    - Dataset Attribute Details
      - Dataset Details
    - Multivolume Dataset Attribute Details
- DFSMS Class Summary
  - SMS Data Class Details
  - SMS Management Class Details
  - SMS Storage Class Details
  - SMS Storage Group Details
  - Largest Datasets in SMS Class (appears for each class)
    - Dataset Attribute Details
      - Dataset Details
    - Multivolume Dataset Attribute Details
- Uncataloged Dataset Summary
  - Largest Uncataloged Datasets on Volume
    - Dataset Attribute Details
      - Dataset Details
- Dataset Attributes Group Summary

- Dataset Attributes Group Details
  - Dataset Attribute Details
    - Dataset Details
  - Multivolume Dataset Attribute Details

Dataset Attributes System Summary workspace

This topic describes the main components of the principal workspace for this node.

This workspace provides the principal view of the Dataset Attribute Database, which maintains data on dataset attributes related to space, DCB, and cataloging. For more information about the Dataset Attribute Database feature, including how to configure it and how to work with groups to monitor a collection of datasets based on specific attributes, see "Dataset Attributes Database" on page 204.

> **Note:** Data collection is turned off by default, so this workspace is initially empty.

Views

This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

Table view

The table view displays one row of data for each resource. The columns show the following attributes:

- Volumes online
- Volumes processed
- Datasets processed
- Available tracks
- Allocated tracks
- Used tracks
- Unused tracks
- Percent of available allocated
- Percent of allocated used
- Percent of allocated free
- Name of largest dataset
- Size of largest dataset
- Name of dataset with largest unused space
- Amount of space unused by largest dataset
- Name of dataset with oldest reference date
- Number of unreferenced days for oldest dataset
- Number of uncataloged datasets
- Collection start & end date/time

Workspaces that are linked to this workspace

See "Navigating the workspaces of the Dataset Attributes System Summary node" on page 112.
Learn how to navigate to workspaces here: "Workspaces" on page 14.

See "Workspaces" on page 14 for a list of all nodes provided in IBM® Z OMEGAMON® AI for Storage®.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the Dataset Attributes System Summary node.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

See “Dataset Attributes System Summary workspace” on page 114 for more information on the primary workspace for this node.

Dataset Attributes System Summary workspace		
View	**View type**	**Attribute group**
Dataset Attribute Summary	TABLE	Dataset Attributes System Summary
Overall Space Utilization	PIE	Dataset Attributes System Summary
Dataset Attribute Dataset Extremes	TABLE	Dataset Attributes System Summary
Overall Space Allocation	BAR	Dataset Attributes System Summary

Dataset Space Summary workspace		
View	**View type**	**Attribute group**
Top Datasets by Allocated Space	TABLE	Dataset Attributes Space Allocated Summary
Top Datasets by Extents	TABLE	Dataset Attributes Extent Summary
Top Datasets by Unused Space	TABLE	Dataset Attributes Space Unused Summary
Top Datasets by CA Splits	TABLE	Dataset Attributes CA Split Summary
Top Datasets by CI Splits	TABLE	Dataset Attributes CI Split Summary

DCB Summary workspace		
View	**View type**	**Attribute group**
DSORG Summary Report	TABLE	Dataset Attributes DSORG Summary
Blocksize Distribution Report	TABLE	Dataset Attributes Blocksize Summary
Blocksize Track Allocation Distribution	BAR	Dataset Attributes Blocksize Summary
Creation Date Distribution Report	TABLE	Dataset Attributes Creation Date Summary
Creation Date Track Allocation Distribution	BAR	Dataset Attributes Creation Date Summary
Reference Date Distribution Report	TABLE	Dataset Attributes Reference Date Summary

Reference Date Track Allocation Distribution	BAR	Dataset Attributes Reference Date Summary

Catalog Summary workspace		
View	**View type**	**Attribute group**
Catalog Summary Report	TABLE	Dataset Attributes Catalog Summary
Datasets per Catalog	BAR	Dataset Attributes Catalog Summary

DFSMS Class Summary workspace		
View	**View type**	**Attribute group**
Storage Group Allocation	BAR	Dataset Attributes SMS Class Summary
Storage Groups	TABLE	Dataset Attributes SMS Class Summary
Data Classes	TABLE	Dataset Attributes SMS Class Summary
Data Class Allocation	BAR	Dataset Attributes SMS Class Summary
Management Classes	TABLE	Dataset Attributes SMS Class Summary
Management Class Allocation	BAR	Dataset Attributes SMS Class Summary
Storage Classes	TABLE	Dataset Attributes SMS Class Summary
Storage Class Allocation	BAR	Dataset Attributes SMS Class Summary

Dataset Attribute Details workspace		
View	**View type**	**Attribute group**
Space Attributes	TABLE	Dataset Attributes Dataset Detail
DCB and Extended Attributes	TABLE	Dataset Attributes Dataset Detail
GDG Information	TABLE	Dataset Attributes Dataset Detail
Dataset SMS Constructs	TABLE	Dataset Attributes Dataset Detail
Dataset Space	BAR	Dataset Attributes Dataset Detail
VSAM Attributes	TABLE	Dataset Attributes Dataset Detail
Encryption Information	TABLE	Dataset Attributes Dataset Detail

Largest Datasets in SMS Class workspace		
View	**View type**	**Attribute group**
Largest Datasets in SMS Class by Allocated Space	TABLE	Dataset Attributes SMS Class Detail
Space Allocation for Largest Datasets in SMS Class	BAR	Dataset Attributes SMS Class Detail

Largest Datasets in DSORG workspace		
View	**View type**	**Attribute group**

Space Allocated for Largest Datasets in DSORG	BAR	Dataset Attributes DSORG Detail
Largest Datasets in DSORG by Allocated Space	TABLE	Dataset Attributes DSORG Detail

Largest Inefficient Blocksize Datasets workspace		
View	**View type**	**Attribute group**
Largest Datasets where Blocksize < 128 by Allocated Space	TABLE	Dataset Attributes Inefficient Blocksize Summary
Space Allocation for Largest Datasets where Blocksize < 128	BAR	Dataset Attributes Inefficient Blocksize Summary

Newborn Dataset Summary - Datasets 0 or 1 Day Old workspace		
View	**View type**	**Attribute group**
Largest Datasets 0 or 1 Day Old by Allocated Space	TABLE	Dataset Attributes Newborn Dataset Summary
Space Allocation for Largest Datasets 0 or 1 Day Old	BAR	Dataset Attributes Newborn Dataset Summary

Mature Dataset Summary - Largest Datasets Unreferenced > 366 Days workspace		
View	**View type**	**Attribute group**
Largest Datasets Unreferenced for > 366 Days by Allocated Space	TABLE	Dataset Attributes Mature Dataset Summary
Space Allocation for Largest Datasets Unreferenced for > 366 Days	BAR	Dataset Attributes Mature Dataset Summary

Largest Never Referenced Datasets workspace		
View	**View type**	**Attribute group**
Largest Never Referenced Datasets by Allocated Space	TABLE	Dataset Attributes Never Referenced Dataset Summary
Space Allocation for Largest Never Referenced Datasets	BAR	Dataset Attributes Never Referenced Dataset Summary

Datasets in Catalog workspace		
View	**View type**	**Attribute group**
All Datasets in Catalog	TABLE	Dataset Attributes Dataset Detail
Space Utilization for All Datasets in Catalog	PIE	Dataset Attributes Catalog Summary

Uncataloged Dataset Summary workspace		
View	**View type**	**Attribute group**
Uncataloged Dataset Summary Report	TABLE	Dataset Attributes Uncataloged Summary
Uncataloged Dataset Allocation by Volume	BAR	Dataset Attributes Uncataloged Summary
Uncataloged Dataset Total	BAR	Dataset Attributes Catalog Summary
Uncataloged Dataset Allocation Total	BAR	Dataset Attributes Catalog Summary

Largest Uncataloged Datasets on Volume workspace		
View	**View type**	**Attribute group**
Largest Uncataloged Datasets on Volume by Allocated Space	TABLE	Dataset Attributes Uncataloged Detail
Space Utilization for Largest Uncataloged Datasets on Volume	BAR	Dataset Attributes Uncataloged Detail
Uncataloged Dataset Space Allocation	PIE	Dataset Attributes Uncataloged Detail

Dataset Attributes Group Summary workspace		
View	**View type**	**Attribute group**
Dataset Attributes Group Summary	TABLE	DSN Attr Group Summary
Used and Unused Tracks by Group	BAR	DSN Attr Group Summary

Dataset Attributes Group Details workspace		
View	**View type**	**Attribute group**
Dataset Attributes Group Details	TABLE	DSN Attr Group Detail
Used and Unused Tracks by Dataset	BAR	DSN Attr Group Detail

Multivolume Dataset Attribute Details workspace		
View	**View type**	**Attribute group**
Total Space Utilization	BAR	DSN Attr Group Detail
Space Utilization By Volume	BAR	DSN Attr Group Detail
Dataset SMS Constructs	TABLE	DSN Attr Group Detail
Dataset Basic Attributes	TABLE	DSN Attr Group Detail
VSAM Cluster Attributes	TABLE	DSN Attr Group Detail
Dataset Attributes By Volume	TABLE	DSN Attr Group Detail
Encryption Information	TABLE	DSN Attr Group Detail

Dataset HLQ Summary		
View	**View type**	**Attribute group**
Dataset HLQ Summary	TABLE	Dataset Attributes HLQ Summary Report

Monitoring dataset groups summary data

This topic describes the workspaces that monitor the status of Dataset Group Summary data. You access these workspaces through the Dataset Group Summary node in the Tivoli® Enterprise Portal.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

You can also issue dataset-level Storage Toolkit commands from the dataset group level. To access these commands, open the Dataset Group Summary workspace from the Navigator, right-click a row on the table view and select the **Dataset Actions...** menu item, then select the action you want to perform.

Navigating the workspaces of the Dataset Group Summary node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the Dataset Group Summary node to navigate to the workspaces that you see in "Dataset Group Summary node: bulleted list showing the hierarchy of linked workspaces" on page 119. You can view details regarding the following types of dataset groups: Non-VSAM, VSAM, and GDG (generation data group). You can access the workspaces in this table using the standard path and method for navigating workspaces.

Dataset Group Summary node: bulleted list showing the hierarchy of linked workspaces

- Dataset Group Summary
  - Dataset Group Details Primary
    - Dataset Group Details **Non-VSAM**
      - SMS Storage Class Details
      - Dataset Performance Detail Report [1]
    - Dataset Group Details **VSAM**
      - Dataset Group VSAM Attributes
        - Catalog Details
        - Catalog Dataset Attributes
        - Last Catalog Backup
        - SMS Storage Class Details
        - Dataset Performance Detail Report [1]
    - Dataset Group Details **GDG**
      - Dataset Group Details Non-VSAM
        - SMS Storage Class Details
        - Dataset Performance Detail Report [1]
  - Dataset Group RLS Datasets
    - Dataset Group RLS Dataset Performance

Footnotes for "Dataset Group Summary node: bulleted list showing the hierarchy of linked workspaces" on page 119:

1. This group of workspaces is organized in the following nested pattern:
   - Dataset Performance Detail Report
     - Dataset Details
     - Dataset Performance Detail Trend

Dataset Group Summary workspace

This topic describes the main components of the principal workspace for this node.

The following sections describe various components of this workspace.

Views

This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

Dataset group space utilization graph

This graph displays total free space and group total used space information for each dataset group.

Dataset group summary table view

The tabular portion covering the entire width of the workspace at the bottom of the screen displays one row of data for each defined dataset group. The rows contain:

- The dataset group name.
- The total number of datasets.
- The dataset with the largest allocation.
- The largest space allocation.
- The dataset with the largest free space.
- The largest space free.
- The dataset with the oldest referenced date within the group.
- The oldest referenced date for the identified dataset within the group.
- The dataset with the worst MSR within the group, and the highest MSR of the dataset.
- The total used space for all of the datasets within the group.
- The total number of I/Os for all the active datasets within the group.
- The total average MSR for all the active datasets within the group.
- The dataset group used space percentage.
- The total allocated space for all the datasets within the group.
- The dataset group percentage of free space.
- The total free space for all the datasets within the group.

Dataset group RLS summary table view

The tabular portion covering the entire width of the workspace at the bottom of the screen displays one row of data for each defined dataset group. The rows contain:

- The name of the dataset group.
- The maximum response time for any dataset in the group in .01 milliseconds.
- The maximum average response time for all direct access requests in .01 milliseconds for any dataset in the dataset group across the sysplex.
- The maximum average response time for all sequential access requests in .01 milliseconds for any dataset in the dataset group across the sysplex.
- The maximum percent of BMF false invalid requests for any dataset in the dataset group across the sysplex in tenths of a percent.
- The maximum number of BMF false invalid requests for any dataset in the dataset group across the sysplex.
- The maximum number of BMF false invalid requests per minute in hundredths for any dataset in the dataset group across the sysplex.
- The maximum number of BMF requests for any dataset in the dataset group across the sysplex.
- The maximum number of BMF requests per minute in hundredths for any dataset in the dataset group across the sysplex.
- The maximum number of control area splits for any dataset in the group across the sysplex.
- The maximum percent of write requests that caused a control area split for any dataset in the dataset group across the sysplex in tenths of a percent.
- The maximum number of control area splits per minute in hundredths for any dataset in the dataset group across the sysplex.
- The maximum number of control interval splits for any dataset in the group across the sysplex.

- The maximum percent of write requests that caused a control interval split for any dataset in the dataset group across the sysplex in tenths of a percent.
- The maximum number of control interval splits per minute in hundredths for any dataset in the dataset group across the sysplex.
- The maximum number of Data Insert Work Area (DIWA) lock requests per minute in hundredths for any dataset in the dataset group across the sysplex.
- The maximum percent of write requests that had Data Insert Work Area (DIWA) lock requests for any dataset in the dataset group across the sysplex in tenths of a percent.
- The maximum number of Data Insert Work Area (DIWA) lock requests for any dataset in the dataset group across the sysplex.
- The maximum percent of write requests that had ESDS Add to End lock requests for any dataset in the dataset group across the sysplex in tenths of a percent.
- The maximum number of lock requests in contention per minute in hundredths for any dataset in the dataset group across the sysplex.
- The maximum percent of lock requests in contention for any dataset in the dataset group across the sysplex in tenths of a percent.
- The maximum number of read requests for any dataset in the dataset group across the sysplex.
- The maximum number of read requests per minute in hundredths for any dataset in the dataset group across the sysplex.
- The maximum number of requests per minute in hundredths for any dataset in the dataset group across the sysplex.
- The maximum total requests for any dataset in the dataset group across the sysplex.
- The maximum number of write requests for any dataset in the dataset group across the sysplex.
- The maximum number of write requests per minute in hundredths for any dataset in the dataset group across the sysplex.

Workspaces that are linked to this workspace

See "Navigating the workspaces of the Dataset Group Summary node" on page 119.
Learn how to navigate to workspaces here: "Workspaces" on page 14.

See "Workspaces" on page 14 for a list of all nodes provided in IBM® Z OMEGAMON® AI for Storage®.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the Dataset Group Summary node.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

See "Dataset Group Summary workspace" on page 119 for more information on the primary workspace for this node.

Dataset Group Summary workspace		
View	**View type**	**Attribute group**
Dataset Group Summary	TABLE	Dataset Group Summary
Dataset Group RLS Summary	TABLE	RLS Dataset Group Summary
Dataset Group Space Utilization	BAR	Dataset Group Summary

Dataset Group Details Primary workspace		
View	**View type**	**Attribute group**
Total Allocations	BAR	Dataset Group Details
Response Time	BAR	Dataset Group Details
Dataset Group Primary Entries	TABLE	Dataset Group Details
I/O Activity	BAR	Dataset Group Details

Dataset Group Details VSAM workspace		
View	**View type**	**Attribute group**
VSAM Cluster Report	TABLE	Dataset Group Details
Response Time	BAR	Dataset Group Details
CI/CA Splits	BAR	Dataset Group Details
Record Access	BAR	Dataset Group Details

Dataset Group Details Non-VSAM workspace		
View	**View type**	**Attribute group**
Response Time	BAR	Dataset Group Details Volume
Dataset Information by Volume	TABLE	Dataset Group Details Volume
Space Utilization	BAR	Dataset Group Details Volume
Aggregate Dataset Information	TABLE	Dataset Group Details

Dataset Group Details GDG workspace		
View	**View type**	**Attribute group**
Aggregate GDG Dataset Information	TABLE	Dataset Group Details
Response Time	BAR	Dataset Group Details
I/O Activity	BAR	Dataset Group Details
Total Allocations	BAR	Dataset Group Details

Dataset Group VSAM Attributes workspace		
View	**View type**	**Attribute group**
Aggregate Dataset Information	TABLE	Dataset Group Details
Dataset Information by Volume	TABLE	Dataset Group Details Volume
Dataset Attributes	TABLE	Dataset Group Detail Attrs

Dataset Group RLS Datasets workspace		
View	**View type**	**Attribute group**
Average Response Time	BAR	RLS Dataset Group Details
DIWA Lock Requests per Minute	BAR	RLS Dataset Group Details
Dataset Group Summary	TABLE	RLS Dataset Group Details
Dataset Performance Summary	TABLE	RLS Dataset Group Details

Dataset Group RLS Dataset Performance workspace		
View	**View type**	**Attribute group**

System Level Data Component Average Response Time	BAR	RLS Dataset Group Details
System Level Index Component Average Response Time	BAR	RLS Dataset Group Details
Sysplex Level RLS Dataset Performance	TABLE	RLS Dataset Group Details
System Level RLS Dataset Performance	TABLE	RLS Dataset Group Details

Monitoring SMS configuration

This topic describes the workspaces that monitor the status of SMS configuration. You access these workspaces through the SMS Configuration node in the Tivoli® Enterprise Portal.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

Navigating the workspaces of the SMS Configuration node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the SMS Configuration node to navigate to the workspaces that you see in "SMS Configuration node: bulleted list showing the hierarchy of linked workspaces" on page 123. You can access the workspaces in this table using the standard path and method for navigating workspaces.

SMS Configuration node: bulleted list showing the hierarchy of linked workspaces
SMS Configuration • SMS Base Configuration ◦ SMS Configuration Details ◦ SMS Tape Storage Group Details ◦ SMS Data Class Details ◦ SMS Management Class Details ◦ SMS Storage Class Details

SMS Configuration workspace

This topic describes the main components of the principal workspace for this node.

The following sections describe various components of this workspace.

Views

This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help, including in the hover help.

SMS base configuration table view

The tabular report at the top of this workspace contains a single row of data describing SMS status in the reporting LPAR. Listed are:

- The SMS status and level
- The MVS™ level

- The TSO ID of the last user to update the SMS configuration
- The date and time of the last update
- The default management class
- The default number of tracks per cylinder, track size and unit type

This tabular report links to SMS Configuration Details workspace.

Storage groups table view

This tabular report contains one row for each storage group defined in the reporting LPAR. Each row lists:

- The storage group name and type
- The TSO ID of the last user to update the group
- The date and time of that update
- A description of the group

This tabular report links to SMS Storage Group Details and SMS Tape Storage Group Details workspaces.

Data classes table view

This tabular report contains one row for each defined SMS data class. Each row lists:

- The name of the data class
- The ID of the user who last updated the data class definition
- The date and time when the data class definition was updated
- A description of the data class

This tabular report links to SMS Data Class Details workspace.

Management classes table view

This tabular report contains one row for each defined SMS management class. Each row lists:

- The management class name
- The user ID that last updated this management class
- The date and time when this management class was last updated
- The storage administrator supplied description of the management class

This tabular report links to SMS Management Class Details workspace.

Storage classes table view

This tabular report contains one row for each defined SMS storage class. Each row lists:

- The name of the storage class
- The last user ID to update this storage class definition
- The date and time when the storage class definition was last updated
- The storage administrator supplied description of the storage class

This tabular report links to Storage Class Details workspace.

Workspaces that are linked to this workspace

See "Navigating the workspaces of the SMS Configuration node" on page 123.
Learn how to navigate to workspaces here: "Workspaces" on page 14.

See "Workspaces" on page 14 for a list of all nodes provided in IBM® Z OMEGAMON® AI for Storage®.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the SMS Configuration node.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

See “SMS Configuration workspace” on page 123 for more information on the primary workspace for this node.

SMS Configuration workspace		
View	**View type**	**Attribute group**
SMS Base Configuration	TABLE	SMS Configuration
Data Classes	TABLE	SMS Data Class
Management Classes	TABLE	SMS Management Class
Storage Classes	TABLE	SMS Storage Class
Storage Groups	TABLE	SMS Storage Group

SMS Storage Group Details workspace		
View	**View type**	**Attribute group**
Storage Group Details Report	TABLE	SMS Storage Group
Storage Group Details Report	TABLE	SMS Storage Group
Storage Group Status	TABLE	SMS Storage Group Status

SMS Tape Storage Group Details workspace		
View	**View type**	**Attribute group**
Storage Group Status	TABLE	SMS Storage Group Status
Tape Storage Group Details Report	TABLE	SMS Storage Group
Tape Storage Group Libraries	TABLE	SMS Storage Group

SMS Configuration Details workspace		
View	**View type**	**Attribute group**
SMS Configuration Members	TABLE	SMS Systems
Cache Sets	TABLE	SMS Cache Sets
Base Configuration	TABLE	SMS Configuration

SMS Data Class Details workspace		
View	**View type**	**Attribute group**
Data Class Allocation	TABLE	SMS Data Class
Data Class VSAM Attributes	TABLE	SMS Data Class
Data Class Details	TABLE	SMS Data Class

SMS Management Class Details workspace		
View	**View type**	**Attribute group**
Management Class Details	TABLE	SMS Management Class
Management Class Details	TABLE	SMS Management Class
Management Class Details	TABLE	SMS Management Class

SMS Storage Class Details workspace		
View	**View type**	**Attribute group**
Storage Class Details	TABLE	SMS Storage Class
Millisecond Response Time	BAR	SMS Storage Class

Monitoring Storage Toolkit action requests

This topic describes the workspaces that monitor the status of Storage Toolkit action requests. You access these workspaces through the Storage Toolkit Action Requests node in the Tivoli® Enterprise Portal.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

Navigating the workspaces of the Storage Toolkit node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the Storage Toolkit Action Requests node to navigate to the workspaces that you see in "Storage Toolkit Action Requests node: bulleted list showing the hierarchy of linked workspaces" on page 126. You can access the workspaces in this table using the standard path and method for navigating workspaces.

Storage Toolkit Action Requests node: bulleted list showing the hierarchy of linked workspaces
Storage Toolkit Action Requests • Storage Toolkit Action Requests ◦ Storage Toolkit Result Summary ◦ Storage Toolkit Result Detail

Storage Toolkit Action Requests workspace

This topic describes the main components of the principal workspace for this node.

The following sections describe various components of this workspace. For a description of the data used as column headings in the workspace table view, see the Storage Toolkit section.

Views

This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

Storage Toolkit Action Requests table view

The table view displays one row of data for each action created by the user. The view contains the following information:

- Name of user who created action request
- Name of request as given by user
- Description of request as given by user
- Type of action
- Number of times request has been run
- Number of results for the action

- Return code for most recently completed action
- Date/time when request was created
- Date/time of the most recently completed action
- Command issued for execution
- Scheduled dates/time at which action request is to be run

Workspaces that are linked to this workspace

See “Navigating the workspaces of the Storage Toolkit node” on page 126.
Learn how to navigate to workspaces here: “Workspaces” on page 14.

See “Workspaces” on page 14 for a list of all nodes provided in IBM® Z OMEGAMON® AI for Storage®.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the Storage Toolkit node.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

See “Storage Toolkit Action Requests workspace” on page 126 for more information on the primary workspace for this node.

Storage Toolkit Action Requests workspace		
View	**View type**	**Attribute group**
Action Requests	TABLE	Storage Toolkit Request
Requests with Last Return Code > 0	BAR	Storage Toolkit Request

Storage Toolkit Result Summary workspace		
View	**View type**	**Attribute group**
Action Results Summary	TABLE	Storage Toolkit Result Summary
Return Code	BAR	Storage Toolkit Result Summary
Output Line Count	BAR	Storage Toolkit Result Summary

Storage Toolkit Result Detail workspace		
View	**View type**	**Attribute group**
Resource List	TABLE	Storage Toolkit Result Resource
Action Results Detail	TABLE	Storage Toolkit Result Detail

Monitoring Copy Services

This topic describes the workspaces that monitor the status of Copy Services. You access these workspaces through Copy Services node in the Tivoli® Enterprise Portal.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.

- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

Navigating the workspaces of Copy Services node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources. See "Workspaces" on page 14 to learn more about navigating from workspace to workspace.

You access the primary workspaces of this node as follows:

1. Click the node in the Tivoli® Enterprise Portal. The node is activated and the first workspace is displayed in the portal.
2. Right-click the node to display the pop-up menu.
3. Select the **Workspace** submenu to see a list of the available linked workspaces. These linked workspaces, in turn, link to many other, related workspaces as shown in "Copy Services node: bulleted list showing the hierarchy of linked workspaces" on page 128

You start in the "z/GlobalMirror (XRC) Sessions workspace" on page 128 to access all the other workspaces that are listed in "Copy Services node: bulleted list showing the hierarchy of linked workspaces" on page 128. You can access the workspaces in this table using the standard path and method for navigating workspaces.

Copy Services node: bulleted list showing the hierarchy of linked workspaces

Copy Services Node
• z/GlobalMirror (XRC) Sessions
  ◦ XRC Volume Pairs
• Flash Copy Volumes [1]
  ◦ Flash Copy Relations [2]
• PPRC Volumes [1]
  ◦ PPRC Connections [2]
• Link Statistics [1]

Footnotes for "Copy Services node: bulleted list showing the hierarchy of linked workspaces" on page 128:

1. These workspaces are accessible only through the following workspaces:
   - Cache CU Performance
   - Cache CU Status
   - SMS Storage Groups Performance
   - SMS Storage Groups Space
2. These workspaces are accessible only through their parent workspaces in the hierarchy or the following workspaces:
   - SMS Storage Group Volume Performance
   - SMS Storage Group Volume Cache
   - SMS Storage Group Volume Space

z/GlobalMirror (XRC) Sessions workspace

This topic describes the main components of the principal workspace for this node.

When you click the Copy Services node in the Navigator of the Tivoli® Enterprise Portal, you see the z/GlobalMirror (XRC) Sessions workspace. The main components of the workspace are listed here.

Views

This workspace provides the following view. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

z/GlobalMirror (XRC) Sessions table view

This table contains a single row that summarizes information for the z/GlobalMirror (XRC) Sessions, and contains the following attributes. Notice that many of these same attributes are shown on the linked-to Summary workspaces described below. In those cases the attribute values are relative to the specific object that you are summarizing (for example, a location, an owner, a creating job, and so on).

- The session name.
- The HLQ.
- The storage subsystem ID associated with the specified storage control session.
- The serial Number.
- The license level of the storage array.
- The logical storage subsystem ID.
- The last record processed time.
- The state.
- The utility volume config.
- The configured timeout in milliseconds.
- The session number within the storage array.
- The total volume pairs.
- The duplex volume pairs.
- The copying volume pairs.
- The volume pair that XRC has not yet synchronized or resynchronized.
- The volume pair that contains non-timestamped writes.
- The established volume pair in an XRC session where XRC is no longer copying data from the primary volume to the secondary volume.
- The volume pair remains in the XRC session.
- The fix utility enabled.
- The device block enabled.
- The write pacing enabled.
- The suspend on long busy pending.
- The enhanced reader support enabled.
- The incremental resync enabled.

Workspaces that are linked to this workspace

See "Navigating the workspaces of Copy Services node" on page 128.
Learn how to navigate among workspaces here: "Workspaces" on page 14.

See "Workspaces" on page 14 for a list of all nodes provided in IBM® Z OMEGAMON® AI for Storage®.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the Copy Services node.

See "Workspaces" for the following information:

- Predefined workspaces and a list of all nodes that are included in the IBM® Z OMEGAMON® AI for Storage® product. Each node corresponds to a specific area of storage resource monitoring. And each node includes links to a unique set of storage resource workspaces.
- Navigating among workspaces.
- Linking to the workspaces of other IBM® Tivoli® Monitoring products.

z/GlobalMirror (XRC) Sessions workspace		
View	**View type**	**Attribute group**
z/GlobalMirror (XRC) Sessions	TABLE	S3 XRC Session

XRC Volume Pairs workspace		
View	**View type**	**Attribute group**
XRC Volume Pairs	TABLE	S3 XRC Volume Pair

Flash Copy Volumes workspace		
View	**View type**	**Attribute group**
Flash Copy Volumes	TABLE	S3 Flash Copy Volume

Flash Copy Relations workspace		
View	**View type**	**Attribute group**
Flash Copy Relations	TABLE	S3 Flash Copy Relation

PPRC Volumes workspace		
View	**View type**	**Attribute group**
PPRC Volumes	TABLE	S3 PPRC Volume

PPRC Connections workspace		
View	**View type**	**Attribute group**
PPRC Connections	TABLE	S3 PPRC Connection

Link Statistics workspace		
View	**View type**	**Attribute group**
Link Statistics	TABLE	S3 Link Statistics

Monitoring ADABAS Database Details

With these workspaces you can monitor the status of the ADABAS Database Details data. You access the workspaces through the ADABAS Summary node in the Tivoli® Enterprise Portal.

For general information on predefined workspaces and a list of all nodes that are provided by IBM® Z OMEGAMON® AI for Storage®, see the *IBM® Z OMEGAMON® AI for Storage® User's Guide.*. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. To learn how to navigate among workspaces and to link the workspaces of other IBM® Tivoli® Monitoring products, see "Workspaces" on page 14.

For more detail, see ADABAS DB files.

Navigating the workspaces of the ADABAS Summary node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the ADABAS Summary node to navigate to the workspaces that you see in the following hierarchy of linked workspaces. You can access the workspaces in the following using the standard path and method for navigating workspaces:

ADABAS Summary

- ADABAS Database Details

ADABAS Summary workspace

This topic describes the main components of the principal workspace for this node.

The following sections describe various components of this workspace.

Views

This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

ADABAS Databases table view

The table view displays one row of data for each ADABAS database configured in the reporting ADABAS subsystem. The view contains the following information:

- The SMS name of the system
- Collection time
- ADABAS version
- Database name and ID
- SVC number
- Max files
- The most recent data protextion log
- Database load time
- The number of files loaded to the database
- Length of the blocks in the database
- File extents allocated for the data storage space table in blocks
- The total number of bytes, blocks, and cylinders allocated, used, and unused, for the DATA and ASSO components
- Percent of the storage allocated for DATA and ASSO usage
- The checkpoint file, security file, and trigger file numbers
- The host name
- The SYSPLEX name

ADABAS Volume Report

The ADABAS Volumes Report provides a list of the volumes in the ADABAS subsystem, with information on each volume:

- Collection time
- Host
- Dataset name
- Volume
- Allocated bytes
- Unused bytes
- Reference date
- Database identification number

ADABAS DB Files Report

The ADABAS DB Files Report provides a list of the database files in the ADABAS subsystem, with information on each:

- Collection time
- Host

- SVC number
- Database identification number
- File number
- File name
- Data allocated bytes
- Data used bytes
- Address converter allocated bytes
- Upper index allocated bytes
- Upper index used bytes
- Normal index allocated bytes
- Normal index used bytes
- Total allocated bytes
- Total used bytes
- Database name
- File locked update
- File replication
- File refreshed
- Loaded with USERISN
- File updated
- Increasing MAXISN allowed
- Privileged usage
- ISN reuse allowed
- Highest ISN
- Highest index
- File loaded and updated date and time
- Minimum ISN
- ISN byte value
- Max compressed record length
- Number of file updates
- Number of loaded records
- Data storage secondary extent max blocks
- Upper index secondary extent max blocks
- Normal index secondary extent max blocks
- DATA padding factor
- ASSO padding factor
- Sysplex
- Unused DATA, address converter, upper index, and normal index storage bytes
- Allocated storage for DATA, address converter, upper index, and normal index in cylinders
- Used storage for DATA, address converter, upper index, and normal index in cylinders:

- Unused storage for DATA, address converter, upper index, and normal index in cylinders
- Data storage, address convert, upper index, and normal index logical extents
- File options (DE,UQ)

Workspaces that are linked to this workspace

See "Navigating the workspaces of the ADABAS Summary node" on page 130.
See the Workspace section for a list of all nodes provided in IBM® Z OMEGAMON® AI for Storage®.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the ADABAS Summary node.

For more information on the primary workspace for this node, see "ADABAS Summary workspace" on page 131.

ADABAS Summary workspace		
View	**View type**	**Attribute group**
ADABAS Databases	TABLE	ADABAS Databases

View	**View type**	**Attribute group**
ADABAS Volumes Report	TABLE	ADABAS Volumes
ADABAS DB Files Report	TABLE	ADABAS Files

Directly monitoring situations

With direct situation analysis, you navigate directly from a situation event workspace to information regarding the original situation that triggered a situation event.

Problem solved by direct situation analysis:

Before direct situation analysis existed, there were many cases in which only the default situation event workspaces were available when a situation is triggered. To investigate the metrics that triggered a situation, the user traversed the Navigator view to access information on the true situation. Direct situation analysis gives you quick access to this information.

When direct situation analysis is available:

Direct situation analysis is available after a situation becomes true, a situation event is opened, and a situation event workspace is created.
When a situation resolves to true and is associated with a Navigator item, a virtual light or situation event indicator overlays the item and all related items above it in the Navigator. As you hover with the mouse over the indicator, a list is displayed. When you click one of the situations in this list, the situation event workspace opens. This workspace shows values of attributes for the situation that was triggered and expert advice, if any.

Realtime Dataset Metrics

With Realtime Dataset Metrics (RDM), you can efficiently look up for detailed dataset attribute information.

> **Note:** The Realtime Dataset Metrics feature supports 5698-ZWG IBM Z Distribution for Zowe 2.16 with installed PTFs UO90057 and UO90058 or 5698-ZWH IBM Z Distribution for Zowe 3.0. The 5698-ZWH IBM Z Distribution for Zowe 3.0 is supported with installed IBM® Z OMEGAMON® AI for Storage® PTF UJ96869. Zowe version 1.*xx* is not supported anymore. Make sure that you use Zowe version 2.16 and later.

Realtime Dataset Metrics overview

Realtime Dataset Metrics is a browser user interface that you can launch from the Dataset Attributes System Summary in the Tivoli® Enterprise Portal (TEP) user interface to OMEGAMON® for Storage. The Realtime Dataset Metrics link is displayed in TEP when dataset collection is completed.

The key features of Realtime Dataset Metrics are the following:

- Summarization by key attributes
- Flexible filtering to quickly build ad hoc queries
- Quick selection of the required columns
- Sorting across the entire result set
- Access to all data through on-demand pagination

Launching Realtime Dataset Metrics

To launch RDM, right-click **Dataset Attributes System Summary**, then select **Dataset Metrics**.

> **Tip:** You can run Dataset Metrics in Zowe 2.16 after applying PTFs UO90057 and UO90058 or Zowe 3.0. The 5698-ZWH IBM Z Distribution for Zowe 3.0 is supported with installed IBM® Z OMEGAMON® AI for Storage® PTF UJ96869.
>
> 1. Log in to Zowe.
> 2. Log in to RDM for the chosen system.
> 3. Choose or add the RDM services that provide data from TEMS to RDM.
>   You can use the DNS or IP address where TEMS with completed Dataset Attribute collection is located and specify KS3_AS_PROTO as port.
>
>   You can also find the port number in the TEMS RKLVLOG by looking for `KS3S0166I KS3SDADW: HTTP SERVER ON PORT`. This message appears after Dataset Attribute collection is completed.

With PTF UJ04274, the Storage Toolkit functionality is available in RDM. You can use Storage Toolkit in RDM the same way, as you do in the TEP:

1. Right-click a dataset for available actions. The actions for saved filters are available in the saved filter set modal window.
2. To view the results of Storage Toolkit actions, open the Storage Toolkit view in the **View By** drop-down list, and click **Run Query**.

Selecting target system

Learn how to select a target system for Realtime Dataset Metrics (RDM).

PREREQUISITE

The target system is the system software libraries and other datasets that you are installing. You must select the target system to configure RDM for data retrieval.

Follow the next steps to specify the target system:

1. Log in to Zowe.
2. Open Dataset Metrics to see the server list of currently added RDM systems.
3. If your system is present in the list, select it and specify the credentials from the target system.
4. If the system is not present in the list, click the **Gear** button to open the **Configuration** dialog.
5. Click **Add server +**.
6. Specify the DNS or IP address and port for the system that has RDM configured.
   The port is also printed in the `RKLVLOG: KS3S0166I KS3SDADW: HTTP SERVER ON PORT = xxxxx STARTED.`

Realtime Dataset Metrics basic features

With the intuitive browser-based user interface of Realtime Dataset Metrics (RDM), you can easily retrieve and sort information about a dataset with its various fields and sections.

View By

With the **View By** drop-down menu, you can view a list of datasets organized as follows:

- **DASD Dataset**
  - Dataset (Dataset attributes)
  - Dataset Groups
  - Cluster Summary
  - HLQ (Dataset High-level Qualifier)
  - SMS Class (Management Class, Data Class, or Storage Class)
  - Storage Group
  - Catalogs
- **All Datasets**
  - All Datasets (summarization of data from DASD, tapes, and HSM)
  - HLQ (Dataset High-level Qualifier)
- **DASD Volumes**
  - Volumes
  - Storage Groups
  - User DASD Groups
  - User Storage Groups
- **Other**
  - Storage Toolkit

> **Note:** If you do not specify any filters, the **Volumes** option might cause long response time and high CPU usage. Add filters to reduce the result row count.

To review dataset groups and the summarized data for them in one place, select the **Dataset Groups** view from the **View By** drop-down menu.

To review custom groups for storage groups and the summarized data in one place, select the **User Storage Groups** view from the **View By** drop-down menu. Use this option to apply, copy, or delete saved filters. You can also submit Storage Toolkit actions for saved user storage groups the same way as with the **Dataset (Dataset attributes)** view. For more information, see:

- "Monitoring user storage groups" on page 79.
- "S3 User Storage Groups" on page 301.

Column Set

With the **Column Set** drop-down menu, you can switch between columns based on the context. For example, you can examine columns related to the following characteristics:

- Space
- VSAM
- VSAM Statistics
- GDG
- DFSMS
- Allocation
- Migration

When you select some **View By** options, such as **All Datasets** or **Dataset**, you can add your own column set.

To do so, click the **+** (plus) button and specify the column set name, and select the interesting attributes in the required order.

Filters

The **Filters** link opens a pop-up window where you can build ad hoc queries to filter the list of datasets by criteria including size, data last referenced, and other specifications. For example, you might want to view datasets larger than 100MB that are not referenced in the past 30 days. After selecting **Filters**, select **Add filter rule**.

You can then select a column to filter by (for example, **Percent Used** or **Allocated Tracks**), an operator (for example, **>** or **=**), and a threshold value. Finally, click **Save**. If you decide not to add the filter rule, click **Cancel**.

When **View By › All Datasets** is selected, additional Advanced Filters Mode is available for you to build filters by using AND and OR operators, brackets, and masks. Thus, you can build a filter with any complexity. Click **Load Example** for usage examples.

You can specify multiple filters, and Realtime Dataset Metrics will run the data through all of them before displaying the filtered results.

To remove a filter, click the **-** (minus sign) button.

To clear all filters, click the **Clear** button.

Save filters

To save the filter for future use, do the following:

1. Click **Save filters** in **All Datasets** or **Dataset**.
2. Specify the filter set name.
3. Optionally, specify the **Filter set group** to simplify future sorting of filters.

To create a group, enter **Filter set group** manually or select it from the drop-down list.

To review previously saved filter sets, click the **X saved** link under the **Save filters** button.

You can use the **Search** input to look for filters by name, and you can reduce the list of output by selecting groups in the **Filter set group** drop-down list.

The following actions are available with saved filters:

- **Copy to clipboard:** This action is useful when you copy filters from one Zowe instance to another or if you want to share the filter definition as text.

- **Delete:** This action deletes the saved group. You can also delete the **Filter set group** if this filter was the only one in that group.
- **Storage Toolkit:** This action can be issued against all datasets applicable to the saved filter.
- **Apply:** You can apply the saved filter and see datasets for it.

Run Query

To retrieve dataset information, specify the filters you need and click **Run Query**.

Results

After you run a query, the query results are displayed. If you want to see more data, you can page through the results either with the mouse wheel or with the page counter and arrows at the lower-left side, or with the **Page Up** and **Page Down** keys. You can sort by column headers, either in ascending or descending order.

Initially, only a subset of the results can be displayed to save processing time and display results more quickly. When this happens, the page counter says, for example, `1 of 15*`. This indicates that only the first 15 pages of data have been retrieved from a larger set of query results.

To see the total number of results, move the mouse pointer over the **of** indicator; in this example, `15*`. This step displays the following link, **Load more data** and **Get result count**.

When you click this link, the total number of results matching your query is displayed, and you can page further down through the result set.

If you reach the bottom of the displayed pages and want more data, select **Load more data** again.

You can enter your page, even if it was not loaded, to navigate the data on that page. Press **CTRL+ L** to open the**Select Line** dialog and go to the page where the requested line of data is located.

You can use **CTRL + Click**, **SHIFT + CTRL**, **CTRL + SHIFT + Click**, and **CTRL + A** to select the data on the page. With the **CTRL** button, you can select data on different pages. The following actions are available for selected rows:

- **Storage Toolkit:** You can issue Storage Toolkit actions for selected rows.
- **Copy Values:** You can right-click a column and select **Copy column value/s** to copy values to the clipboard. In **View By › All Datasets** , copied data can be added as values for a multi value filter.

Export

To save the displayed results to the dataset, click the **Export** button and enter the dataset name. The dataset is then allocated, and the data is exported in comma-separated values (CSV) format to the dataset, similar to using the **Dataset Attribute Group Extract** dialog in the Tivoli® Enterprise Portal.

When **View By › All Datasets** is selected, you can export data to the USS directory in `CSV`, `JSON`, or `XML` format.

When **View By › All Datasets** is selected and the **Summarization** dialog is open, you can export data to the USS directory in `CSV`, `JSON`, or `XML` format.

Inline editing

With inline editing, you can edit fields directly from a table, without using extra forms. This editing method is available for the Dataset Name, SMS Storage Class, and SMS Management Class attributes. Instead of using the menu and selecting **`Rename/Alter`** commands, you can make all changes from the same view. You can inline-edit a single line, or multiple lines at a time.

Single line

To inline-edit a single line, do the following:

1. Double-click an attribute field. If it is editable inline, it will be highlighted.
2. Make your edits.
3. Press **Enter** and confirm the changes to issue the **`Rename/Alter`** command.
   If an error was made during the edit, a corresponding dialog is displayed with an error message.

Multi line

To inline-edit multiple lines at once, do the following:

1. Use **SHIFT + Click** or **CTRL + Click** to select the lines to be edited.

2. Double-click the attribute field.
   If it is an editable attribute, a corresponding **Datasets to be changed** form is displayed with the selected datasets and attributes.
3. Edit the attribute values.
4. click **Apply**.

History mode

With the **History mode** option, you can review previously collected data for the following attribute groups:

- **DASD Dataset**
  - HLQ (Dataset High-level Qualifier)
  - SMS Class (Management Class, Data Class, or Storage Class)
  - Storage Group
  - Catalogs
- **DASD Volumes**
  - Storage Groups
  - User Groups

> **Note:**
> - To retrieve historical data by using **History mode**, history collection must be enabled for the corresponding attribute groups in History Collection Configuration via the enhanced 3270UI or Tivoli® Enterprise Portal interface.
> - Retrieving historical data for volumes can produce a lot of records. Consider using a filter with the Volser, Device Hex Address, Storage Group Name or Hardware Serial Number attributes to lower the record count and CPU usage.

To set **History mode**, do the following:

1. Click the Gear icon in the upper-right corner of the screen.
2. From the menu that opens, select **History mode**.
   The **History mode settings** pop-up window appears. In **History mode settings**, you can specify the time period for which you want the data to be retrieved.
3. Select one of the two options, **Last** or **Custom**:
   - Select **Last** to retrieve the data collected over the last one or more hours, days, weeks, months, quarters, or years (if the data is warehoused).
     You can specify the values as needed. The shortest available time period is one hour.
   - Select **Custom** if you want to specify a more precise time period.
     You can specify the **Start Date**, **End Date**, and **Time**. To do so, click the corresponding field, select the required date from the calendar, and enter the required time.
4. Click **Apply**.

For more information about history collection configurations, see:

- Mechanisms for collecting historical data
- Historical data collection
- Configuring near-term history

Storage Toolkit for Realtime Dataset Metrics

With the Storage Toolkit feature, you can use actions in the Realtime Dataset Metrics (RDM) interface.

Storage Toolkit actions provide access to mainframe operations to use target resources, such as volumes and mainframe datasets. For example, you use the actions to migrate datasets or release free space. You can use this feature with most of the RDM views.

> **Note:** These actions are only available in the Realtime Dataset Metrics application, they are not available in TEP.

To access user actions, do the following:

1. Select and right-click rows in the workspaces that have action menus.
2. Select **User Actions** from the context menu.
3. Choose the appropriate action from the submenu.
   When you select an action, a dialog box is displayed in which you configure the action and submit the request.

User action commands	
Name of menu item	**Corresponding dialog box**
New Batch Job	Create Batch Job template dialog box
New Command	New Command template dialog box

User action subcommands	
Name of menu item	**Corresponding dialog box**
Apply	Apply Create Batch Job or Command dialog box filled with template values
Edit	Edit Batch Job or Command template dialog box filled with template values
Delete	Confirmation dialog box

Create Batch Job template dialog box

Use this dialog box to create a custom Job template for later use. With this action you can prefill the values for the **Create Batch Job** dialog box as described in "Create Batch Job dialog box" on page 162 and save it with a custom name typed in the **Action name** field.

When saved, the action is displayed under the **User Action** context menu group with the following options:

- Apply - Submit user action.
- Edit - Change user action template.
- Delete - Delete the user action template.

New Command template dialog box

Use this dialog box to create a custom Command template for later use. With this action you can prefill the values for the Create Command dialog box.

When saved, the action is displayed under the **User Action** context menu group.

- Apply - Submit user action.
- Edit - Change user action template.
- Delete - Delete the user action template.

For related information, see "Using the Storage Toolkit" on page 146.

Session Jobs

The RDM interface has a special view of Session Jobs that represents and lists all the jobs committed during the current login session. If the redirect to result is enabled after submitting Storage Toolkit action, it automatically opens the Session Jobs view.

You can also access the Session Jobs view from the Storage Toolkit view for easy viewing of last job results. For more information, see "Storage Toolkit" on page 143.

User actions in Storage Toolkit for Realtime Dataset Metrics only

Learn about the dialog boxes that are related to Storage Toolkit user actions that you can create and save for later use in Realtime Dataset Metrics (RDM).

> **Note:** These actions are only available in the RDM application, they are not available in TEP.

To access user actions, do the following:

1. Select and right-click the rows in the workspaces that have action menus.
2. Select **User Actions** from the menu.
3. Choose the appropriate action from the submenu.
   When you select an action, a dialog box is displayed in which you configure the action and submit the request.

User actions commands	
Name of menu item	**Corresponding dialog box**
New Batch Job	Create Batch Job template dialog box
New Command	New Command template dialog box

User action subcommands	
Name of menu item	**Corresponding dialog box**
Apply	Apply Create Batch Job or Command dialog box filled with template values
Edit	Edit Batch Job or Command template dialog box filled with template values
Delete	Confirmation dialog box

Create Batch Job template dialog box

Use this dialog box to create a custom Job template for later use. With this action, you can prefill the values for the Create Batch Job dialog box, as described in "Create Batch Job dialog box" on page 162, and save it with a custom name typed in the **Action name** field.

When saved, the action is displayed under the **User Action** context menu group with the following options:

- Apply - Submit user action.
- Edit - Change user action template.
- Delete - Delete the user action template.

New Command template dialog box

Use this dialog box to create a custom Command template for later use. With this action, you can prefill the values for the Issue Command dialog box.

When saved, the action is displayed under the **User Action** context menu group.

- Apply - Submit user action.
- Edit - Change user action template.
- Delete - Delete the user action template.

For related information, see "Using the Storage Toolkit" on page 146.

All Datasets feature

Enabling the All Datasets feature

With the All datasets feature, you can view the summarization of data from DASD, tapes, HSM, and Dataset High Level Qualifier (HLQ). It is available in the View-by menu of Realtime Dataset Metrics.

PREREQUISITE

Ensure that the **Path to database** field exists in the USS directory to which the TEMS address space User must have access to.

1. Log in to Realtime Dataset Metrics.
2. Open the **View By › All Datasets**.
3. Open **Collection control**.
4. Set the value to the **Path to database** field.
5. Select at least one source for data collection:
    - Datasets from DASD
    - Datasets from the Tape Management System (DFSMSrmm or CA 1)
    - Datasets from DFSMShsm

Collection control and History mode

Depending on your collection mode, you can choose to collect real-time data, schedule data collection, or review previously collected data.

Collection control

In the Collection control dialog, you can schedule the collection to run on the specific days at a specific time and list the interested collection types.

- **Datasets from DASD:** This collection triggers the existing Dataset Attribute collection to get information about datasets from the DASD farm.
  If you select the **Datasets from DASD** source type, Dataset Attribute collection starts before the All Datasets collection.

  If Dataset Attribute collection was not started earlier, the new Full collection gets started.

  If Dataset Attribute collection is completed - the Incremental collection starts to get fresh data.

- **Datasets from Tape Management System:** Depending on the system configuration, this collection triggers the existing DFSMSrmm or CA 1 collections.
  If you select **Datasets from Tape Management System**, the corresponding DFSMSrmm or CA 1 collections gets started.

- **Datasets from DFSMShsm:** This option starts the collection which reads the DFSMShsm Control Datasets for the backed-up and migrated datasets.

All Datasets collection consumes on average ~ 1 CPU Second per ~ 22'000 datasets.

You can see the status of the previous collection. You can use **Force start** to run the collection immediately. If you move the mouse to the **Force start** section, you see a time estimation for the next collection.

History mode

To review previously collected data in **View By › All Datasets**, click the **Gear** icon and select **History mode** to open the History mode setting dialog where you can perform the following actions:

- **Filter data sources by type:** Find all previous collections that contain selected collection types.
- **Max count of data sources:** Limit the number of stored collections by changing this number. 0 means no limit.

- **Select data source time:** Select the collection from the specified date and time and click **Apply** to see the data from the selected collection.

Storage Toolkit

With the Storage Toolkit feature, you can use actions in the Tivoli® Enterprise Portal (TEP), enhanced 3270UI, or the Realtime Dataset Metrics (RDM) interface.

By using Storage Toolkit actions, you get convenient access to mainframe operations and can target resources such as volumes and mainframe datasets. For example, you can use actions to migrate datasets or release free space.

With the Situation Storage Toolkit Extension enhancement, you can associate a Storage Toolkit command with a situation. This means that the Storage Toolkit command will be run when the situation becomes true. All of the capabilities of Storage Toolkit can be associated with a situation including the ability to issue custom commands, batch JCL, and Storage Toolkit generated commands. See "Associating Storage Toolkit commands with situations" on page 192 for more information.

You can initiate a broad range of commonly used mainframe commands by using product-provided menus and dialog boxes. You can also issue custom commands. Using Storage Toolkit, you can invoke commands from the following facilities in a z/OS® system:

- DFSMSdss (Device Support Services)
- DFSMShsm (Hierarchical Storage Manager)
- ICKDSF (Device Support Facilities)
- IDCAMS (Access Method Services)
- DFSMSrmm™ (DFSMS Removable Media Manager)
- Mainframe (z/OS® console command)
- TSO (Time Sharing Option)

With Storage Toolkit, you can also create and submit user-defined JCL streams.

Storage Toolkit menus and dialog boxes

This section describes the actions that are available in the context menus of Storage Toolkit.

To access the actions, select and right-click rows in the workspaces that have action menus. A context menu is displayed where you select an item such as **Dataset Actions**. Then you choose the appropriate action from the submenu. When you select an action, a dialog box is displayed where you configure the action and submit the request.

The items on the context menu vary depending on the selected resources, however the following items are always available:

- Create Command
- Create Batch Job
- Submit Command or Job

These items do not have submenus. When you select them, you proceed directly to the corresponding dialog box.

Some actions can process multiple resources in a single request. In other words, you can run a single action against multiple resources. For example, to migrate multiple datasets, you can proceed as follows:

1. Select the target resources in the Tivoli® Enterprise Portal or the RDM interface.
2. Right-click and select **Dataset Actions > Migrate**.
3. Configure the command.
4. Submit the request.
   The **`Migrate`** command is issued against each resource.

When you select an action, Storage Toolkit launches a dialog box where you configure the action. The dialog boxes contain a set of standard tabs, which are described in "Standard tabs in dialog boxes of Storage Toolkit" on page 157.

Custom commands and jobs

Name of menu item	Corresponding dialog box
Create Batch Job	"Create Batch Job dialog box" on page 162
Create Command	"Issue Command dialog box" on page 165
Submit Command or Job	"Submit Command or Job dialog box" on page 166

User action commands

Name of menu item	Corresponding dialog box
New Batch Job	Create Batch Job template dialog box
New Command	New Command template dialog box

User action subcommands

Name of menu item	Corresponding dialog box
Apply	Apply Create Batch Job or Command dialog box filled with template values
Edit	Edit Batch Job or Command template dialog box filled with template values
Delete	Confirmation dialog box

Actions related to volumes that are displayed in workspaces of the Tivoli® Enterprise Portal

Name of menu item	Corresponding dialog box
Backup	"Volume Backup dialog box" on page 167
Convert	"Volume Conversion dialog box" on page 167
Migrate	"Volume Migration dialog box" on page 168
Release Space	"Volume Space Release dialog box" on page 168
Compress	"Compress dialog box" on page 168
Defrag	"Defrag dialog box" on page 169

Actions related to datasets that are displayed in workspaces of the Tivoli® Enterprise Portal

Name of menu item	Corresponding dialog box
Backup	"Dataset Backup dialog box" on page 170
Migrate	"Dataset Migration dialog box" on page 171
Move & Copy	"Dataset Copy and Move dialog box" on page 171
Recall	"Dataset Recall dialog box" on page 172
Recover	"Dataset Recover dialog box" on page 172
Release Space	"Dataset Space Release dialog box" on page 172
Compress	"Dataset Compress dialog box" on page 173
Print	"Print dialog box" on page 184
Allocate	"Dataset Allocate dialog box" on page 173
Rename	"Dataset Rename dialog box" on page 182
Alter	"Dataset Alter dialog box" on page 174
Delete	"Dataset Delete dialog box" on page 179
Listcat	"Dataset Listcat dialog box" on page 180
Repro	"Dataset Repro dialog box" on page 182
Verify	"Dataset Verify dialog box" on page 184
Catalog	"Dataset Catalog dialog box" on page 179
Uncatalog	"Dataset Uncatalog dialog box" on page 184
Define Cluster	"Dataset Define Cluster dialog box" on page 175

Actions related to dataset attribute groups that are displayed in workspaces of the Tivoli® Enterprise Portal and the enhanced 3270UI

Name of menu item	Corresponding dialog box
Extract Attributes	"Dataset Attribute Group Extract dialog box" on page 170

Actions related to DFSMShsm

"Cancel HSM Request dialog box" on page 186.

The menu item that corresponds to this dialog box is available when all the selected items are either queued or active. If the state of any selected item is Complete, Abended, or Cancelled, the menu is not available.

Actions related to DFSMSrmm

Name of menu item	Corresponding dialog box
Actions related to DFSMSrmm-related datasets	
Add Dataset	"Add Dataset dialog box" on page 187
Change Dataset	"Change Dataset dialog box" on page 187
Delete Dataset	"Delete Dataset dialog box" on page 188
Actions related to DFSMSrmm-related owners	
Add Owner	"Add Owner dialog box" on page 188
Change Owner	"Change Owner dialog box" on page 188
Delete Owner	"Delete Owner dialog box" on page 188
Actions related to DFSMSrmm-related volumes	
Add Volume	"Add Volume dialog box" on page 188
Change Volume	"Change Volume dialog box" on page 189
Confirm Actions	"Confirm Actions for All Volumes dialog box" on page 189
Delete Volume	"Delete Volume dialog box" on page 189
Actions related to DFSMSrmm-related VRS resources	
Add VRS	"Add VRS dialog box" on page 189
Change VRS	"Change VRS dialog box" on page 190
Delete VRS	"Delete VRS dialog box" on page 190

The documentation for the z/OS® operating system includes reference information for the commands and parameters that Storage Toolkit uses. You can access documentation for the version of z/OS® that you use at http://www-03.ibm.com/systems/z/os/zos/

How Storage Toolkit processes action requests

This section briefly describes how the IBM® Z OMEGAMON® AI for Storage® monitoring agent handles an action request.

Note: You can use Storage Toolkit to target resources in a separate z/OS® environment. However, this activity is valid only when that separate environment is located on DASD that is shared between the separate z/OS® environments.

All actions, not only the actions that submit user-defined JCL streams, run as batch jobs on the z/OS® system. Each type of command has a batch job interface which Storage Toolkit uses to execute the action.

The JCL for each command type is defined by PARMGEN and is subsequently used to run the command. PARMGEN is also used to define a default JOB card which is applied to each batch job but can be overridden with a custom JOB card that you provide on the Storage Toolkit dialog box.

When you submit an action request, Storage Toolkit starts a thread on z/OS® to handle the execution of the request. The thread generates the JCL, substituting values for variables and expanding `)DOT` `)ENDDOT` clauses, submits the batch job for execution using the MVS™ Internal Reader, and waits for its completion.

At the beginning of each JCL stream, Storage Toolkit inserts a step that executes code to monitor the rest of the steps in the job. As each step completes, this code collects command output or, in the case when your request is to run user-defined JCL, it collects the contents of corresponding datasets and DD names that you specified to be copied for later viewing. DD names that are routed to a SYSOUT class are not collected until the final job step.

At the end of each JCL stream, Storage Toolkit adds a final step that also executes the Storage Toolkit code. The code in the final step collects SYSOUT from each step where it was requested, collects the return code of the batch job, and optionally, collects the JES output. Additionally, it notifies the Storage Toolkit thread that the batch job is complete. The thread then stores the results of the execution into the checkpoint database and terminates.

To accomplish this, the Storage Toolkit thread creates temporary datasets to contain the JCL, the results of the batch job, and other files. The datasets are transient and are shared between the thread and the batch job. When the batch job completes, the datasets are deleted.

Activities required to process your request are executed under a security environment that Storage Toolkit establishes using your Tivoli® Enterprise Portal user ID. The security environment is established on the z/OS® system where the monitoring agent runs, that is why your portal user ID must match your mainframe ID on that system. If this ID is not a valid mainframe user ID, a security environment cannot be established, and the action request fails.

Storage Toolkit performs the following operations within the context of your security environment:

- Create and access all temporary datasets required by the batch job.
- Submit the batch job.

When your action request runs a user-defined JCL, Storage Toolkit also uses your security context to access the JCL dataset and any other datasets that you identify as containing substitution variables. If you are not authorized to the datasets, the action request fails.

> **Note:** To ensure the security of Storage Toolkit commands, some type of security authentication must govern the users who log on to the Tivoli® Enterprise Portal. You set up authentication on the hub monitoring server in the IBM® Tivoli® Monitoring environment. For further information, see "Authentication and authorization for running Storage Toolkit commands" on page 160.

Using the Storage Toolkit

This section includes the scenario for using the Storage Toolkit, the process of using actions in the Storage Toolkit, and the Storage Toolkit limitations and hints.

This section describes the general process of using the Storage Toolkit and a typical scenario for use of the toolkit.

Scenario for using the Storage Toolkit

The Storage Toolkit enables the following scenario:

Phase One:	You create a situation that monitors available free space in your storage groups. Whenever the situation is triggered, you want to find the affected resources and take action. Specifically, you want to find and take action on the datasets with the most unused space.

Phase Two:	You work in the Tivoli® Enterprise Portal, instead of accessing the z/OS® system and performing a manual process on each resource: 1. Navigate to an appropriate workspace for the product, such as the Largest Datasets in SMS Class, and sort on Unused Tracks. 2. Select datasets that can be migrated. **Note:** In addition to the migration command, the Storage Toolkit offers a range of commands that is listed in the "Storage Toolkit menus and dialog boxes" on page 143 topic. You can also submit user-defined JCL streams, and you can issue mainframe commands that are not already available in the dialog boxes that the product provides. 3. Right-click the selected datasets and choose **Dataset Actions > Migrate** in the pop-up menu. 4. Use the Storage Toolkit dialog boxes (accessed through a right-click menu) to issue DFSMShsm™ migrate commands for all selected datasets. The DFSMShsm™ migrate command you selected from a single dialog box is issued against each dataset that you selected. 5. View the results of the migration commands in one of the Storage Toolkit workspaces.

Process for using actions in the Storage Toolkit

The following steps describe the general process of using the actions provided by the Storage Toolkit. The options that are displayed in the context menus change, depending on what resources you select in a workspace view in the Tivoli® Enterprise Portal or the RDM interface. For example, if you have selected datasets in a workspace view, you see the **Dataset Actions** option in the context menu.

1. Select one or more rows from a view in a workspace. The rows correspond to resources such as volumes and datasets on z/OS® systems.
2. Right-click within the rows. A context menu is displayed where you can select an appropriate option. The following example steps describe how to invoke the appropriate Storage Toolkit dialog box when you select *datasets* in a workspace view and right-click within the rows:
   a. Select **Dataset Actions** from the context menu to access the submenu.
   b. Select an appropriate dataset action in the submenu, such as **Backup**. The corresponding dialog box is displayed.

   The menu displays different options if you are working with other types of resources, such as volumes. In most cases, the action can simultaneously target multiple resources that you select in Step 1.
3. Make selections in the dialog box to configure the command that you want to issue. These selections are automatically transformed to command syntax that is valid for the mainframe environment. See "Storage Toolkit menus and dialog boxes" on page 143 for links to descriptions of all available dialog boxes.

Reusing action requests

The Storage Toolkit provides three methods for you to reuse an action. Consider whether you want to rerun the action using the same or different resources before you choose the method.

When you reuse an action, the new execution is processed under your security context on the z/OS® system. In the following descriptions, it is assumed that you are authorized to the resources and command associated with the reuse action.

Method 1: Reusing an existing action using the same resources

1. Navigate to the Storage Toolkit Action Requests workspace (available from the Storage Toolkit Navigator node).
2. Right-click the action in the Action Requests view that you want to reuse, and select **Submit Request** from the context menu. The dialog box corresponding to the selected action is displayed.
3. Modify the settings that control how the request runs as appropriate.

4. Click **OK**. The selected action request is submitted on the z/OS® system to which you are connected, and the same set of resources is affected by the action. If the selected action request is associated with groups, then the current volume or dataset resources in that same set of groups are affected by the action.

Method 2: Reusing an existing action using the same or different resources

1. Select one or more rows from a view in a workspace.
2. Right-click within the rows, and select **Submit Command or Job** from the context menu. From the **Submit Command or Job** dialog box, select the action you want to reuse. The dialog box corresponding to the selected action is displayed.
3. Modify settings as appropriate. If you want to reuse the action using the rows you selected in Step 1, ensure that the **Use request's default data** option in the **General** tab is not selected. If you want to reuse the action using the resources that were used when the action initially ran, select the **Use request's default data** option.
4. Click **OK**. The action request is submitted on the z/OS® system to which you are connected. The resources that are affected depend on the setting of **Use request's default data**. Additionally, if the **Use request's default data** option is selected and the action initially ran against a set of groups, then the current volume or dataset resources in that same set of groups are affected by the action.

Method 3: Reusing a specific execution of an action

When you rerun a specific execution of an action, the resources that are used are the same as those associated with the selected execution.

1. Navigate to the Storage Toolkit Action Requests workspace (available through the Storage Toolkit Navigator node).
2. Find the action that you want to reuse, and link to the Storage Toolkit Result Summary workspace.
3. Right-click the specific result in the Action Results Summary view that you want to reuse, and select **Submit Request** from the context menu. The dialog box that corresponds to the selected action is displayed.
4. Modify settings as appropriate.
5. Click **OK**. The selected action request is submitted on the z/OS® system to which you are connected. The same set of resources used by the selected execution are affected by the action. If the selected execution is associated with groups, then the current volume or dataset resources in that same set of groups are affected by the action.

Action requests in workspaces that you create

If you create a workspace view, the Storage Toolkit action menus are available in the new view when the following conditions are true:

- Your workspace uses one of the queries that come with the product.
  AND
- The query originally had a Storage Toolkit action menu associated with it. These menu options are listed in "Storage Toolkit menus and dialog boxes" on page 143.

About results for action requests

The result for action requests can contain output from one or more steps or locations. For this reason, the **Step Name** and **Dataset or DD Name** attributes identify the source of the details.

- When you run a command from a command dialog or issue a command from the **Issue Command** dialog, the values for the **Step Name** and **Dataset or DD Name** are predefined by the Storage Toolkit.

Command type	Step Name	Dataset or DD Name
DFSMSdss	KS3DSS1	SYSPRINT
DFSMShsm	KS3HSM1	KS3OUT
IDCAMS	KS3AMS1	SYSPRINT
ICKDSF	KS3DSF1	SYSPRINT
DFSMSrmm	KS3RMM1	KS3OUT

Command type	Step Name	Dataset or DD Name
TSO	KS3TSO1	KS3OUT

- When you create JCL using the **Create Batch Job** dialog box, the values for the **Step Name** and **Dataset or DD Name** reflect the values that you specified in the **Files whose contents should be copied for later viewing** fields on the **Options** tab.
- When you select **Copy JCL and JES logs for later viewing** on the **JCL** tab, the **Step Name** values are predefined by the Storage Toolkit. The **Dataset or DD Name** values are blank.

Step Name	Description
*JCL	The JCL that was submitted
*JESMSGLG	The contents of the JESMSGLG log
*JESJCL	The contents of the JESJCL log
*JESYSMSG	The contents of the JESYSMSG log

No output from Mainframe commands is returned, because execution of the command is not synchronized with the execution of the batch job. Other reasons why you might not receive output include the following:

- The command that you are running might not return output.
- A file or log to be copied for later viewing is empty.
- You specified `0` as the value for the **Maximum output lines** option.
- The output is not available or the Storage Toolkit cannot capture it.

In some cases, you might find that command output can be captured using a different technique to run the command. For example, if you use the **Issue Command** dialog box to run the DFSMShsm **LIST** command, the Storage Toolkit is not able to capture the output. However, you can write JCL that uses `HSENDCMD WAIT` to issue the **LIST** command and route the output to a dataset. Using the **Create Batch Job** dialog box, you can then run this JCL and request the dataset to be copied for later viewing.

Checking for the completion of action requests

You might need to check the status of the toolkit action request between the time you click **OK** in the toolkit dialog box and the time that the job completes in the mainframe environment. In some cases, it can be unclear whether a job is taking a long time to run or has failed. The status of the job might be **EXECUTING** *in both cases*. You can respond to these possible conditions as described in this scenario:

Conditions of the scenario, in sequential order	Your response options
You notice a request on the Tivoli® Enterprise Portal that remains in **EXECUTING** state without completing.	Check the status of the batch job in the z/OS® system.
Problem scenario 1: The batch job has failed due to a JCL error.	Cancel the request on the Tivoli® Enterprise Portal. This action releases the thread and the resources that have been invoked. (*Optional*) Attempt to fix the JCL error and run the request again.
Problem scenario 2: The batch job requires more time to complete.	• Consider raising the priority of the job. • Consider adjusting the times when the job is allowed to run, if the job has been scheduled to run at off-peak hours.

Storage Toolkit limitations and hints

This topic describes limitations, hints and other considerations for the Storage Toolkit. You can avoid problems by remaining aware of these issues. Also remain aware of other issues documented throughout this chapter.

- **Datasets not locked:** The Storage Toolkit enables you to click the **Edit JCL** button in the Create Batch Job dialog box, and edit a dataset that is located in the mainframe environment. However, you must ensure that no one accesses the dataset at the same time, because the dataset is not locked during your editing session. Another user could edit the dataset, for example, in TSO using ISPF, during your editing session. That user would not see a "dataset in use" message; if he saves his changes before you save yours, your

changes overlay his. If that user is still editing the dataset when you attempt to save your changes, the attempt to save your changes fails. The results of each editing session might be unpredictable depending on the scenarios and editing tools that the other user uses.

- **EBCDIC code page 037 only:** The JCL editor provided by IBM® Z OMEGAMON® AI for Storage® supports EBCDIC code page 037 only. The editing or authoring of JCL in other code pages such as EBCDIC code page 930 (in other words, Japanese EBCDIC) is not supported.
- **Storage Toolkit first and last steps:** In every batch job, the Storage Toolkit inserts a Toolkit-specific first step that sets up a monitor over the steps in the job. The monitor collects targeted output from the steps in the job. The Storage Toolkit also appends a last step to the end of the job. The last step collects targeted SYSOUT from the previous steps along with the JES output and the return code. It also notifies the monitoring agent that the batch job is complete. The Toolkit-specific first step is inserted immediately before the first EXEC, PROC, or INCLUDE statement that it locates in the JCL.

  > **Note:** If your JCL has an INCLUDE statement before an EXEC or PROC statement, the INCLUDE member must not contain JCL statements, such as JCLLIB, that must precede the first job step. Because the Toolkit first step is inserted before the INCLUDE, the batch job fails in this case.

- **Using the null statement to mark the end of a job:** When you use the Create Batch Job dialog box to execute user-defined JCL, the Storage Toolkit generates a copy of your JCL to which it adds a Toolkit-specific first step and a Toolkit-specific last step. If your JCL ends with the null JCL statement denoting end of job,

  ```
  //
  ```

  that null statement is removed from the generated JCL because the job does not end until after the Toolkit-specific last step runs.
- **Conditional processing:** Do not specify a COND parameter on the JOB card or on an EXEC statement that might cause the Storage Toolkit steps that were inserted into the JCL to not run. If you use the COND parameter or you use IF/THEN/ELSE/ENDIF statements, you must ensure that the Storage Toolkit first and last steps run.
- **No support for multiple jobs in user-defined JCL:** When you use the Create Batch Job dialog box to execute user-defined JCL, the JCL must not include multiple jobs. The Storage Toolkit does not support this type of JCL. Results are unpredictable.
- **JOB card:**
  - When you use the Create Batch Job dialog box to execute user-defined JCL, the batch job is submitted using the Replacement JCL JOB card which you specify on the **JCL** tab. This overrides a JOB card that might be present in the JCL. If you do not specify a Replacement JCL JOB card, your installation-specific JOB card is used.
  - Do not specify a CLASS or a TYPRUN option on your JOB card that just copies or scans the job. Because the batch job does not execute, you action request remains in the EXECUTING state. You must cancel the request to release the thread and resources associated with it and to remove it from EXECUTING state.
  - Do not specify a COND parameter on the JOB card that might cause the Storage Toolkit first and last steps inserted into the JCL to not run.
  - When you request the JES output be copied for later viewing, make sure the MSGLEVEL on your JOB card is set to the level of output that you desire.
  - When you specify your JOB card, consider assigning it a unique job name. If the name matches a batch job that is executing on your z/OS® system, your job might be delayed until the executing job completes. To prevent this, assign a unique name to your job.
- **Using a PROC in user-defined JCL:** When you use the Create Batch Job dialog box, you can run JCL that executes a procedure, however the Storage Toolkit might not be able to properly copy the contents of files associated with steps in that procedure:
  - The procedure can be instream or in a system or private procedure library. If you use a private procedure library, you must ensure the JCLLIB statement precedes the Storage Toolkit first step.

- You can request files that are referenced in the procedure be copied for later viewing, but with certain limitations:
  - The step name is the step name in your JCL that executes the procedure. You cannot specify the step names that are in the procedure itself.
  - If the procedure consists of a single step, the contents of the requested files is returned.
  - If there are multiple steps in the procedure, the contents of a requested dataset or DD name that references a dataset is returned for each procedure step (in other words, multiple times). The contents of a DD name that is routed to SYSOUT is returned for each procedure step in which the SYSOUT DD name is defined (ie., one or more times).

- **Use step names in user-defined JCL:** Do not include steps in your user-defined JCL without step names, if you intend to copy files for later viewing that are associated with those steps. The Storage Toolkit requires a step name.
- **/*XMIT:** Do not use the /*XMIT statement in any of your JCL. The Storage Toolkit does not support this. Results are unpredictable.
- **DYNAMNBR:** If you submit user-defined JCL that allocates datasets, be aware that the Storage Toolkit allocates datasets in each step, too. If necessary, you might need to use the DYNAMNBR parameter on your EXEC statement to allow for your datasets and 3 Storage Toolkit datasets.
- **JCL errors:** The Storage Toolkit last step appended to the end of each batch job notifies the monitoring agent when the batch job completes. When this occurs the action request that is pending completion of the job is updated with the results of the batch job. If the last step in the batch job does not run, for example, the batch job failed with a JCL error or conditional processing bypassed the last step, the action request remains in an EXECUTING state. If you determine that an action request is in EXECUTING state longer than you anticipate, you must check the status of the batch job on the z/OS® system. If the job failed such that the last step did not run, you must cancel the execution of the action request in the Tivoli® Enterprise Portal. This releases the thread and resources associated with the request and removes it from EXECUTING state. You can then determine why the job failed, correct the error and resubmit the request.
- **Return codes:** Certain return codes, which are generally paired with a status, are set by the Storage Toolkit when it detects an error processing an action request. The following table lists common return codes and their corresponding status:

Common Storage Toolkit return codes

Return code	Status
117	This status typically indicates that the JCL exceeds 72 characters when the substitution variables are applied. It might also indicate other JCL-related errors, such as a missing JOB card, or another dataset requiring variable substitution that exceeds 80 characters when the substitution variables are applied.
119	The User Data Server has ended abnormally or the batch job has ended, but the Storage Toolkit is unable to determine the return code.
121	Authorization has failed.
123	A dataset error has occurred, such as the dataset containing your JCL does not exist. Messages in the RKLVLOG of the Tivoli® Enterprise Monitoring Server might help you analyze this result.

Return code	Status
125	The execution of the action request was stopped because the action request was associated with a set of groups that did not exist or were empty at the time of execution. The status that is displayed in the Result Summary workspace is set to one of the following values: **NonexistentGroups** Indicates that none of the groups associated with the request existed. **EmptyGroups** Indicates that all of the groups associated with the request were empty. **BadGroups** Indicates that a combination of empty and missing group errors affected the entire set of groups. This status value might also indicate that some other error was detected when the groups were processed. Review the messages in the RKLVLOG to assist your analysis of the results. **Note:** Groups might be empty because a collection is running or has not yet run. If so, retry the request when the collection completes.

If you see a return code that does not look familiar, you might want to convert it to its hexadecimal equivalent, because the code might indicate an abend in the batch job. For example, a return code of `193` is the same as `x'0C1'`.

- **Variable substitution and line limits:** Substitution variables that are defined through the Storage Toolkit are replaced when the action request runs. The Storage Toolkit creates temporary datasets to contain the updated statements.
  There are two basic categories of datasets (JCL and Other) that are updated with substitution variables. The Storage Toolkit processes them as follows:
  - **JCL dataset:** Variable substitution is applied to all components of the batch job, including the JOB card, extra JCL, the Toolkit-specific steps, and the body of your JCL along with any instream datasets. The Storage Toolkit interprets column 72 as a continuation character and preserves its contents. The data between columns 2 and 71 might shift left or right depending on the size of the variable name and its substitution value. If the data shifts beyond column 71, the request fails. The return code for the request is set to 117 and the status for this execution is**InvalidJCL** . You must perform these actions:
    - Verify the substitution variables and values are correct and do not have unintentional consequences to the components of the batch job
    - Correct the JCL to ensure that no line exceeds the limit.

    The action request may also fail with return code 117 if any lines in the original components of the batch job exceed 80 characters.
  - **Other dataset:** Variable substitution is applied to all records in other datasets that you specify as needed by the job that also contains substitution variables. The toolkit makes no assumptions about the contents of the dataset and considers each line, from column 1 to column 80, as a line of data. Variable substitution may cause the data in columns 2 through 80 to shift left or right depending on the size of variables names and their values. If the data shifts beyond column 80 (excluding trailing blanks), the request fails. The return code for the request is set to 117 and the status for this execution is InvalidJCL. You must perform these actions:
    - Verify that the substitution variables and values are correct and do not have unintentional consequences to the contents of the dataset.
    - Correct the contents of the dataset to ensure that no line exceeds the limit.
- **Validating JCL:** When you write JCL for use in the Create Batch Job dialog box, always check the validity of the statements before you submit the batch job. For example, when you edit JCL in the Edit JCL dialog box, consider whether line lengths will exceed the 72-byte limit after variable substitution is performed. When substitution variables are replaced in the JCL at execution time, resultant JCL lines that contain more than

72 bytes causes the JCL to not be submitted. A status of InvalidJCL is displayed in the Result Summary workspace for the action request.

- **Reserved variable names:** The Storage Toolkit reserves the following variable names. You must not use these names for your own variables:

```
%%KS3TK_CMD_DSN%%
%%KS3TK_HSM%%
%%KS3TK_DYNAMNBR%%
```

- **Fully Qualified Datasets needed by the job that also contain substitution variables:** When you use the Create Batch Job dialog box, you can specify additional datasets that contain substitution variables. The Storage Toolkit creates a temporary dataset with the updates and replaces the name of the original dataset with the temporary one in its copy of your JCL. In order for the names to be replaced, the datasets must be referenced in your JCL; they cannot be in cataloged procedures or INCLUDE members that your JCL might use.
- **JES output:**
  - The techniques that the Storage Toolkit uses to collect JES logs and system-handled output datasets (SYSOUT) require your z/OS® operating system use JES2.
  - Because the Storage Toolkit last step collects JES output just before the batch job ends, some of the messages you normally see in the JES logs such as the job start (IEF375I) and job end (IEF376I) messages are not included in the JES output.
- **Mainframe commands:**
  - Mainframe console commands are submitted through an SDSF batch job interface. A forward slash (/) must precede the command, as in this example, which cancels a time-sharing user (*tso_user_ID*):

    ```
    /C U=tso_user_ID
    ```

  - Command output is not returned for Mainframe console commands because execution of the command is not synchronized with execution of the batch job.
  - Because execution of the command is not synchronized with execution of the batch job, the return code associated with the action request reflects the submission of the command to the SDSF batch job interface. It does not reflect the execution of the command itself.
  - Because the Storage Toolkit uses SDSF, your z/OS® operating system must use JES2.
- **Shared DASD:** The temporary datasets that the Storage Toolkit creates to contain the generated JCL, the results dataset, and other files are shared between the Toolkit and the batch job. Because the batch job can run on a z/OS® system in your SYSPLEX different from the one where the monitoring agent runs, the temporary datasets must be created on DASD shared across the systems. Your installation can control the location of the temporary datasets using options in PARMGEN. These options also control the location of datasets created using the Edit JCL option in the Create Batch Job dialog box.
  In addition, when you use the Create Batch Job dialog box, you specify the dataset containing the JCL you wish to submit and, optionally, specify datasets needed by the job that also contain substitution variables. These datasets must be cataloged and located on on-line DASD that is accessible to the z/OS® system where the monitoring agent runs.
- **APF-authorized load library on remote systems:** The Storage Toolkit inserts a first step and last step into every batch job. These steps run Toolkit code that is located in the TKANMODL load library for your installation's run time environment. The load library must be APF-authorized. If the batch job runs on the same z/OS® system as the monitoring agent, the load library is normally already APF-authorized. If your batch job runs on another z/OS® system in your SYSPLEX, you must ensure the load library is APF-authorized on that system as well. The load library must also be located on DASD that is shared across the systems.
- **Unprintable characters:** You must ensure that the files you specify as Files whose contents should be copied for later viewing on the Create Batch Job dialog box or the output from a command on the Issue Command dialog box will contain character data only. If the files or command output contain unprintable characters (for example, hexadecimal data), these characters might not display properly in the Storage Toolkit Result Detail workspace.

- **Checkpoint dataset storage exhausted:** When you submit an action request, information about the request is stored in the checkpoint database. When the request completes, results of the execution are also stored. The information stored in the checkpoint database includes elements such as:
  - The name and description of the request
  - The time the request was submitted and the time it completed
  - The resources associated with the request
  - The return code from the execution of the request
  - The output produced by the execution of the request, which might include:
    - Command output
    - Files copied for later viewing
    - The submitted JCL
    - The JES files produced by the batch job.

  The Storage Toolkit checkpoint database is allocated and initialized when the runtime environment is created. This database is never extended. You must ensure that sufficient space exists for the activities of the database. Otherwise, the database can run out of space. You can remove the requests and the results when you no longer need them. You can also use options on the **General** tab of the Storage Toolkit dialog boxes to help you manage the results.

  > **Note:** If the results from the execution of an action request exceed the free space available in the checkpoint database, the output is lost entirely. The error message KS3T830E SERVICE CHECKPOINT DATASET STORAGE EXHAUSTED in the RKLVLOG of the Tivoli® Enterprise Monitoring Server indicates this condition. The *IBM® OMEGAMON® for Storage on z/OS®: Troubleshooting Guide* provides further information about this issue.

Using substitution variables in the Storage Toolkit

This section explains how to use substitution variables with the Storage Toolkit and provides examples.

You can use substitution variables to indicate portions of a Storage Toolkit command or batch job that require substitution with volume names, dataset names, or other values. "Using the Storage Toolkit" on page 146 provides helpful information about using substitution variables.

Setting substitution variables for toolkit commands

Many of the action requests that the Storage Toolkit generates can include substitution variables. You can define these variables in the **Command** tab in Step "6" on page 155 as part of the following overall suggested procedure:

> **Note:** A **Command** tab is not present in the Create Batch Job dialog box. In that dialog box, you set substitution variables in the **Options** tab.

1. Right-click rows in one of the views of a workspace to access the pop-up menu.
2. Select the Storage Toolkit command that you want to invoke. The **Options** tab is selected by default in the dialog box.

   > **Note:** An **Options** tab is not provided in the Issue Command dialog box, because you use this dialog box to create arbitrary free-form command and parameter strings. If you are using this dialog box, go to Step "6" on page 155

3. Make selections and enter values in the **Options** tab, as appropriate.

4. Select the **Command** tab. The **Command** field shows how the options that you selected in the **Options** tab are rendered in raw text within a command.
5. Before you make any changes in the **Command** tab, inspect the command parameters and syntax that are displayed to ensure that they match your expectation. If necessary, modify the command using either of the following approaches:
   - *Before you make any changes in the **Command** tab*, return to the **Options** tab and modify your selections.

     > **Note:** If you make changes in the **Command** tab, the **Options** tab is disabled, and you cannot modify the selections that you made there. To be able to work with the **Options** tab again, you can open a new instance of the toolkit dialog box. Click **Cancel** to dismiss the current instance of the dialog box, and access the dialog box again through the Tivoli® Enterprise Portal.

   - If you are familiar with command syntax, edit the raw text that you see in the **Command** field of the **Command** tab. Keep in mind that after you make changes in the **Command** tab, you cannot make further modifications to the **Options** tab.
6. If you want to define substitution variables, access the **Command** tab and define them in the **Substitution variables and their runtime replacement values** area. In many cases, the substitution variables and their runtime replacement values table are primed with variables and values defined by the Storage Toolkit based on the action and resources that you selected.

   > **Note:** A **Command** tab is not present in the Create Batch Job dialog box. In that dialog box, you set substitution variables in the **Options** tab.

7. Click the **Show Data** button to access the "Data View dialog box" on page 161 to confirm that the correct volumes or datasets are targeted by the substitution variables.

Standard substitution variables in the toolkit

The following general guidelines apply to the substitution variables that the toolkit uses:

- IBM® Z OMEGAMON® AI for Storage® uses a pair of percent signs (%) surrounding the `SUB`*nn* name to identify the substitution variables that it provides, such as `%SUB01%`.
- When you create variables for JCL files, it is recommended that you use the percent-sign (%) convention, as in this example: `%MY_VAR%`. This convention enables the Storage Toolkit to find these variables and list them automatically in the "Create Batch Job dialog box" on page 162. Variables that do not follow this convention are not listed automatically in the dialog box.
- Do not create a variable that could be a subset of some other variable. For example, do not name a variable **MY_VARNAME** and another variable **MY_VAR**.
- Do not create a variable that could match data in your command, JCL files, or generated JCL that should not be replaced. For example, in the following scenario a value might be replaced that should not be replaced:
  a. You define an **ADDR** variable and assigned to it a value of **01AE**.
  b. You use the variable in the following command:

     ```
     ANALYZE UNITADDRESS(ADDR) SCAN
     ```

  c. Automatic substitution of the **ADDR** variable would generate the following command. The generated command includes a modification to the **UNITADDRESS** parameter that might have effects that you do not intend:

     ```
     ANALYZE UNIT01AEESS(01AE) SCAN
     ```

DISP is another example of a variable name that is not recommended because that string is likely to appear on a JCL DD statement, where it should not be replaced.

- The product provides generic attribute names, such as ***DSN** and ***ROWNUMBER**.

To view the scope of substitution variables that you define in the **Options** tab, click the **Show Data** button.

How commands run when substitution variables are defined

You specify substitution variables and their mappings in the Substitution variables and their run-time replacement values area of the dialog box. Depending on your settings, the action request issues the command one or more times as follows:

Number of rows selected in the workspace	Are substitution variables defined in the Command tab?	How the command runs
One	Yes or no	One time for the selected item.
Multiple	Yes	Multiple times; one time for each selected item.
One or more groups	Yes. Group attribute is specified.	Once for each volume or dataset resource in each selected group.

Examples of substitution variables

The following example table shows typical substitution variables for a command that is issued by the Storage Toolkit. A similar, two-column table is provided in the dialog boxes of the toolkit.

Command Variable[1]	Attribute or String[2]
`%dsname%`	`ENTRYNAME`
`my-volser`	`VOLUME`
`%job-name%`	`"USER01A"`
`%sysout-class%`	`"X"`
`%rownum%`	`*ROWNUMBER`
`%user-account%`	`*USERID`

1. The heading for this column is "**Variable**" when you are defining substitution variables for a batch job in the Create Batch Job dialog box. See “About the "Command Variable" (or "Variable") column” on page 156 for more information.
2. You can enter raw text or use a drop-down menu to select the entries in the Attribute or string column. See “About the "Attribute or String" column” on page 157 for more information.

About the "Command Variable" (or "Variable") column

An example of the **Command Variable** (or **Variable**) column for substitution variables is shown in “Examples of substitution variables” on page 156. These examples reflect names of substitution variables that might be present in user-defined JCL. The examples illustrate the following guidelines and characteristics for entries in the **Command Variable** column:

- Strings can consist of up to 64 characters from the following character collection: [`a-zA-Z0-9_.%-`]. The strings are not case sensitive when processed by the mainframe agent.
- The ampersand character (&) is not supported for a substitution variable. Otherwise, there could be confusion with other use of ampersands within the JCL.
- Do not allow variable names to be a substring of another substitution variable. For example, do not use **my_dsn** and **my_dsname** as substitution variables.
- Do not create a variable that could match data in your command, JCL files, or generated JCL that should not be replaced. For example, **DISP** is not recommended as the name of a substitution variable, because that string is likely to appear on a JCL DD statement, where it should not be replaced.

About the "Attribute or String" column

An example of the **Attribute or String** column for substitution variables is shown in "Examples of substitution variables" on page 156. In most cases, the values for substitution variables are taken from columns of data in the workspace where the toolkit action request originates. You use the drop-down list in the **Attribute or String** column to select a value. This value is used as the source of values for the corresponding substitution variable.

The toolkit detects the entries that are available for the row that you select in a workspace prior to invoking a Storage Toolkit dialog box. "Examples of substitution variables" on page 156 illustrates the following guidelines and characteristics for entries in the Attribute or string column:

- **ENTRYNAME:** In this example, the user selected a table attribute, **ENTRYNAME**, from a list box, for the value of the variable.
- **USER01A:** In this example, the user typed a literal name (**USER01A**) to correspond to the %**job-name**% variable.
- **%job-name%, %sysout-class%:** You can enter literal text strings in this column, instead of attribute names. The raw text can reference items in the JCL that are not related to the underlying table, such as a %**job-name**%, %**sysout-class**%, and so on. The toolkit treats a value that you enclose in single or double quotes entered as a literal string rather than the name of a table column. A single occurrence of this string is used in the batch job or command.
- ***ROWNUMBER:** In this example, the user selected the ***ROWNUMBER** value from the drop-down menu that corresponds to the %**rownum**% variable. As a result of this setting a variable number of rows is generated for use in the JCL. The toolkit generates integer numbers ranging from 1 through the total number of data rows being passed from the underlying table. For example, you might select three rows from a workspace view when you create or resubmit an action request. In this case, the values 1, 2, and 3 would be generated.
- ***USERID:** In this example, the user selected the ***USERID** value from the drop-down menu that corresponds to the **%user-account%** variable. This value is filled in with a single occurrence of the name of the user account that is currently active in the Tivoli® Enterprise Portal.
- The following points apply to setting up batch jobs in the Create Batch Job dialog:
  - When the toolkit targets a member of a partitioned dataset or a sequential dataset that already exists, the toolkit scans that file for strings that meet the following naming convention: **%*name*%**, where *name* is the name of a substitution variable. The toolkit populates the dialog box with each of the corresponding variables for **%name%**.
  - The substitution of values for variables within the JCL might produce lines that are longer than 72 characters. If this occurs, this execution of the action request fails. A status of **Invalid JCL** is displayed in the Result Summary workspace.
  - The substitution of values for variables within other datasets that require variable substitution might produce lines that are longer than 80 characters. If this occurs, this execution of the action request fails. A status of **InvalidJCL** is displayed in the Results Summary workspace.
  - When configuration of the batch job is complete, click **OK** in the **Create Batch Job** dialog box. At this point, the JCL file that you provide and any files outside that JCL that contain substitution variables are copied into temporary files. These temporary files are used when the batch job JCL is submitted. The toolkit does not modify the original files, because other users might be referencing the files at the same time.

Standard tabs in dialog boxes of Storage Toolkit

Learn how to use the tabs that are standard in most dialog boxes.

The dialog boxes of Storage Toolkit have similar layouts.

To access dialog boxes, select and right-click one or more rows in the workspaces that have action menus. Then you choose one of the Storage Toolkit options in the pop-up menu to access the corresponding dialog box.

Options tab

Use this tab to choose configuration settings for the action that you want to run. The **Options** tab provides a subset of the available command parameters. The "Command reference guide" describes where to go for complete documentation of the available command parameters.

Command reference guide
The documentation for the z/OS® operating system includes reference information for the commands and parameters that the Storage Toolkit uses. You can access documentation for the version of z/OS® that you use through the following Web page: http://www-03.ibm.com/systems/z/os/zos/

The options that you select on the **Options** tab are rendered as parameters on the command. You can view the raw text of the command including your options on the **Command** tab.

> **Note:** If you make any modifications in the **Command** tab, further modifications in the **Options** tab are disabled. If you want to use the **Options** tab again, click **Cancel** to dismiss the dialog box, and begin constructing a new command in a new invocation of the dialog box.

Refer to "Setting substitution variables for toolkit commands" on page 154 for a scenario that describes how the **Options** tab and the **Command** tab interrelate.

The following exceptions exist regarding the **Options** tab of the Storage Toolkit:

- The **Options** tab for the Create Batch Job dialog box enables you to define substitution variables. You can click the **Show Data** button to access the Data View dialog box, where you can review the scope of the variables that you specify.
- The **Options** tab is not provided in the following dialog boxes, because they have no options to set: Dataset Space Release dialog box, Dataset Verify dialog box, Dataset Uncatalog dialog box, and Cancel HSM Request dialog box.
- The **Options** tab is not provided in the Issue Command dialog box, because you specify the command and its parameters as raw text in the **Command** tab of that dialog box.

General tab

Use the **General** tab to set name and description values for the action. Also use this tab to set the following options:

- **Run Options:**
  - **Run now, if no schedule:**
    - Select this option to cause the action request to run when you click **OK** in the dialog box. Keep in mind that any schedule that you specify in the **Schedule** tab takes precedence over the **Run now, if no schedule** option.
    - To save a Storage Toolkit action for later use, clear the **Run now, if no schedule** check box. The toolkit action request is saved, but it does not run. You can run the action later using the **Submit Request** or the **Submit Command or Job** context menu items. See "Reusing action requests" on page 147 for related information.
  - **Delete request and results after one run:** Select this option to cause deletion of the action request and results when the action request completes. If you do not need to review the results of a request and only intend to run the request one time, you might want to use this option. For example, you might want to delete the request and results for an MVS™ operator command that you need to run only one time. If you do not select this option, the request and results, persist until you delete them, as described in "Managing Actions" on page 192.
  - **Use request's default data:** When you reuse an existing action, select this option to resubmit it using the same resources that were used when the action was initially created. See "Reusing action requests" on page 147 for additional details.
  - **Run on situation firing:** This option is disabled and not checked for all non-situational requests. It is checked only for situational requests. It is included for information purposes only.

- **Results:**

  > **Note:** Use the dialog box options that are listed here to manage the results of requests. You must manage the buildup of old results and unused requests. Otherwise, this data can build up and actually cause the Toolkit Checkpoint Database to run out of space.

  - **Delete after (days):** This value determines the number of days after the toolkit action request completes that the results are automatically deleted. If the value is `0`, no results are saved. If you do not specify a value, the results persist until you delete them.
  - **Maximum output lines:** This value determines the number of lines that are saved from the execution of the action request. If the value is `0`, no output lines are saved. If this field is blank, all of the output lines are saved, up to a maximum defined by configuration options that are chosen for your production environment. If you specify a value other than `0`, the number of output lines that are saved may also be reduced based on the configuration options. See "Using the Storage Toolkit" on page 146 for additional information.

Scheduling tab

Use this tab to schedule the action in one of the following ways:

> **Note:** Time is in 12-hour format, so you must specify morning (**am**) or afternoon/night (**pm**).

- *At* a specific time and date OR
- *At* a specific time and days of the week. (Select the **Every** check box to enable the action to run every week on the schedule you specify.) OR
- *In* a specific time frame, when a specific amount of time has passed.

For details regarding the syntax of scheduling specifications, see the description of the **Schedule** attribute in the section "Storage Toolkit Requests" on page 363.

JCL tab

Use this tab to specify changes to JCL (job control language) characteristics as follows:

- **Replacement JCL JOB card:** Use this text box to define a new JCL JOB card specification to override the installation-specific JOB card. This job card can include substitution variables.
- **Extra JCL statements:** Use this text box to specify JCL statements to append to the JCL stream. For example, you can add DD statements to supplement a command in addition to control statements. This JCL can include substitution variables.
- **Dataset prefix for temp files:** Use this field to specify the prefix to use for the names of the temporary datasets that the Storage Toolkit creates to process the action request. The Storage Toolkit assigns unique names to the temporary datasets using this prefix as the high level qualifier. The Storage Toolkit accesses the temporary datasets under your security context. Therefore, your user ID must have authority to create, read, write, and delete datasets whose names begin with the prefix. If you do not specify a value, the default is your user ID.
- **Copy JCL and JES logs for later viewing:** Select this option to save a copy of the submitted JCL as well as the various JES files (**JESMSGLG**, **JESJCL**, and **JESYSMSG**) that are produced when the batch job runs. You can view these files in the **Action Results Details** view of the Storage Toolkit Result Detail workspace.

Command tab

Use this tab to review and modify the command that you construct using the Storage Toolkit.

For most dialog boxes of the toolkit, this tab displays a command that summarizes the selections that you make in the **Options** tab. However, there are two exceptions from the standard behavior of the **Command** tab:

- When you use the "Issue Command dialog box" on page 165 (by selecting the **Create Command** option in the pop-up menu of a product workspace) the **Command** tab is blank, ready for you to manually enter all the specifications for a new command.
- When you use the "Create Batch Job dialog box" on page 162, the **Command** tab is not present. The functionality to run user-defined JCL is provided in the **Options** tab of that dialog box.

1. The **Command type** drop-down list displays the command type that corresponds to the type of toolkit action that you are constructing:
   - DFSMSdss (Device Support Services)
   - DFSMShsm (Hierarchical Storage Manager)
   - ICKDSF (Device Support Facilities)
   - IDCAMS (Access Method Services)
   - DFSMSrmm™ (Removable Media Manager)
   - Mainframe (z/OS® console command)
   - TSO (Time Sharing Option)

   The type of command is selected automatically while you are using any of the dialog boxes of the toolkit. (The exception is the Issue Command dialog box, where the command type field can be configured. By default, **DFSMSdss (Device Support Services)** is selected, but you can select any of the command types in the drop-down list.)
2. Experienced users can modify the command and parameters that are displayed in the **Command** text box.

   > **Note:** If you make any modifications in the **Command** tab, further modifications in the **Options** tab are disabled. If you want to use the **Options** tab again, click **Cancel** to dismiss the dialog box, and begin constructing a new command in a new invocation of the dialog box.

3. Define substitution variables, as needed, in the **Substitution variables and their run-time replacement values** table. The following guidelines apply to the use of substitution variables:
   - IBM® Z OMEGAMON® AI for Storage® uses the percent sign (%) for the default substitution variables that it provides.
   - The product provides generic attribute names, such as *DSN and *ROWNUMBER.

   "Setting substitution variables for toolkit commands" on page 154 includes a description of how the **Options** tab and the **Command** tab interrelate.

Authentication and authorization for running Storage Toolkit commands

The user ID for the Tivoli® Enterprise Portal is fundamental to Storage Toolkit functionality. The following principles determine the success of Storage Toolkit operations:

- The user ID for the Tivoli® Enterprise Portal enables toolkit activity in the mainframe environment.
- Authorization is achieved when the portal user ID matches a valid user ID in the z/OS® system.
- You ensure the security of the z/OS® environment by enabling user authentication on the hub monitoring server.

Keep in mind the following details regarding the portal user ID:

- The following list provides examples of the types of authentication that you could use to secure the use of the Storage Toolkit:
  - For a z/OS® hub monitoring server, RACF®, ACF2, Top-Secret, or equivalent hub-server authentication.
  - For a distributed hub monitoring server, user IDs can be authenticated either by the local system registry or by an external LDAP-enabled central registry.
- The Storage Toolkit does not have access to lists of valid user IDs in the mainframe environment. The toolkit depends on you to use a portal ID that matches a valid ID in the z/OS® environment.

- For UNIX-based monitoring environments, you must ensure that the user ID exactly matches the user ID in the z/OS® environment, including case sensitivity. For example, if the user ID in the z/OS® system is `USER1` (and not `user1`), you must log in to UNIX-based Tivoli® Enterprise Portal with a `USER1` user account. If you log in to the Tivoli® Enterprise Portal with a `user1` user account, Storage Toolkit actions would fail because `user1` is not equivalent to `USER1` on the z/OS® system.
- You have the option to use various methods to resubmit actions that you defined previously. Authorization of a resubmission depends on the currently active user ID. (Resubmission does not depend on the ID used when the action was originally run.)
- To enable use of the temp files that are configured in the **JCL** tab, see the guidelines that are provided in the JCL tab section. The names of the target datasets start with the dataset prefix that is specified for temp files.

  Important: Keep in mind that the user ID must have the required authority to work with these datasets in the mainframe environment.

Note: You can map the Tivoli® Enterprise Portal user to any other user from which jobs are submitted. For more information, see https://www.ibm.com/docs/en/om-shared?topic=commands-setting-up-user-id-mapping-capability.

Data View dialog box

You use this dialog box to inspect the actual data that will substitute the substitution variables that you define in the **Command** or **Options** tabs.

You access this dialog box by clicking the **Show Data** button in the **Command** tab (or in the **Options** tab, when you are using the Create Batch Job dialog box). The **Command** tab is one of standard tabs that is described in "Standard tabs in dialog boxes of Storage Toolkit" on page 157. If you see inappropriate data in the Data View dialog box, respond as follows:

- If the substitution variables that you defined generate results that are not valid, click **Cancel** to return to the **Command** tab, and adjust your definitions.
- If resources are referenced in the dialog box that you did not intend to include, close the Data View dialog box, and click **Cancel** in the toolkit dialog box that you are configuring. Select the resources that you want to target, and access the toolkit dialog box again.

Note: If you select the **Use request's default data** option in the **General** tab, the substitution uses the data from the original request. The data that is displayed in the Data View dialog box will reflect this fact.

The general process of setting substitution variables and evaluating their scope is described in "Using substitution variables in the Storage Toolkit" on page 154. Column headings in the Data View dialog box display the ***ATTRIBUTE*(%*variable*%)** format, where the attribute name is followed by the corresponding substitution variable name in parenthesis, as in the following example headings:

- ***DSN (%SUB01%)**
- **Managed System (%name%)**

In these example headings, ***DSN** and **Managed System** are the attribute names and **%SUB01%** and **%name%** are the names of the substitution variables.

Custom actions in Storage Toolkit

In Storage Toolkit, you can create custom actions by using corresponding dialog boxes.

The **Custom actions and corresponding dialog boxes** table provides the names of the custom actions from the menu and the dialog boxes related to those custom actions.

Custom actions and corresponding dialog boxes	
Name of menu item	**Corresponding dialog box**
Create Batch Job	“Create Batch Job dialog box” on page 162
Create Command	“Issue Command dialog box” on page 165
Submit Command or Job	“Submit Command or Job dialog box” on page 166

Create Batch Job dialog box

You can use the Create Batch Job dialog box to submit user-defined JCLs.

This action request is a more flexible equivalent to the toolkit commands that the product provides.

“Using the Storage Toolkit” on page 146 provides helpful information about using this dialog box.

To learn about the standard capabilities of a dialog box, see Standard tabs in dialog boxes of Storage Toolkit. The topic includes a link to the reference guide for commands that are used in the Storage Toolkit.

See “Conventions for JCL files that you use with the Storage Toolkit” on page 164. Also see “Standard tabs in dialog boxes of Storage Toolkit” on page 157 regarding the authorization that is required for the action requests of the Storage Toolkit.

- **Create Batch Job dialog panes**: The right-hand pane is comprised of two tables. The first table includes the substitution variables and the value of their run-time replacements. The second table includes files whose contents you should copy to view later. The left-hand pane features a navigation tree view. Use the left-hand pane to view a collection of datasets. Each node on the tree represents a dataset. A context menu item may be disabled based on the kind or number of datasets that are selected. Right-click on a dataset node in the navigation tree, to display a context menu and the following list of menu options:
  - **Add**: To add a dataset to the tree view. You can add only one dataset at a time. Once you add a dataset, the dataset will be scanned for predefined substitution variables.
  - **Remove**: To remove a dataset from the tree view. You can remove multiple datasets at a time. If you right-click and remove the parent dataset node in the tree, then you will remove the parent dataset as well as all the children datasets nodes that are associated with the parent node.

    > **Note:** All substitution variables included in the substitution variable table that are associated with and are unique to the dataset being removed will be deleted from the table. For example, you may select any number of datasets in the tree, then right-click and remove them.

  - **Rename**: To rename the dataset in the tree view *not* on the z/OS managed system.
  - **Edit**: To edit a dataset in the tree view. You can edit only one dataset at a time.
  - **Refresh**: To rescan the contents of the dataset in the tree view. You can refresh multiple datasets at a time. Revalidates to make sure the substitution variables that are located in the substitution variable table still exist in the dataset on the managed system.
  - **Submit**: To submit the batch JCL. This menu item is enabled only if the selected dataset is a Master dataset.
- **Mainframe fully qualified dataset containing the batch JCL**: Specifies the dataset that contains the batch job. The dataset must be a sequential dataset or a member of a partitioned dataset (PDS). The logical record length of the dataset must be 80 and the dataset must be a fixed length format. To edit or create the dataset, click the **Edit JCL** button, which causes the “Edit Dataset dialog box” on page 165Edit Dataset dialog box to be displayed.

  > **Note:** The user ID for the Tivoli Enterprise Portal that you use must be authorized in the mainframe system to access the dataset. When the user ID is not authorized, you see an error message when you click **Edit JCL**. For more information, see “Standard tabs in dialog boxes of Storage Toolkit” on page 157.

- **Substitution variables and their run-time replacement values**: Specifies replacement values to be used at run time. The replacement values allow the batch job to reference specific objects, such as datasets and volumes. For example, these objects might correspond to values in the table rows that you select prior to invoking this dialog box.
  In the **Variable** column, you can enter a unique name for each substitution variable that exists in the dataset. In the cells of the **Attribute or string** column, you can select from the available names. These names reflect attributes from the table rows that you selected in a workspace, prior to invoking the Create Batch Job dialog box. Alternatively, you can manually enter values in the cells. For example, you might want to enter values such as job name, **sysout** class, and other values that might not be present in the workspace table from which you are selecting rows.

  > **Tip:** Use the percent-sign convention for the substitution variables (*%name%*) that are present in your JCL dataset. That way, this dialog box automatically finds and lists the JCL substitution variables.

  Refer to "Using substitution variables in the Storage Toolkit" on page 154 for guidelines that apply to the use of substitution variables.

- **Files whose contents should be copied for later viewing**: Specifies the DD names and datasets to be preserved as a record of the results of the batch job. The files are available for viewing in the Storage Toolkit Result Details workspace.
  In the **Step name** column, you specify the name of the step whose data you want to save. In the **Fully Qualified Dataset or DD name** column, you specify the sequential datasets, members of partitioned datasets, or DD names that you want to copy. The following table shows some sample settings for this area of the dialog box:

Sample settings for "Files whose contents should be copied for later viewing"	
Step name	**Dataset or DD name**
STEPX	USER02.LOG.LISTING
STEPY	SYSPRINT

  The specification in row one preserves data from the **USER02.LOG.LISTING** dataset.

  The specification in row two preserves the **SYSPRINT** file in **STEPY**. This output is saved in the checkpoint file. The following restrictions apply:

  - Only sequential datasets or members of a PDS are supported.
  - VSAM datasets are not supported.
  - Datasets or files that contain unprintable characters (for example hexadecimal code) might display improperly in the Storage Toolkit Result Details workspace.
  - Output lines longer than 256 characters are truncated.

- **Fully Qualified Datasets needed by the job that also contain substitution variables:** Specifies datasets that are located outside the JCL job stream. Such datasets might be needed to dynamically supply substitution variables that correspond to objects (such as datasets and volumes) that are needed when the batch job runs. These additional datasets can be sequential datasets or members of partitioned datasets. The logical record length of these datasets must be 80 and these datasets must be a fixed length format. For example, an IEHLIST job might use a SYSIN DD statement that points to a DASD dataset rather than in-stream control cards. That dataset could contain substitution variables that need replacement at execution time.

  If no datasets that contain substitution variables are referenced from the JCL, this area of the dialog box can remain empty.

  You might encounter the following return codes when the attempt to access datasets fails:

Return code	Description
119 (DataSetError)	The User Dataset server (UDS) has failed.
121 (AuthorizationFailure)	You are attempting to access a dataset for which you do not have RACF authority.

Return code	Description
123 (DataSetError)	The dataset does not exist or the dataset is not a supported format or record length.

You might also receive a return code (status) of 117 (**InvalidJCL**). This status typically indicates that the JCL exceeds 72 characters when the substitution variables are applied. It might also indicate other JCL-related errors, such as a missing JOB card, or that another dataset requiring variable substitution exceeds 80 characters when the substitution variables are applied.

Be sure to consult the RKLVLOG for additional messages regarding these results.

When configuration of the batch job is complete, you click **OK** in the Create Batch Job dialog box. At this point, the JCL file that you provide and any files outside that JCL that contain substitution variables are copied into temporary files. These temporary files are used when the batch job JCL is submitted.

Conventions for JCL files that you use with the Storage Toolkit

The dataset for the JCL that you create or edit must be a sequential dataset or a member of a partitioned dataset (PDS). Any JCL dataset that you use with the Storage Toolkit must follow conventions to enable the Tivoli® Enterprise Portal to substitute data values at run time. The following elements (marked in bold text in "Figure: Variable elements for JCL files used with the Storage Toolkit" on page 164) are substitution variables that will be altered by the Storage Toolkit at job run time:

- `%job-name%` (see line 000100)
- `%sysout%` (see line 000100)
- `%my-volser%` (see line 000800)
- `%dsname%` (see line 000800)

Figure: Variable elements for JCL files used with the Storage Toolkit

```
****** ***************************** Top of Data ******************************
000100 //%job-name% JOB (),'DHODS',REGION=4M,MSGCLASS=%sysout%,NOTIFY=&SYSUID,
000200 //         CLASS=A
000300 //* ---------------------------------------------------------------------
000400 //LISTVTOC EXEC PGM=IEHLIST
000500 //SYSPRINT DD SYSOUT=*
000600 //SYSIN    DD *
000700 )DOT
000800   LISTVTOC DUMP,VOL=3390=%my-volser%,DSNAME=(%dsname%)
000900 )ENDDOT
001000 /*
001100 //
****** **************************** Bottom of Data ****************************
```

Tip: Use the percent-sign convention for the substitution variables (*%name%*) that are present in your JCL dataset. That way, this dialog box automatically finds and lists the JCL substitution variables.

The **)DOT** and **)ENDDOT** lines in "Figure: Variable elements for JCL files used with the Storage Toolkit" on page 164 indicate the start and end of table processing. The JCL lines between the **)DOT** and **)ENDDOT** lines are repeated for each table row that is passed into the mainframe agent at JCL execution time. Typically, these lines contain substitution variables that are replaced with new values at run time. For example, the output for these SYSIN JCL lines might look similar to this excerpt:

```
000600 //SYSIN    DD *
000700   LISTVTOC DUMP,VOL=3390=TSO201,DSNAME=(SYS2.USER.CLIST)
000800   LISTVTOC DUMP,VOL=3390=TSO202,DSNAME=(SYS2.USER.PROCLIB)
000900   LISTVTOC DUMP,VOL=3390=MVSA12,DSNAME=(SYS2.PROD.LOADLIB)
001000 /*
```

Note: Use of the)DOT and)ENDOT to aggregate results from multiple workspace rows into a single batch job is supported for direct toolkit actions, such as performing a multi-row select in a workspace view, and launching a toolkit action from the right-click menu. You must always use the)DOT and)ENDOT in situation-driven toolkit actions.

Edit Dataset dialog box

You can use the Edit Dataset dialog box to create, edit, and save JCL, which is displayed in a scrollable ASCII text box in the dialog box.

The user ID that you use in the Tivoli® Enterprise Portal must be a valid mainframe user ID that provides access to the mainframe JCL dataset.

Click the **Edit JCL** button in the Create Batch Job dialog box to access the Edit dataset dialog box. The name of the dataset to edit appears in the title of the dialog box.

If the dataset that you specify does not exist, the dataset is created when you click **OK**. This dialog box can contain ASCII text only.

Line and column indicators displayed below the ASCII text box identify the current position of the cursor. A ruler is also available to assist with column identification. By default, the ruler is hidden. When displayed, it appears above the ASCII text box.

Use **Ctrl**+**R** to toggle the display of the ruler.

The Edit dataset dialog box also displays your editing mode. The editing mode appears below the ASCII text box:

- REP indicates that you are typing in Replace mode (overwriting existing characters).
- INS indicates that you are typing in Insert mode.

Press the **Insert** key to toggle between insert and replace modes. Replace mode is the default.

A vertical gray bar in the ASCII text box is positioned at the logical record length of the dataset. If the dataset does not exist, a length of 80 is used. You cannot type beyond this bar, nor can characters shift past the bar when you are in insert mode. When the line on which you are typing reaches the bar, additional characters that you type are ignored.

Use the **Enter** key to create new lines.

Issue Command dialog box

You use this dialog box to issue arbitrary commands, including IDCAMS, ICKDSF, DFSMSrmm™, DFSMShsm™, DFSMSdss™, TSO, and mainframe console commands in text format.

Note: An MVS command that you issue through the Issue Command dialog box is submitted through SDSF. A forward slash (/) must precede MVS commands, as in this example, which cancels a time-sharing user (*tso_user_ID*):

```
/C U=tso_user_ID
```

Also see “Standard tabs in dialog boxes of Storage Toolkit” on page 157 regarding the authorization that is required for the action requests of the Storage Toolkit.

Typically you use this feature when you are familiar with the syntax and parameters of a command. You construct and issue a mainframe command directly from the Tivoli® Enterprise Portal, without having to directly access the mainframe environment. See “Standard tabs in dialog boxes of Storage Toolkit” on page 157 for a description of the relevant requirements for your portal user ID.

You supply the command details in the following fields of the dialog box:

- Command Type drop-down list

- Command text box
- Substitution variables and their run-time replacement values table

For additional information, see the description of the generic **Command** tab.

In the generic **Command** tab you enter *supplemental* parameters to a basic command that you specify in a graphical user interface. In contrast, the **Command** tab of the Issue Command dialog box is used to enter a complete command in text form. Because you specify all options for the command in the text box, the Issue Command dialog box does not have an **Options** tab.

Command tab

Use this tab to construct the raw text of the command that you want to issue.

1. Select one of the following types of command in the **Command type** drop-down list:
   - DFSMSdss (Device Support Services)
   - DFSMShsm (Hierarchical Storage Manager)
   - ICKDSF (Device Support Facilities)
   - IDCAMS (Access Method Services)
   - DFSMSrmm™ (Removable Media Manager)
   - Mainframe (z/OS® console command)
   - TSO (Time Sharing Option)

   You must select the type of command (corresponding to the type of facility in the mainframe environment) that matches the command type that you enter in the Command text box. Otherwise, the command is not processed successfully in the mainframe environment.
2. Type the command and its parameters in the **Command** text box. Use substitution variables as needed.
3. Define substitution variables, as needed, in the **Substitution variables and their run-time replacement values** table.
4. Click the **Show Data** button to access the "Data View dialog box" on page 161, where you can review the scope of the substitution variables that you specify.

To learn about the standard capabilities of a dialog box, see Standard tabs in dialog boxes of Storage Toolkit. The topic includes a link to the reference guide for commands that are used in the Storage Toolkit.

Submit Command or Job dialog box

You use this dialog box to reuse a command that you previously issued through the Storage Toolkit.

The relevant types of commands include the following:

- A standard command like the **Migrate** command for datasets that invokes the "Dataset Migration dialog box" on page 171.
- A batch job generated in the "Create Batch Job dialog box" on page 162.
- A command generated in the "Issue Command dialog box" on page 165.

You select data rows in a workspace, such as dataset names or volume names, and right-click and select **Submit Command or Job** to invoke the dialog box. The Submit Command or Job dialog box is displayed and includes the following features:

- A table of data regarding commands that have previously run, including the ID of the creator, name and type of command, and the syntax of the command.
- The **Show only actions applicable to the selected data row(s)** check box. When you select this option, the table updates to display only the toolkit actions that are relevant to the selected data rows or commands that do not require data from the underlying tables. (These rows are the ones that you select prior to right-clicking and selecting **Submit Command or Job** in the pop-up menu.)

When you select a toolkit action and click **OK** in the Submit Command or Job dialog box, the relevant toolkit dialog box is displayed. For example, if you click on an action that migrates datasets and click **OK**, the "Dataset Migration dialog box" on page 171 is displayed. At the same time, the relevant values of the data rows that you

selected in the portal workspace are captured for use in the dialog box. When you submit the action, all the values are combined into a single batch file. The current Tivoli® Enterprise Portal user ID is used for this transaction. For more information, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157.

See Reusing action requests to learn how to resubmit a toolkit action through a Storage Toolkit workspace.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Volume actions in Storage Toolkit

Learn about the dialog boxes that are related to volume actions in Storage Toolkit.

You access the following dialog boxes in the **Volume Actions** submenu. This submenu is available in the pop-up menu that you access by right-clicking specific views in a workspace.

Name of menu item	Corresponding dialog box
Backup	"Volume Backup dialog box" on page 167
Convert	"Volume Conversion dialog box" on page 167
Migrate	"Volume Migration dialog box" on page 168
Release Space	"Volume Space Release dialog box" on page 168
Compress	"Compress dialog box" on page 168
Defrag	"Defrag dialog box" on page 169

You can also create a custom action. For more information, see "Custom actions in Storage Toolkit" on page 161.

Volume Backup dialog box

You use this dialog box to issue the DFSMShsm BACKVOL command to back up all eligible datasets on the selected volumes.

You access this dialog box in the **Volume Actions** submenu. This submenu is available in the pop-up menu that you access by right-clicking specific views in a workspace.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

The action supplies the following information to the service:

- Device Type (default device category, Type or DASD)
- Backup Type (incremental or full)
- Backup datasets not backed up in the last *n* days. (Range of *n*: 0-999)

For datasets not managed by SMS, you can target datasets that have not been backed up in the last *n* days, where *n* ranges from 0 to 999.

With Tivoli OMEGAMON for Storage on z/OS Version 4.2.0 Interim Feature 3 or later, you can select *SMSVOLG or *USRVOLG as an attribute of the command variable on the Command tab of the Volume Backup dialog box. *SMSVOLG represents a dynamic set of volumes within the storage group which will be resolved at job submission time. *USRVOLG represents a dynamic set of volumes within the USER DASD storage group which will be resolved at job submission time.

Volume Conversion dialog box

You use this dialog box to issue the DFSMSdss CONVERTV command to convert the selected volumes to or from system-managed storage (SMS).

You access this dialog box in the **Volume Actions** submenu. This submenu is available in the pop-up menu that you access by right-clicking specific views in a workspace.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

The following parameters are available:

- **Action:** Prepare, Test, or Convert.
- Convert multi-volume datasets.
- Redetermine classes for SMS datasets.

Volume Migration dialog box

You use this dialog box to issue the DFSMShsm MIGRATE VOLUME command to migrate all eligible datasets on the selected volumes.

You access this dialog box in the **Volume Actions** submenu. This submenu is available in the pop-up menu that you access by right-clicking specific views in a workspace.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

For datasets not managed by SMS, you can target datasets that have not been referenced in the last *n* days, where *n* ranges from 0 to 999.

Volume Space Release dialog box

You use this dialog box to issue the DFSMSdss RELEASE command to release allocated but unused space from all datasets on the selected volumes.

You access this dialog box in the **Volume Actions** submenu. This submenu is available in the pop-up menu that you access by right-clicking specific views in a workspace.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

The following parameters are available:

- Minimum secondary allocation for release (*n* number of tracks, where *n* ranges from 0 to 99999999)
- Minimum unused space prior to release (*n* number of tracks, where *n* ranges from 0 to 99999999)
- Wait time in seconds between retry attempts (*n* seconds, where *n* ranges from 0 to 255)
- Number of attempts allowed to obtain volume access (*n* tries, where *n* ranges from 0 to 99)
- A fully qualified dataset name or pattern for datasets that you want to exclude from the space release operation. Use the **Fully Qualified Dataset Name or Pattern** table to manually enter dataset names or DSN patterns.

Compress dialog box

You use this dialog box to issue the DFSMSdss™ COMPRESS command to compress all eligible datasets on the volumes that you selected in the Tivoli® Enterprise Portal.

You access this dialog box in the **Volume Actions** submenu. This submenu is available in the pop-up menu that you access by right-clicking specific views in a workspace.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

The dialog box provides the following options:

- **Exclude datasets qualified by:** You can add a qualifier to exclude specific datasets on the volume. Enter a fully-qualified name or pattern that identifies the datasets to exclude.

- **Include datasets qualified by:** You can add a qualifier to include specific datasets on the volume. Enter a fully-qualified name or pattern that identifies the datasets to include.
- **Run as DFSMSdss™ Administrator:** Select this option to cause the command to run with Administrator authorization in DFSMSdss™. The user ID that you are using must be authorized as a storage administrator in DFSMSdss™.
- **Dynamic Allocation:** Select this option to cause the command to use dynamic allocation for the affected datasets.

Defrag dialog box

You use this dialog box to issue the DFSMSdss™ DEFRAG command to defragment the volumes that you selected in the Tivoli® Enterprise Portal.

You access this dialog box in the **Volume Actions** submenu. This submenu is available in the pop-up menu that you access by right-clicking specific views in a workspace.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

The dialog box provides the following options:

- **Exclude Datasets qualified by:** You can add a qualifier to exclude datasets. Enter a fully-qualified name or pattern that identifies the datasets to exclude.
- **Fragmentation Index:** Specify an integer *n*. When fragmentation drops below the value of *n*, the defragmentation operation ends. The DFSMSdss facility places a decimal point before the value that you specify for *n*, as in these examples:
  - If you specify `1`, the value becomes `.1`.
  - If you specify `999`, the value becomes `.999`.
  - If you specify `001`, the value becomes `.001`.
- **Fast Replication:** Use the drop-down menu to set your preference for use of the DFSMSdss™ fast replication feature. If the **Fast Replication** option is blank (the default), fast replication is not enabled.
- **Run as DFSMSdss™ Administrator:** Select this option to cause the command to run with Administrator authorization in DFSMSdss™. The user ID that you are using must be authorized as a storage administrator in DFSMSdss™.
- **Dynamic Allocation:** Select this option to cause the command to use dynamic allocation for the affected datasets.
- **Peer-To-Peer Remote Copy for FlashCopy®:** Select this option to specify that a Peer-to-Peer Remote Copy (PPRC) primary volume is allowed to become a FlashCopy® target volume. This option is ignored when FlashCopy® is not used for defragmentation. RACF authorization might be required for this option.
- **Consolidate:** Select this option to enable the command to consolidate multiple extent datasets when possible.
- **Write Check:** Select this option to enable verification for the data moved by defragmentation. Verification increases the time required for the defragmentation operation.

Dataset actions in Storage Toolkit

You can find information about dialog boxes that are related to dataset actions for Storage Toolkit.

You access the following dialog boxes in the **Dataset Actions** submenu. This submenu is available in the pop-up menu that you access by right-clicking specific views in a workspace.

Name of menu item	Corresponding dialog box
Backup	"Dataset Backup dialog box" on page 170
Attribute Group Extract	"Dataset Attribute Group Extract dialog box" on page 170
Migrate	"Dataset Migration dialog box" on page 171

Name of menu item	Corresponding dialog box
Move & Copy	“Dataset Copy and Move dialog box” on page 171
Recall	“Dataset Recall dialog box” on page 172
Recover	“Dataset Recover dialog box” on page 172
Release Space	“Dataset Space Release dialog box” on page 172
Compress	“Dataset Compress dialog box” on page 173
Print	“Print dialog box” on page 184
Allocate	“Dataset Allocate dialog box” on page 173
Rename	“Dataset Rename dialog box” on page 182
Alter	“Dataset Alter dialog box” on page 174
Delete	“Dataset Delete dialog box” on page 179
Listcat	“Dataset Listcat dialog box” on page 180
Repro	“Dataset Repro dialog box” on page 182
Verify	“Dataset Verify dialog box” on page 184
Catalog	“Dataset Catalog dialog box” on page 179
Uncatalog	“Dataset Uncatalog dialog box” on page 184
Define Cluster	“Dataset Define Cluster dialog box” on page 175

You can also create a custom action. See “Custom actions in Storage Toolkit” on page 161.

Dataset Backup dialog box

You use this dialog box to issue the DFSMShsm BACKDS command to create a backup version of specific datasets.

The action that you define targets the datasets that you select.

You access this dialog box in the **Dataset Actions** submenu. This submenu is available in the pop-up menu that you access by right-clicking specific views in a workspace.

To learn about the standard capabilities of a dialog box, see “Standard tabs in dialog boxes of Storage Toolkit” on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

You can also create a custom action. See “Create Batch Job dialog box” on page 162, “Issue Command dialog box” on page 165, and “Submit Command or Job dialog box” on page 166.

This dialog box works with the following information:

- **Target for the backup operation:** Default, Tape, or DASD.
- **Concurrent copying:** Standard, Preferred, or Required.

You can also select *DSNG or *DSNAG as an attribute of the command variable on the Command tab of the Dataset Backup dialog box. *DSNG represents a dynamic set of datasets within the dataset group which will be resolved at job submission time. *DSNAG represents a dynamic set of datasets within the dataset attribute group which will be resolved at job submission time.

Dataset Attribute Group Extract dialog box

Use this dialog box to issue the **`Dataset Attribute Group Extract`** command to extract attribute data.

You access the Dataset Attribute Group Extract dialog box in the Tivoli® Enterprise Portal (TEP) interface and in the enhanced 3270UI differently:

- In TEP, you access this dialog box in the **Dataset Actions** submenu. This submenu is available in the pop-up menu that you open by right-clicking specific views in a workspace.
- In the enhanced 3270UI, you access this dialog box in the **Storage Toolkit Actions** submenu.

The action that you define targets the dataset attribute groups that you select. This panel prepares the output dataset that will contain the extracted data. The dataset will be allocated if it does not exist.

The **Space Units** section initially displays the calculated block size for the output dataset. When you change the **Units** field to either tracks or cylinders both in TEP and the enhanced 3270UI, the **Value** field will show N/A. But in the enhanced 3270UI, this value is not displayed on the **Options** tab.

The application saves the values that you set and uses them when you resubmit the action (either manually via the Storage Toolkit workspace, or via a schedule or situation).

> **Note:** As the contents of an attribute group change over time, the primary and secondary size allocations used in the original action may no longer support extracting the data.
> Consider this when creating the original action. Also consider changing the secondary allocation to accommodate the maximum number of records possible for the attribute group or groups.

To learn about the standard capabilities of a dialog box, see Standard tabs in dialog boxes of Storage Toolkit. The topic includes a link to the reference guide for commands that are used in the Storage Toolkit.

You can also create a custom action. See “Create Batch Job dialog box” on page 162, “Issue Command dialog box” on page 165, and “Submit Command or Job dialog box” on page 166.

Dataset Migration dialog box

You use this dialog box to issue the DFSMShsm MIGRATE DATASETNAME command to migrate the selected datasets.

You access this dialog box in the **Dataset Actions** submenu. This submenu is available in the pop-up menu that you access by right-clicking specific views in a workspace.

To learn about the standard capabilities of a dialog box, see “Standard tabs in dialog boxes of Storage Toolkit” on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

You can also create a custom action. See “Create Batch Job dialog box” on page 162, “Issue Command dialog box” on page 165, and “Submit Command or Job dialog box” on page 166.

Convert (recall immediately to a different volume) is allowed.

To recall datasets to a different volume immediately, select the **Convert** checkbox.

Dataset Copy and Move dialog box

You use this dialog box to issue the DFSMSdss COPY command to copy or move the selected datasets to new locations.

You access this dialog box in the **Dataset Actions** submenu. This submenu is available in the pop-up menu that you access by right-clicking specific views in a workspace.

To learn about the standard capabilities of a dialog box, see “Standard tabs in dialog boxes of Storage Toolkit” on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

You can also create a custom action. See “Create Batch Job dialog box” on page 162, “Issue Command dialog box” on page 165, and “Submit Command or Job dialog box” on page 166.

The dialog box sets the following COPY command parameters for an action that starts a DFDSS service:

- For source dataset the dialog box can define the following behavior:
  - Process multi-volume datasets
  - Process VSAM sphere components
  - Force unmovable dataset
  - Disposition of the source dataset: keep, delete, or uncatalog (Uncat)
- For the destination of the move/copy operations the dialog box can set the following parameters:

- Volser name
- Fully Qualified Dataset name
- Name of the storage class
- Name of the management class
- Other options:
  - Bypass the SMS ACS routine, Catalog (default), and Reblock
  - Conversion to PDS, PDSE, or no conversion

Dataset Recall dialog box

You use this dialog box to issue the DFSMShsm RECALL command to recall the selected migrated datasets.

You access this dialog box in the **Dataset Actions** submenu. This submenu is available in the pop-up menu that you access by right-clicking specific views in a workspace.

The dialog box defines the DA option for a DFSMShsm™ service to the default, RELBLK, RELTRK or SAMETRK. You also have the option to specify the target volume.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

You can also create a custom action. See "Create Batch Job dialog box" on page 162, "Issue Command dialog box" on page 165, and "Submit Command or Job dialog box" on page 166.

Dataset Recover dialog box

You use this dialog box to issue the DFSMShsm RECOVER command to restore the selected datasets from backed up versions of those datasets.

You access this dialog box in the **Dataset Actions** submenu. This submenu is available in the pop-up menu that you access by right-clicking specific views in a workspace.

The action targets a single dataset to be recovered from a backup operation:

- A new, fully qualified name for the recovered dataset
- The volume to which the recovered dataset is sent
- The option to replace the dataset
- The backup version to be used in restoring the dataset

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

You can also create a custom action. See "Create Batch Job dialog box" on page 162, "Issue Command dialog box" on page 165, and "Submit Command or Job dialog box" on page 166.

Dataset Space Release dialog box

You use this dialog box to issue the DFSMSdss RELEASE command to release allocated but unused space from the selected datasets.

You access this dialog box in the **Dataset Actions** submenu. This submenu is available in the pop-up menu that you access by right-clicking specific views in a workspace.

This dialog box does not have options that you can set.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

You can also create a custom action. See “Create Batch Job dialog box” on page 162, “Issue Command dialog box” on page 165, and “Submit Command or Job dialog box” on page 166.

Dataset Compress dialog box

You use this dialog box to issue the DFSMSdss™ COMPRESS command to compress the datasets that you selected in the Tivoli® Enterprise Portal.

You access this dialog box in the **Dataset Actions** submenu. This submenu is available in the pop-up menu that you access by right-clicking specific views in a workspace.

The dialog box provides the following options:

- **Run as DFSMSdss™ Administrator:** Select this option to cause the command to run with Administrator authorization in DFSMSdss™. The user ID that you are using must be authorized as a storage administrator in DFSMSdss™.
- **Dynamic Allocation:** Select this option to cause the command to use dynamic allocation for the affected datasets.

You can also select *VOL@*DSNG or *VOL@*DSNAG as an attribute of the command variable on the Command tab of the Dataset Compress dialog box. *VOL@*DSNG represents the volume attribute of each dataset in the dataset group and is resolved at job submission. *VOL@*DSNAG represents the volume attribute of each dataset in the dataset attribute group and is resolved at job submission. Note that *VOL@*DSNG and*VOL@*DSNAG attributes should be used in conjunction with *DSNG and *DSNAG, respectively. When you click the **Show Data** button, the dialog box displays the group name (for the extension subject) and the follower subject name.

To learn about the standard capabilities of a dialog box, see “Standard tabs in dialog boxes of Storage Toolkit” on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

You can also create a custom action. See “Create Batch Job dialog box” on page 162, “Issue Command dialog box” on page 165, and “Submit Command or Job dialog box” on page 166.

Dataset Allocate dialog box

You use this dialog box to issue an IDCAMS ALLOCATE request to create a dataset. You can optionally model it based on the selected dataset.

You access this dialog box in the **Dataset Actions** submenu. This submenu is available in the pop-up menu that you access by right-clicking specific views in a workspace.

To learn about the standard capabilities of a dialog box, see “Standard tabs in dialog boxes of Storage Toolkit” on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

You can also create a custom action. See “Create Batch Job dialog box” on page 162, “Issue Command dialog box” on page 165, and “Submit Command or Job dialog box” on page 166.

IDCAMS ALLOCATE options that are available in the Allocate dialog box of the Storage Toolkit	
Option	**Description**
Fully Qualified Dataset name	Type the fully-qualified name of the dataset that you want to allocate.
Model name	Type the name of a model dataset that you want to use.
Use model to allocate new dataset	Select this option when you want to use the value of Model name in the allocation of the new dataset.
Disposition	You have the option to select a disposition value in this drop-down menu.
Record format	You have the option to select a value for the format of the record in this drop-down menu. Options include fixed-format records, variable-format records, and undefined-format records.

Option	Description
Dataset organization	You have the option to select the dataset organization type in this drop-down menu. Options include **DA** (direct access); **DAU** (direct access unmovable); **PO** (partitioned organization); **POU** (partitioned organization unmovable); **PS** (physical sequential); and **PSU** (physical sequential unmovable).
Dataset type	You have the option to select a dataset type in this drop-down menu. Options include partitioned dataset (PDS) and the Hierarchical File System (HFS) for z/OS® UNIX®.
Data class	Type the name of the data class for either SMS or non-SMS-managed datasets.
Record length	Specify the length, in bytes, of the largest logical record in the dataset.
Management class	Type the name of the management class for a new dataset. When possible, do not specify this value. Allow the value to default through the ACS routines.
Block size	Specify the block size of the data control block (DCB) for the dataset.
Storage class	*For SMS-managed datasets* you can type the storage class name. When possible, allow this value to default through the ACS routines established by the system administrator.
Volser(s)	Choose one of the following values in the list: • The serial number of an eligible direct access volume on which a new dataset is to be located. OR • The serial number of a direct access volume on which an old dataset is located. If you use VOLUME for an old dataset, the dataset must be on the specified volume for allocation to take place. If you do not include VOLUME, the system allocates new datasets to any eligible direct access volume.
Unit	Enter the name of the unit type to which a file or dataset is to be allocated. You can list an installation-defined group name, a generic device type, or a specific device address.
Directory blocks	Specify the number of 256-byte records for the directory of a new partitioned dataset. Use this parameter to allocate a new partitioned dataset.
Space Units	Specify the units and amount of space to allocate to a new dataset. Options include the following combinations units and values: • BLOCK: An integer value to be multiplied by the **Primary** value that you specify in the Space area of this dialog box. • AVBLOCK: An integer value to be multiplied by the **Primary** value that you specify in the Space area of this dialog box. • CYLINDERS: The **Primary** value that you specify in the Space area of this dialog box is given in cylinders. (No **Value** is required for this option.) • TRACKS: The **Primary** value that you specify in the Space area of this dialog box is given in tracks. (No **Value** is required for this option.)
Release space	Select this option when you want to delete unused space when the dataset is closed.
Space	Allocates the amount of space for a new dataset. The **Primary** value specifies the initial number of units of space that is allocated for a dataset. For a partitioned dataset, a directory quantity is not necessary. The **Secondary** value specifies the number of units of space to be added to the dataset each time the previously allocated space has been filled. You must provide the primary quantity along with the secondary value.

Dataset Alter dialog box

You use this dialog box to issue an IDCAMS ALTER request to alter the datasets and catalogs that you select in the Tivoli® Enterprise Portal.

You access this dialog box in the **Dataset Actions** submenu. This submenu is available in the pop-up menu that you access by right-clicking specific views in a workspace.

The dialog box provides the following options:

- **Fully Qualified Name:** Specifies a new, fully-qualified name for the dataset that you have selected. If you have selected multiple datasets, this field is disabled. Other fields in the dialog box continue to be available.
- **Catalog:** Specifies a new name for the catalog that contains the entry.
- **Management class:** *For SMS-managed datasets only*, this option specifies a new name for the management class for a dataset. The name must come from a list of names that the system administrator defines.
- **Storage class:** *For SMS-managed datasets only*, this option specifies a new name for the storage class. The name must come from a list of names that the system administrator defines.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

You can also create a custom action. See "Create Batch Job dialog box" on page 162, "Issue Command dialog box" on page 165, and "Submit Command or Job dialog box" on page 166.

Dataset Define Cluster dialog box

You use this dialog to issue an IDCAMS DEFINE CLUSTER command to define attributes for a cluster as a whole and for the components of a cluster.

You access this dialog box in the **Dataset Actions** submenu. This submenu is available in the pop-up menu that you access by right-clicking specific views in a workspace.

The **Options** tab of this dialog box provides the following options in three tabbed subsections: the **Cluster** tab, the **Data** tab, and the **Index** tab. These tabs correspond to the cluster that you define and its data and index components. You specify many of the same attributes for each of the three entities (cluster, and its data and index components). The values that you specify are unique to each entity.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

You can also create a custom action. See "Create Batch Job dialog box" on page 162, "Issue Command dialog box" on page 165, and "Submit Command or Job dialog box" on page 166.

IDCAMS DEFINE CLUSTER options that are available in the Dataset Define Cluster dialog box of the Storage Toolkit. Part 1 of 3, Cluster options

Option	Description
Fully Qualified Name	Type the fully-qualified entryname for the cluster. This name is not propagated to the DATA and INDEX components of the cluster. For SMS and non-SMS-managed clusters, the component names must resolve to the same catalog as the cluster name of the dataset. You can define a separate entryname for the cluster, its data component, and its index component. If no name is specified for the data and index component, a name is generated. When the cluster, data component, and index component are individually named, each can be addressed.
Fully Qualified Model	Type the fully-qualified name of the cluster that you want to use as a model.
Space Units	Select the unit of space that you want to allocate in the volume. Options include cylinders, KB, MB, records, or tracks. Then specify a corresponding value for the unit that you selected. The **Primary** value that you specify allocates the initial amount of space to the cluster. The **Secondary** value that you specify allocates an amount of space for each time the cluster extends.

Option	Description
Volser(s)	Specifies the volumes on which a cluster is to have space. The maximum number of VOLSERs is 59. Choose one of the following values in the list: • The serial number of an eligible direct access volume on which a new cluster is to be located OR • The serial number of a direct access volume on which an old cluster is located. If you choose an old cluster for **Volser**, the cluster must be on the specified volume for allocation to take place. If you do specify a value for **Volser**, the system allocates new clusters to any eligible direct access volume.
Control interval size	The size of the control interval for the cluster.
Buffer space	Specifies the minimum space for buffers. If you do not specify a value, VSAM attempts to get enough space to contain two data component control intervals and, if the data is key-sequenced, one index component control interval.
Data class	Identifies the name (1-to-8 characters) of the data class for the dataset. This metric provides the allocation attributes for new datasets.
Account	Specify up to 32-bytes of accounting information and user data for the dataset. The information is supported for only SMS-managed VSAM and non-VSAM datasets. The information is used only for the dataset level (not member level) of PDSE/PDS.
Management class	*For SMS-managed datasets*, you can type the name of the management class for a new dataset.
Owner	Type the name of the owner of the cluster.
Storage class	*For SMS-managed datasets*, you can type the name of the storage class. Your storage administrator defines the range of names that you can use for storage classes.
Catalog	Type the name of the catalog in which you want to define the entry.
Erase	Specifies that the cluster is overwritten with binary zeros when its entry in the catalog is deleted.
Recatalog	Select this option to cause the catalog entries for the cluster components to be re-created from information in the VVDS, if valid VVDS entries are found on the primary VVDS volume.
Reuse	Select this option to enable the cluster to be opened repeatedly as a reusable cluster.
Spanned	Select this option to enable a data record to cross the boundaries of a control interval. If the maximum length of a data record (as specified with RECORDSIZE) is larger than a control interval, the record is contained on more than one control interval. This option allows VSAM to select a control interval size that is optimal for the DASD. You cannot use this option in the definition of a linear dataset cluster.
Speed	Select this option to prevent preformatting of the control areas of the data component during loading (creation) of a dataset.
Data organization type	By default, the data organization type is set to Indexed. Use this option to specify the type of data organization for the cluster: Indexed, Linear, Non-indexed, or Numbered.

IDCAMS DEFINE CLUSTER options that are available in the Dataset Define Cluster dialog box of the Storage Toolkit. Part 2 of 3, (optional) Data component options

Option	Description
Fully Qualified Name	Type the fully-qualified name of the data component.
Fully Qualified Model	Type the fully-qualified name of the data component that you want to use as a model.
Space Units	Select the unit of space that you want to allocate in the volume. Options include cylinders, KB, MB, records, or tracks. Then specify a corresponding value for the unit that you selected. The **Primary** value that you specify allocates the initial amount of space to the data component. The **Secondary** value that you specify allocates an amount of space for each time the data component extends.

Option	Description
Volser(s)	Specifies the volumes on which a component is to have space. The maximum number of VOLSERs is 59. Choose one of the following values in the list: • The serial number of an eligible direct access volume on which a new component is to be located OR • The serial number of a direct access volume on which an old component is located. If you choose an old component for **Volser**, the component must be on the specified volume for allocation to take place. If you do specify a value for **Volser**, the system allocates new components to any eligible direct access volume.
Control interval size	The size of the control interval for the component.
Buffer space	Specifies the minimum space for buffers. If you do not specify a value, VSAM attempts to get enough space to contain two data component control intervals and, if the data is key-sequenced, one index component control interval
Record Size	Specifies the average and maximum lengths, in bytes, of the records in the data component. The minimum record size is 1 byte. • **Nonspanned records:** For nonspanned records, the (maximum record size + 7) total cannot exceed the size of the control interval of the data component. In this case, the default values for Average and Maximum are 4089 and 4089, respectively. • **Spanned records:** When you use a record size that is larger than one control interval, you must also select the **Spanned** option in the **Cluster** tab. In this case the default values for Average and Maximum are 4086 and 32600, respectively.
Free Space	Specifies the percentage of each control interval (CI) and control area (CA) to be set aside as free space when the cluster is initially loaded or when a mass insert is done. The value of CI% is the percentage of space to be preserved for adding new records and updating existing records with an increase in the length of the record. Because a CI is split when it becomes full, the CA might also need to be split when it is filled by CIs created by a CI split. The empty space in the control interval and control area is available for data records that are updated and inserted after the cluster is initially loaded. This parameter applies only to key-sequenced clusters, and variable-length relative records with variable-length records. CI-percent is the number of bytes that is equal to, or slightly less than, the percentage value of CI-percent. CA-percent is the number of control intervals equal to, or less than, the percentage of CA-percent.
Share Options	These settings determine how a component or cluster can be shared among users. To ensure integrity, you should be sure that share options specified at the DATA and INDEX levels are the same. • **Cross Region:** Specifies the amount of sharing allowed among regions within the same system. • **Cross System:** Specifies the amount of sharing allowed among systems. Job steps of two or more operating systems can gain access to the same VSAM dataset regardless of the disposition indicated in the DD statement for the dataset in each step. For information regarding the options, see the description of the IDCAMS DEFINE CLUSTER command consult *DFSMS Access Method Services for Catalog*s. A general link to this documentation is provided in "Standard tabs in dialog boxes of Storage Toolkit" on page 157.

Option	Description
Keys	In this area, you can specify information about the prime key field of a key-sequence data records of a dataset. This parameter applies only to key-sequenced clusters. The default is a key field of 64 bytes, beginning at the first byte (byte 0) of each data record. The key field of the cluster's index is called the prime key to distinguish it from other keys, called alternate keys. • **Length:** Specifies the length of the key. • **Offset:** Specifies the displacement of the key (in bytes) from the beginning of the record. The sum of length plus offset cannot exceed the length of the shortest record. The length of the key can be 1 to 255 bytes.
Erase	Select this option to cause components to be overwritten with binary zeros when their entries in the catalog are deleted.
Reuse	Select this option to enable the component to be opened repeatedly as a reusable cluster.
Spanned	Select this option to enable a data record to cross the boundaries of a control interval. If the maximum length of a data record (as specified with **Record Size**) is larger than a control interval, the record is contained on more than one control interval. This option allows VSAM to select a control interval size that is optimum for the DASD.
Speed	Select this option to prevent preformatting of the control areas of the data component during loading (creation) of a dataset.

IDCAMS DEFINE CLUSTER options that are available in the Dataset Define Cluster dialog box of the Storage Toolkit. Part 3 of 3, (optional) Index component options

Option	Description
Fully Qualified Name	Type the fully-qualified name of the index component.
Fully Qualified Model	Type the fully-qualified name of the index component that you want to use as a model.
Space Units	Select the unit of space that you want to allocate in the volume. Options include cylinders, KB, MB, records, or tracks. Then specify a corresponding value for the unit that you selected. The **Primary** value that you specify allocates the initial amount of space to the data component. The **Secondary** value that you specify allocates an amount of space for each time the index component extends.
Volser(s)	Specifies the volumes on which a component is to have space. The maximum number of VOLSERs is 59. Choose one of the following values in the list: • The serial number of an eligible direct access volume on which a new component is to be located OR • The serial number of a direct access volume on which an old component is located. If you choose an old component for **Volser**, the component must be on the specified volume for allocation to take place. If you do specify a value for **Volser**, the system allocates new components to any eligible direct access volume.
Control interval size	The size of the control interval for the component.
Share Options	These settings determine how a component or cluster can be shared among users. To ensure integrity, you should be sure that share options specified at the DATA and INDEX levels are the same. • **Cross Region:** Specifies the amount of sharing allowed among regions within the same system. • **Cross System:** Specifies the amount of sharing allowed among systems. Job steps of two or more operating systems can gain access to the same VSAM dataset regardless of the disposition indicated the DD statement for the dataset in each step. For information regarding the options, see the description of the IDCAMS DEFINE CLUSTER command consult *DFSMS Access Method Services for Catalogs*. A general link to this documentation is provided in "Standard tabs in dialog boxes of Storage Toolkit" on page 157.

Option	Description
Reuse	Select this option to enable the component to be opened repeatedly as a reusable cluster.

Dataset Catalog dialog box

You use this dialog box to define a catalog entry for a single non-VSAM dataset you have selected in the Tivoli® Enterprise Portal.

You access this dialog box in the **Dataset Actions** submenu. This submenu is available in the pop-up menu that you access by right-clicking specific views in a workspace.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

You can also create a custom action. See "Create Batch Job dialog box" on page 162, "Issue Command dialog box" on page 165, and "Submit Command or Job dialog box" on page 166.

The action request defines a catalog entry for a non-VSAM datasets or a collection of objects. The dialog box provides the following options:

- **Catalog:** Identifies the catalog in which the non-VSAM dataset or OAM object is to be defined.
- **Device types:** Specifies the device types of the volumes containing the non-VSAM dataset being defined. If the non-VSAM dataset is located on different device types, the device types must be specified in the same order as the volume serial numbers listed in the **Volser(s)** parameter.
  You can specify a generic device name that is supported by your system, using up to six characters. For example, `3380`, `3390`, `3400-2`, and `3400-5` are valid device types.

  If you expect to change the device type of the system residence volume, you can enter four zeroes (`0000`) in one of the rows of the Device Types table. At LOCATE and DELETE time, this field is resolved to the device type. This method allows you to use the non-VSAM datasets without having to recatalog them to point to the new volume. A Device Type value of `0000` causes the actual device type to be determined from the current system residence volume (or its logical extension) at the time the catalog entry is retrieved.

  To avoid an erroneous result, you must use this wildcard convention (`0000`) with a **Volser(s)** parameter of the following types:

  - An indirect volume serial (`******`) identifier.
  - An extended indirect volume serial (a system symbol) identifier in the form *&xxxxx*, where *&xxxxx* is a symbol contained in some system library. The volume must be mounted and on line at the time the catalog entry is retrieved from the catalog. If it is not mounted and on line, the catalog request is terminated with a return and reason code.
- **Volser(s):** Specifies the volumes that contain the non-VSAM dataset. A **Volser(s)** value is required when you define a non-OAM non-VSAM dataset. As mentioned in the previous field description, you can also specify indirect volume serial identifiers and extended indirect volume serial identifiers.
  When you use the wildcard specification of six asterisks (`******`), the system dynamically resolves this specification to the system residence volume serial number, whenever the catalog entry is used. This value is not resolved when the DEFINE NONVSAM command is processed. This behavior allows you to later change the volume serial number of system residence volume without also having to recatalog the non-VSAM datasets on that volume.
- **Recatalog:** Specifies that the catalog entries are recreated if valid VVDS entries are found on the primary VVDS volume. RECATALOG can be specified only for an SMS-managed dataset.

Dataset Delete dialog box

You use this dialog box to issue an IDCAMS DELETE request to delete the datasets that you select in the Tivoli® Enterprise Portal.

You access this dialog box in the **Dataset Actions** submenu. This submenu is available in the pop-up menu that you access by right-clicking specific views in a workspace.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

You can also create a custom action. See “Create Batch Job dialog box” on page 162, “Issue Command dialog box” on page 165, and “Submit Command or Job dialog box” on page 166.

The dialog box provides the following options:

- **Entry type:** Select an option in this drop-down list to identify the type of object or entry to be deleted. If the object to be deleted is a catalog, truename entry, or VSAM volume record, you must specify its entry type as USERCATALOG, TRUENAME, NVR, or VVR. If you delete a migrated dataset without specifying the entry type, DFSMShsm™ deletes the dataset without recalling it.
- **Catalog:** Type the name of the catalog that contains the entries to be deleted.
- **Erase:** Select this option to cause the deleted entries to be overwritten with binary zeros.
- **Force:** Select this option to enable deletion of entries, even if they are not empty.
- **Purge:** Select this option to enable deletion of entries, even when a retention date has been specified.
- **Recovery:** Select this option to enable retention of information regarding the deleted records, in anticipation of a recovery operation. When you select **Recovery**, the **Erase**, **Force**, **Purge**, and **Noscratch** options are disabled.
- **Noscratch:** Select this option to enable deletion of the catalog entry from the catalog without mounting the volume that contains the object defined by the entry.

Dataset Listcat dialog box

You use this dialog box to issue an IDCAMS LISTCAT request for the datasets that you select in the Tivoli® Enterprise Portal.

You access this dialog box in the **Dataset Actions** submenu. This submenu is available in the pop-up menu that you access by right-clicking specific views in a workspace.

The dialog box provides the options shown in “IDCAMS LISTCAT options that are available in the Dataset Listcat dialog box of the Storage Toolkit” on page 180.

To learn about the standard capabilities of a dialog box, see “Standard tabs in dialog boxes of Storage Toolkit” on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

You can also create a custom action. See “Create Batch Job dialog box” on page 162, “Issue Command dialog box” on page 165, and “Submit Command or Job dialog box” on page 166.

IDCAMS LISTCAT options that are available in the Dataset Listcat dialog box of the Storage Toolkit

Option	Description
Target Entries	• **Use selected row entries:** Select this option to run the LISTCAT command against the datasets that you selected in the Tivoli® Enterprise Portal. • **Level:** Select this option to run the LISTCAT command against datasets that match the level of qualification that you specify in the text field.

Option	Description
Entry Type:	• **Alias:** Select this option to cause alias datasets to be listed. • **Alternate index:** Select this option to specify that entries for alternate indexes are to be listed. If **Alternate index** is specified and **Data** and **Index** are not also specified, entries for the data and index components of the alternate index are not listed. • **Cluster:** Select this option to specify that cluster entries are to be listed. If **Cluster** is specified and **Data** and **Index** are not also specified, entries for the data and index components of the cluster are not listed. • **Data:** Select this option to specify that entries for data components of clusters and alternate indexes are to be listed. • **Generation group:** Select this option to specify that entries for generation data groups are to be listed. GDSs in the active state, existing at the time the LISTCAT command is entered, are identified as such when ALL is specified • **Index:** Select this option to specify that entries for index components of key-sequenced clusters and alternate indexes are to be listed. If a VSAM name of the object is specified and INDEX is coded, only the index component entry of the object is listed. When INDEX is the only entry type parameter coded, the index component of the catalog is not listed. • **Non-VSAM:** Select this option to specify that entries for non-VSAM datasets are to be listed. • **Page space:** Select this option to specify that entries for page spaces are to be listed. • **Path:** Select this option to specify that entries for paths are to be listed. • **User catalog:** Select this option to specify that catalog connectors are to be listed. The user catalog connector entries are in the master catalog. (User catalog connector entries can also be in a user catalog, but the operating system does not recognize them when searching for a user catalog.)
Creation (days)	Type a **days** value to cause entries of the indicated type (CLUSTER, DATA, and so on) to be listed only if they were created in the specified number of days ago or earlier. The maximum number that can be specified is 9999. Zero indicates that all entries are to be listed.
Expiration (days)	Type a **days** value to cause entries of the indicated type (CLUSTER, DATA, and so on) to be listed only if they expire in the specified number of days or earlier. The maximum number that can be specified is 9999 and it indicates that all entries are to be listed. Zero indicates that only entries that have already expired are to be listed.
Listing type	Select this option to specify the fields to be included for each entry listed. The following list shows the options that are available: • **Name** specifies that the name and entry type of the entries are to be listed. Some entry types are listed with their associated entries. The entry type and name of the associated entry follow the name of the listed entry. • **History** specifies that only the following information is to be listed for each entry: name, entry type, ownerid, creation date, expiration date, and release. • **Volume** specifies that the information provided by specifying HISTORY, plus the volume serial numbers and device types allocated to the entries, are to be listed. • **Allocation** specifies that the information provided by specifying VOLUME plus detailed information about the allocation are to be listed. • **All** specifies that all fields are to be listed.
Catalog	Specifies the name of the catalog that contains the entries that are to be listed. If **Catalog** is coded, only the entries from that catalog are listed. See the "Catalog Search Order for LISTCAT" topic in the IDCAMS documentation for information about the order in which catalogs are searched.

Dataset Rename dialog box

You use this dialog to issue an IDCAMS ALTER request to rename a specific dataset that you select in the Tivoli® Enterprise Portal.

You access this dialog box in the **Dataset Actions** submenu. This submenu is available in the pop-up menu that you access by right-clicking specific views in a workspace.

The dialog box provides the following option: **Fully Qualified New name**. Type a new, fully-qualified name for the dataset in this field. This dialog box is a simplified version of the "Dataset Alter dialog box" on page 174.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

You can also create a custom action. See "Create Batch Job dialog box" on page 162, "Issue Command dialog box" on page 165, and "Submit Command or Job dialog box" on page 166.

Dataset Repro dialog box

You use this dialog box to issue an IDCAMS REPRO request for the dataset that you select in the Tivoli® Enterprise Portal.

You access this dialog box in the **Dataset Actions** submenu. This submenu is available in the pop-up menu that you access by right-clicking specific views in a workspace.

The dialog box displays options in the Repro and Merge/**DBCS** tabs that are located in the **Options** tab.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

You can also create a custom action. See "Create Batch Job dialog box" on page 162, "Issue Command dialog box" on page 165, and "Submit Command or Job dialog box" on page 166.

*IDCAMS REPRO options that are available in the Dataset Repro dialog box of the Storage Toolkit, Part 1 of 2, **Repro** tab*

Option	**Description**
Target Dataset section	Specify the fully qualified name of the target dataset in the **Fully Qualified Dataset** field.
Source Dataset Copy Options section	
Volume entries	Use this option to copy/merge a tape volume catalog to another tape volume catalog. Specify the entries of the source tape volume catalogs to be merged or copied in this field.
Error limit	Type a failure limit in this field. Use this parameter to set a limit to the number of errors that REPRO copy tolerates.
Replace	Select this option to indicate that a record in the source cluster is to replace a record in the target cluster when the source cluster is copied into the target cluster.
Reuse	Select this option if the target dataset is to be opened as a reusable dataset. This parameter is valid only for VSAM datasets.

Option	Description
Repro Range Attributes	Choose a **From/To** range or a **Skip/Count** range or **All** of the records. • **From/To:** When you choose the From/To range type, specify the following values: ◦ **Range Type:** This drop-down list offers the following options: ◦ **Key:** To specify the keys of the first and the last record you want to copy. ◦ **Address:** To specify the relative byte address (RBA) of the first and the last record to be copied. ◦ **Number:** To specify the relative record number of the first and the last record to be copied. ◦ **From and To fields:** Enter values that correspond to the **Range Type** that you specified: ◦ To target multiple items that have similar key names, you can specify an initial character string followed by a wildcard character (*). ◦ Other valid range values include decimal, hexadecimal, and binary numeric. • **Skip/Count:** When you choose the **Skip/Count** range type, specify the following values: ◦ **Number of records to skip:** Zero (0) is the default. ◦ **Count of records to copy:** The number of subsequent records to copy. All records, except those skipped, are copied by default. • **All:** Copy all of the records. To copy all of the records, choose **All** or either of the **From/To** and **Skip/Count** radio buttons, and leave the other fields blank.

*IDCAMS REPRO options that are available in the Dataset Repro dialog box of the Storage Toolkit, Part 2 of 2, **Merge**/**DBCS** tab*

Option	Description
Merge Entries section	**Mergecat:** You cannot select this option when **Use DBCS support** is selected. Select **Mergecat** to specify that the source catalog entries are to be merged with the target catalog entries and that the merged entries are to be deleted from the source catalog upon a successful merge operation. You can specify the entries to be merged either as the list of **Entries** (entry names or generic names) or as a single **Level** of qualification. • You can specify a list of names and generic names in the **Entries** field to match all the entries to be merged. When using a generic name with the ENTRIES parameter, entries must have one qualifier in addition to those specified in the command. • You can specify a **Level** value in order that all entries matching the level of qualification you indicated are to be merged.
Use DBCS support	Use this option to specify that bytes in the logical record contain DBCS characters. The **Repro** command checks to ensure the DBCS data meets DBCS criteria. This parameter cannot be specified with MERGECAT.
Skip section	Use the **Skip DBCS check** options in this section to indicate that characters between **Offset1** and **Offset2** are not to be checked for DBCS criteria during **Repro** command processing. You can specify several **Offset1**-**Offset2** pairs here. • **Offset1** indicates the byte offset in the logical record to be copied at which checking is to cease until offset2 is reached. • **Offset2** indicates the byte offset in the logical record after which checking is to resume. **Offset2** must be greater than **Offset1**. **No DBCS check:** Select this option to disable DBCS verification checking.

Option	Description
Shift section	Specify insertion of a DBCS shift in this section. Select **Insert Shift** to specify offsets. If SO and SI characters are not present, you must specify shift values. (By default, the logical record is assumed to already contain SO and SI characters, and the **Repro** command checks during copying to ensure that the DBCS data meets DBCS criteria.) • **Offset1** indicates the byte offset in the logical record to be copied before which a SO character is to be inserted. • **Offset2** Indicates the byte offset in the logical record to be copied after which a SI character is to be inserted. **Offset2** must be greater than **Offset1** and the difference must be an odd number. Select **Insert all** to specify that the entire logical record is assumed to contain only DBCS characters.

Dataset Uncatalog dialog box

You use this dialog to issue an IDCAMS DELETE NOSCRATCH request to uncatalog the datasets that you select in the Tivoli® Enterprise Portal.

You access this dialog box in the **Dataset Actions** submenu. This submenu is available in the pop-up menu that you access by right-clicking specific views in a workspace.

An **Options** tab is not provided, because there are no options to set.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

You can also create a custom action. See "Create Batch Job dialog box" on page 162, "Issue Command dialog box" on page 165, and "Submit Command or Job dialog box" on page 166.

Dataset Verify dialog box

You use this dialog box to issue an IDCAMS VERIFY request for the datasets that you select in the Tivoli® Enterprise Portal.

You access this dialog box in the **Dataset Actions** submenu. This submenu is available in the pop-up menu that you access by right-clicking specific views in a workspace.

The VERIFY command causes a catalog to correctly reflect the end of a VSAM dataset after an error occurs while closing a VSAM dataset. This dialog box does not have the **Options** tab, because there are no options to set.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

You can also create a custom action. See "Create Batch Job dialog box" on page 162, "Issue Command dialog box" on page 165, and "Submit Command or Job dialog box" on page 166.

Print dialog box

You use this dialog box to issue an IDCAMS PRINT request to print the datasets that you select in the Tivoli® Enterprise Portal.

You access this dialog box in the **Dataset Actions** submenu. This submenu is available in the pop-up menu that you access by right-clicking specific views in a workspace.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

You can also create a custom action. See "Create Batch Job dialog box" on page 162, "Issue Command dialog box" on page 165, and "Submit Command or Job dialog box" on page 166.

The dialog box provides the following options:

Print tab

Access the **Print** tab in the **Options** tab of this dialog box to set attributes for format and range of printing:

- **Format of Output:** Select **Character**, **Dump**, or **Hex**.
- **Print Range Attributes:** Choose a **From/To** range or a **Skip/Count** range or **All** of the records.
  - **From/To:** When you choose the **From/To** range type, specify the following values:
    - **Range Type:** This drop-down menu offers the following options:
      - **Key**: To specify the keys of the first and the last record you want to print.
      - **Address**: To specify the relative byte address (RBA) of the first and the last record to be printed.
      - **Number**: To specify the relative record number of the first and the last record to be printed.
    - **From** and **To** fields: Enter values that correspond to the **Range Type** that you specified:
      - To target multiple items that have similar key names, you can specify an initial character string followed by a wildcard character (*).
      - Other valid range values include decimal, hexadecimal, and binary numeric.
  - **Skip/Count:** When you choose the **Skip/Count** range type, specify the following values:
    - **Number of records to skip**: Zero (0) is the default.
    - **Count of records to print**: The number of subsequent records to print. All records, except those skipped, are printed by default
  - **All:** Print all of the records.

To print all of the records, choose **All** or either of the **From/To** and **Skip/Count** radio buttons, and leave the other fields blank.

> **Note:** When you are working in the Print dialog box and you select the **From/To** or **Skip/Count** radio buttons, the radio buttons remain selected. If you decide to deselect one of the radio buttons, you must click **Cancel** to dismiss the dialog box. Access the dialog box again and make the revised selections that you want.

DBCS tab

Access the **DBCS** tab to configure printing of dual-byte character sets (DBCS):

- **Use DBCS handling:** Select this check box to enable the printing of dual-byte characters.
  - **Skip Options:** Specify skip options in this area. Select the **Skip DBCS check** option to specify a table of **Offset1** and **Offset2** pairs, which are byte offsets in the logical record that define the beginning and the end of an area where DBCS checking is not performed. Select the **No DBCS check** option to disable validity checking for dual-byte characters.
    To omit any **Skip Options** in the IDCAMS PRINT command select the **Skip DBCS check** option and leave the offset table blank.
  - **Shift Options:** Specify shift options in this area. These options serve the scenario in which SO (shift out) and SI (shift in) characters are not present already in the record that you are printing. Select the **Insert shift** option to specify a table of **Offset1** and **Offset2** pairs, which are byte offsets in the logical record that indicate where the DBCS characters begin (**Offset1**) and end (**Offset2**).
    Select **Insert all** when the record to be printed consists of only dual-byte characters. This option causes a single SO character to be inserted at the beginning of each record and SI to be inserted at the end of the record.

    To omit any **Shift Options** in the IDCAMS PRINT command select the **Insert shift** option and leave the offset table blank.

Cancel HSM Request dialog box

Use the Cancel HSM Request dialog box to issue the Cancel HSM Request command to cancel active or queued DFSMShsm™ ("HSM") requests.

Unlike other dialog boxes for the Storage Toolkit, the **Options** tab is not displayed in the Cancel HSM Request command dialog box. This action has no HSM command parameter options to set.

To learn about the standard capabilities of a dialog box, see Standard tabs in dialog boxes of Storage Toolkit. The topic includes a link to the reference guide for commands that are used in the Storage Toolkit.

Accessing the Cancel HSM Request feature

You can access the Cancel HSM Request feature through the DFSMShsm Status workspace for IBM® Z OMEGAMON® AI for Storage®. You access and use the feature as follows:

1. Navigate to the DFSMShsm Status workspace in the Navigator.
2. Access the DFSMShsm Request Details workspace by clicking one of the blue link icons on the left side of the rows in the DFSMShsm Function Status Report table. This workspace includes the DFSMShsm Requests view that lists HSM requests.
   Or

   Access the CRQplex Requests workspace from the DFSMShsm Host Details, HSMplex CRQplex Details, or CRQplex Details workspaces. The CRQplex Requests workspace includes the CRQplex Request Report view.
3. In the DFSMShsm Requests view or the CRQplex Request Report view, select one or more rows (HSM commands) to cancel.
   Press the Control (**Ctrl**) key and click multiple rows, as needed.
4. Right-click on the selected rows to access the context menu.
5. Select **Cancel HSM Request**. The Cancel HSM Request dialog box is displayed.

   > **Note:** The menu item is available when all the items that you have selected are either queued or active. If the HSM request is queued, then the Toolkit request executes on the system from which the HSM request originated. If the HSM request is active, then the Toolkit request executes on the system on which the HSM request is processing. If the state of any selected item is **complete**, **abended**, or **cancelled**, the menu is not available.

6. As needed, enter options in the fields of the **General** tab, **Schedule** tab, and **JCL** tab of the dialog box. "Standard tabs in dialog boxes of Storage Toolkit" on page 157 describes these tabs.
   Use the **Replacement JCL JOB card** field in the **JCL** tab to assign a unique job ID to this job. If you do not do this, the batch job that is submitted to process the HSM cancel request might not run until any job that is currently executing with that job ID completes.
7. Click **OK** in the dialog box to confirm your settings. A confirmation dialog box is displayed.
8. Read the warnings in the confirmation dialog and then click **OK**, as appropriate.

See "Warning about cancellation of HSM requests" on page 186 for further information.

Warning about cancellation of HSM requests

Cancelling HSM requests can have negative consequences. This section introduces some issues that you should consider before cancelling an HSM request. See *z/OS® DFSMShsm™ Storage Administration Reference* (http://publibz.boulder.ibm.com/epubs/pdf/dgt2s440.pdf) for detailed information about these issues.

- **Active versus stopped:** The Cancel HSM Request feature is intended for use when *active* HSM tasks are suspended. Cancelling an HSM task that is actually moving toward completion might have unintended consequences.
- **Required followup activity:** When an **HSM RECOVER REPLACE** command is cancelled, it might be necessary to reallocate and recatalog the target dataset before you issue the **HSM RECOVER** command again.

- **When duplexing output tapes:** There is a limitation in **HSM BACKDS** processing when output tapes are duplexed. Under some circumstances, processing stops. Cancelling DFSMShsm™ might be required to release the original tape and tape drive.
- **Backup tasks:** Before you cancel a suspended HSM journal backup task, allow any active **HSM CS** backup tasks to run to completion. If an **HSM CDS** backup task is also suspended, cancel it before you cancel the HSM journal backup task that is suspended.
- **Common Recall Queue requests:** When cancelling an HSM request that is on the Common Recall Queue, the HSM cancel command will be executed on the appropriate z/OS system, which might not be the z/OS system you are viewing in the Navigator. If a CRQ request is in queued status, the HSM cancel command will be executed on the z/OS image from which the request originated. If the CRQ request is active, then the HSM cancel command will be executed on the z/OS system on which the request is being processed. To view the results of this command, you must navigate to the Tivoli Enterprise Portal Storage Toolkit node under the z/OS system on which the cancel command was executed.

DFSMSrmm actions in the Storage Toolkit

You find information about dialog boxes that are related to DFSMSrmm actions in the Storage Toolkit.

The DFSMSrmm TSO command is documented in the DFSMSrmm Guide and Reference (SC26-7404-07) available at http://publibz.boulder.ibm.com/epubs/pdf/dgt2r360.pdf. Consult this manual to identify required parameters.

You access the dialog boxes in the pop-up menu that opens when you right-click specific views in a workspace.

You can also create a custom action. See “Custom actions in Storage Toolkit” on page 161.

Add Dataset dialog box

You use this dialog box to issue the DFSMSrmm™ ADDDATASET command from RMM dataset-oriented workspaces.

You must specify a fully-qualified name for the dataset.

You can specify the following details for the dataset that you add:

- Fully Qualified DSN
- Volser
- Retention area:
  - Retention days
  - Expires on
- File sequence
- Block size
- Logical record length
- Record format

The **Options** tab contains the most commonly used keywords of this command. Not all keywords are presented in the **Options** tab. If a command keyword that you need is not presented, access the **Command** tab for this dialog box and enter the appropriate keyword and parameters in the **Command** input area.

To learn about the standard capabilities of a dialog box, see “Standard tabs in dialog boxes of Storage Toolkit” on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Change Dataset dialog box

You use this dialog box to issue the DFSMSrmm™ CHANGEDATASET command from RMM dataset-oriented workspaces.

The **Options** tab contains the most commonly used keywords of this command. Not all keywords are presented in the **Options** tab. If a command keyword that you need is not presented, access the **Command** tab for this dialog box and enter the appropriate keyword and parameters in the **Command** input area.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Delete Dataset dialog box

You use this dialog box to issue the DFSMSrmm™ DELETEDATASET command from RMM dataset-oriented workspaces.

The **Options** tab contains the most commonly used keywords of this command. Not all keywords are presented in the **Options** tab. If a command keyword that you need is not presented, access the **Command** tab for this dialog box and enter the appropriate keyword and parameters in the **Command** input area.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Add Owner dialog box

You use this dialog box to issue the DFSMSrmm™ ADDOWNER command from RMM owner-oriented workspaces.

The **Options** tab contains the most commonly used keywords of this command. Not all keywords are presented in the **Options** tab. If a command keyword that you need is not presented, access the **Command** tab for this dialog box and enter the appropriate keyword and parameters in the **Command** input area.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Change Owner dialog box

You use this dialog box to issue the DFSMSrmm™ CHANGEOWNER command from RMM owner-oriented workspaces.

The **Options** tab contains the most commonly used keywords of this command. Not all keywords are presented in the **Options** tab. If a command keyword that you need is not presented, access the **Command** tab for this dialog box and enter the appropriate keyword and parameters in the **Command** input area.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Delete Owner dialog box

You use this dialog box to issue the DFSMSrmm™ DELETEOWNER command from RMM owner-oriented workspaces.

The **Options** tab contains the most commonly used keywords of this command. Not all keywords are presented in the **Options** tab. If a command keyword that you need is not presented, access the **Command** tab for this dialog box and enter the appropriate keyword and parameters in the **Command** input area.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Add Volume dialog box

You use this dialog box to issue the DFSMSrmm™ ADDVOLUME command from RMM volume-oriented workspaces.

The dialog box presents the most frequently used parameters and options for the command.

The **Options** tab contains the most commonly used keywords of this command. Not all keywords are presented in the **Options** tab. If a command keyword that you need is not presented, access the **Command** tab for this dialog box and enter the appropriate keyword and parameters in the **Command** input area.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Change Volume dialog box

You use this dialog box to issue the DFSMSrmm™ CHANGEVOLUME command from RMM volume-oriented workspaces.

The **Options** tab contains the most commonly used keywords of this command. Not all keywords are presented in the **Options** tab. If a command keyword that you need is not presented, access the **Command** tab for this dialog box and enter the appropriate keyword and parameters in the **Command** input area.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Confirm Actions for All Volumes dialog box

You use this dialog box to issue the DFSMSrmm™ CHANGEVOLUME command from RMM volume-oriented workspaces.

This dialog box exploits the CONFIRM... parameters of this DFSMSrmm™ command.

The **Options** tab contains the most commonly used keywords of this command. Not all keywords are presented in the **Options** tab. If a command keyword that you need is not presented, access the **Command** tab for this dialog box and enter the appropriate keyword and parameters in the **Command** input area.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Delete Volume dialog box

You use this dialog box to issue the DFSMSrmm™ DELETEVOLUME command from RMM volume-oriented workspaces.

The **Options** tab contains the most commonly used keywords of this command. Not all keywords are presented in the **Options** tab. If a command keyword that you need is not presented, access the **Command** tab for this dialog box and enter the appropriate keyword and parameters in the **Command** input area.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Add VRS dialog box

You use this dialog box to issue the DFSMSrmm™ ADDVRS command from RMM VRS-oriented workspaces.

You must specify a fully-qualified name for the dataset.

The **Options** tab contains the most commonly used keywords of this command. Not all keywords are presented in the **Options** tab. If a command keyword that you need is not presented, access the **Command** tab for this dialog box and enter the appropriate keyword and parameters in the **Command** input area.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Change VRS dialog box

You use this dialog box to issue the DFSMSrmm™ CHANGEVRS command from RMM VRS-oriented workspaces.

The **Options** tab contains the most commonly used keywords of this command. Not all keywords are presented in the **Options** tab. If a command keyword that you need is not presented, access the **Command** tab for this dialog box and enter the appropriate keyword and parameters in the **Command** input area.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Delete VRS dialog box

You use this dialog box to issue the DFSMSrmm™ DELETEVRS command from RMM VRS-oriented workspaces.

The **Options** tab contains the most commonly used keywords of this command. Not all keywords are presented in the **Options** tab. If a command keyword that you need is not presented, access the **Command** tab for this dialog box and enter the appropriate keyword and parameters in the **Command** input area.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Situation Editor - Storage Toolkit Action Extension

The Situation Editor allows you to associate a Storage Toolkit requests that you want to run when a situation becomes true.

The **Storage Action** radio button appears when IBM® Z OMEGAMON® AI for Storage® is installed and a storage table is enabled for Storage Toolkit commands. The table is enabled if the `OPTION: SITUATIONACTIONEXT` tag exists for that particular table in the ODI file.

When you select the **Storage Action** radio button, the Storage Action area is displayed. The Storage Action area consists of three display fields:

Request Type

Identifies the type of request (if any) that was specified.

Request Title

Identifies the title of the request (if any) that was specified.

Request Location

Identifies the Storage agent on which the request (if any) will be executed.

and three action buttons:

Clear

Clears any existing Storage Toolkit requests associated with the selected node(s) for this situation.

Storage Toolkit

Begins a set of dialogs to create a Storage Toolkit request to associate with this situation. When you have finished creating the request, the Request Type, and Request Title, and Request Location fields are filled in with the relevant information.

?

Displays this help.

Associate Storage Toolkit Command

The Associate Storage Toolkit Command dialog allows you to create a Storage Toolkit request to associate with selected situation.

The number of Storage Toolkit actions associated with situations is limited by the length of the result command (508 bytes). There should be enough room for 5-8 nodes per situation.

Request creation type

Specify the type of command you want to create:

Create New

Create a new command. The available commands are:

- Create Command
- Create Batch Job
- Submit Command or Job
- Volume Backup
- Volume Convert
- Volume Migrate
- Volume Release Space
- Volume Compress
- Volume Defrag

Create From Existing

Create a command from an existing Storage Toolkit command.

Edit Existing Node

Modify an existing Storage Toolit command, add selected node to current command, or create a command from an existing command to then modify it.

New Request Information

Select a Storage Toolkit command from the list to associate with the selected situation.

Node

Select nodes where the command will be created.

Group Type

Specify which type of group you want to process against. There are different attributes associated with each type of group, whether SMS or User. This field is visible only when a Storage Toolkit command is being issued against the S3_Volume_Group_Summary attribute group.

Overview of common action commands

With Storage Toolkit actions, you can start commands on target resources in thez/OS ® system.

Table 1 shows typical goals for some action commands in Storage Toolkit.

Common commands enabled by Storage Toolkit	
Command type	**Goal**
HSM Volume Services	Migrate datasets on volume not referenced for a certain number of days Perform HSM incremental backup of eligible datasets on volume
HSM Dataset Services	Migrate dataset to level 1 Migrate dataset to level 2 Backup dataset Recover dataset from backup version Recall a migrated dataset

Command type	Goal
DFDSS Volume Services	Release unused space in datasets on volume Test volume conversion to SMS format Prepare volume for conversion to SMS format Convert volume to SMS format
DFDSS Dataset Services	Copy/move dataset Release unused space within dataset
Cancel HSM Request Service	Cancel active or queued DFSMShsm ("HSM") requests

For additional information, see "Using the Storage Toolkit" on page 146.

Managing Actions

You can view and manage action using workspaces.

The following workspaces help you view and manage actions:

- Storage Toolkit Action Requests workspace: The Storage Toolkit Action Requests workspace of the Storage Toolkit enables you to view and maintain reports on actions. You select and right-click items in the workspace to access context menus that allow you to create, submit, delete all, and delete the items. You can also determine if a Situation Storage Toolkit Extension request has been associated with a situation using the columns:
  Situation: Identifies the situation name associated with the toolkit command, if applicable.

  Associated With Situation: Identifies that the request is currently associated with the situation, if applicable.

  Action summaries and the data from the running of actions persist until you delete them. Select **Delete All Results** to delete all of the results associated with the selected action requests. Select **Delete Request** to delete the selected action requests and all of the results associated with them.
- Storage Toolkit Result Summary workspace: The Storage Toolkit Action Requests Summary workspace provides a link to the Storage Toolkit Result Summary workspace. As with the principal workspace, you select and right-click items to select management operations, such as delete. Action summaries and the data from the running of actions persist until you delete them. Select **Delete Result Details** to delete the output data associated with the selected results. Select **Delete** to delete the selected results and the output data associated with the selected results.
- DFSMShsm Request Details workspace: The DFSMShsm Request Details workspace provides context menus that enable you to select and cancel active or queued DFSMShsm requests. See "Cancel HSM Request dialog box" on page 186.

Associating Storage Toolkit commands with situations

With the Situation Storage Toolkit Extension, you specify a Storage Toolkit command to be run when a situation becomes true and the attribute group or table named in the situation is enabled for the Situation Storage Toolkit Extension.

Here is an example of associating a new Storage Toolkit command with a situation:

1. Select a node in the navigation tree to which you want to associate the new situation.

> **Note:**
> - Use)DOT and)ENDDOT constructions inside your JCL because only one action is submitted for all triggered rows in the sample interval.

- Use of the)DOT and)ENDOT to aggregate results from multiple workspace rows into a single batch job is only supported for direct toolkit actions, such as performing a multi-row select in a workspace view and launching a toolkit action from the right-click menu.
- Use of the)DOT and)ENDOT is not supported in situation driven toolkit actions because the platform framework does not support launching a single action for multiple situation events.

2. Right-click the node to display the pop-up menu. Select the **Situations** menu item to display the Situation Editor dialog.
3. Click the **Create New Situation** icon in the top-left corner of the Situation Editor dialog.
4. Specify a name for the new situation.
5. Click **OK** to continue. The Select Condition dialog is displayed.
6. Select an attribute group in the Attribute Group (left) pane of the Select Condition dialog. The members of the attribute group will be displayed in the Attribute Item (right) pane.
7. Select one or more attributes from the Attribute Item pane.
8. Click **OK** to continue. The Select Condition dialog closes, and the **Formula** tab of the Situation Editor dialog is displayed.
9. Enter appropriate values in the Formula pane for the situation attributes that you chose.
10. Select the **Distribution** tab.
11. Select a Managed System to which the situation will be distributed.
12. Select the **Action** tab.
13. Select the **Storage Action** radio button in the Action Selection area. A new Storage Action pane is displayed.
14. Click the **Storage Toolkit** button in the Storage Action pane to display the Associate Storage Toolkit Command dialog box.
15. Select the **Create New** radio button.
16. Select a Storage Toolkit command to associate with the situation.
17. Select nodes where the command will be created.
18. Click **OK** to proceed to the Storage Toolkit Command dialog.
19. Use the Storage Toolkit dialogs to complete your command specifications. See "Storage Toolkit" on page 143 for more information on using the Storage Toolkit dialogs. When complete, click **OK** on the Storage Toolkit Command dialog to return to the Situation Editor.
20. Click **OK** in the Situation Editor dialog to save the situation and close the dialog.
21. Select the Storage Toolkit node in the navigation tree to navigate to the Storage Toolkit Action Requests Workspace where you can view the status of the Storage Toolkit action request that you just created.

User actions in Storage Toolkit for Realtime Dataset Metrics only

Learn about the dialog boxes that are related to Storage Toolkit user actions that you can create and save for later use in Realtime Dataset Metrics (RDM).

Note: These actions are only available in the RDM application, they are not available in TEP.

To access user actions, do the following:

1. Select and right-click the rows in the workspaces that have action menus.

2. Select **User Actions** from the menu.
3. Choose the appropriate action from the submenu.
   When you select an action, a dialog box is displayed in which you configure the action and submit the request.

User actions commands	
Name of menu item	**Corresponding dialog box**
New Batch Job	Create Batch Job template dialog box
New Command	New Command template dialog box

User action subcommands	
Name of menu item	**Corresponding dialog box**
Apply	Apply Create Batch Job or Command dialog box filled with template values
Edit	Edit Batch Job or Command template dialog box filled with template values
Delete	Confirmation dialog box

Create Batch Job template dialog box

Use this dialog box to create a custom Job template for later use. With this action, you can prefill the values for the Create Batch Job dialog box, as described in “Create Batch Job dialog box” on page 162, and save it with a custom name typed in the **Action name** field.

When saved, the action is displayed under the **User Action** context menu group with the following options:

- Apply - Submit user action.
- Edit - Change user action template.
- Delete - Delete the user action template.

New Command template dialog box

Use this dialog box to create a custom Command template for later use. With this action, you can prefill the values for the Issue Command dialog box.

When saved, the action is displayed under the **User Action** context menu group.

- Apply - Submit user action.
- Edit - Change user action template.
- Delete - Delete the user action template.

For related information, see “Using the Storage Toolkit” on page 146.

Dataset Group feature

With the Dataset Group feature, you can quickly navigate to problem areas and automatically apply solutions.

OMEGAMON® for Storage offers several GUI-based features to help you identify and solve storage-related application problems. With the graphical user interface, you can locate problem areas and automatically apply solutions, instead of using more time-consuming methods in the z/OS® operating system.

- To leverage the Dataset Group feature, you should know the following:
  - How to add, edit, delete, and set the collection interval for dataset groups.
  - Dataset masks and the rules for constructing them.
  - How to use dataset groups to resolve storage problems.
- Another problem-solving tool is described in "Storage Toolkit" on page 143.

The Dataset Group feature versus the Dataset Attributes feature

Depending on your requirements, use the Dataset Group feature and the Dataset Attribute feature to manage datasets.

OMEGAMON® for Storage gives you two powerful facilities for managing and monitoring datasets: the Dataset Group feature and the Dataset Attribute feature. Both of these features provide similar (although not same) information, you must know when one approach is better than the other.

- Construct and use Dataset Groups when you have specific sets of similarly named datasets to be monitored, either for space or performance.

  > **Note:** To see dataset-level I/O performance metrics, the volumes on which those datasets are located must have been defined for dataset-level I/O monitoring.

- Enable Dataset Attribute collection when you want to monitor overall space utilization, exceptions, statistics, trends, and other metrics with respect to datasets.
- Define Situations for Dataset Group tables that raise alerts if specific groups of datasets are in danger of running out of space or are performing poorly, as described in this section.
- Define Situations against Dataset Attribute tables to raise alerts when any dataset exceeds a defined threshold (for example, number of extents) or when overall conditions (for example, total system free space) warrant that a warning be raised.

Working with dataset groups

Learn how to define a group of datasets based on filtering rules and set the collection interval for a dataset group.

> **Note:** The availability of menu selections in the following procedures depend on your security permissions.

Dataset Group Properties dialog box

You use this dialog box to define a group of datasets based on filtering rules (masks) and related properties.

You access this dialog box through the Dataset Group Summary workspace, which is described in "Monitoring dataset groups summary data" on page 118. You can change the following settings:

- **RLS Group:** Select this checkbox to enable collection of RLS data for each member of the group.

- **Mask:** You type a mask definition following the rules used by DFDSS dataset name masking. See "Dataset group masking rules" on page 197.
- **Active:** Select this checkbox to cause data collection from the catalog to be performed for the datasets that fall under a mask.
- **Space Data:** Select this checkbox to collect space data for the datasets that fall under a mask.
- **RLS Mask:** Select this checkbox to collect RLS data for the datasets that fall under this dataset mask.

To learn how to add and configure a dataset group, see "Adding a dataset group" on page 196.

Dataset Group Data Collection dialog box

You use this dialog box to set the interval at which data collection occurs.

You access this dialog box through the Dataset Group Summary workspace, which is described in "Monitoring dataset groups summary data" on page 118. You can change the following settings:

- **RMF interval:** Data is collected at the predefined time interval of the resource measurement facility (RMF).
- **On Demand:** Data is collected when you request collection.
- *nnn* Mins (5 - 360). Data is collected at the interval of minutes that you define, where *nnn* is an integer from 5 to 360.

Adding a dataset group

Use the following procedure to add a new dataset group.

1. In the navigator, click **Dataset Group Summary**.
2. Right-click a row within the Dataset Group Summary table.
3. Click **Add Group**.
4. Enter a group name in the **Add New Group** window.
5. Click **OK**.
6. In the **Dataset Group Properties** window, type a mask value.
7. Press **Enter**.
8. Edit the group properties as needed.
   You can specify the following:
   - **Mask:** A mask against which catalogued dataset names are processed. Masks follow the rules used by DFDSS dataset name masking.
     For additional information about mask rules, see "Dataset group masking rules" on page 197.

     To add a new mask, click the **Add a new mask...** button located in the upper-left corner of the dialog box.
   - **Active:** When selected, data collection from the catalog is performed for the datasets that fall under the mask.
   - **Space Data:** When selected, space data is collected for the datasets that fall under the mask.
9. Click **OK** to finish adding the dataset group.

AFTER COMPLETING THE TASK

Editing or deleting dataset groups

You can edit a group definition in the Dataset Summary table. To do so, right-click the group row, and then click **Edit Group**. Then, you can edit the group properties.

To delete a group in the Dataset Summary table, right-click the group row, then click **Delete Group**.

Setting the collection interval

Use the following procedure to specify the collection interval.

1. In the Dataset Summary table, right-click the row for any group.
2. Click **Collection Interval**. You set the interval for all groups in one operation.
3. Specify the group data collection interval by selecting one of the three options:
   - **RMF Interval:** When selected, the collection interval is synchronized to the RMF interval.
   - **On Demand:** When selected, the data is collected in real time. However, data is collected every 5 minutes.
   - **Mins:** When selected, the collection interval can be specified to take place automatically every *N* minutes between 5 and 360 minutes.
4. Click **OK** to set the collection interval.

Dataset group masking rules

Learn how to use dataset group masking rules to filter and retrieve the exact dataset names you need by applying flexible wildcard patterns with asterisks (*) and percent signs (%).

A dataset group mask is a character string that describes the dataset group names for which you want information returned.

Asterisks (*) and percent signs (%) are used to represent qualifiers or parts of qualifiers.

The information in the **Mask characters** table explains the usage of asterisks and percent signs in masks.

Mask characters	
Mask characters	**Description**
*	A single asterisk by itself indicates that either a qualifier or one or more characters within a qualifier can occupy that position. An asterisk can precede or follow a set of characters.
**	A double asterisk indicates that zero or more qualifiers can occupy that position. A double asterisk cannot precede or follow any characters; it must be preceded or followed by either a period or a blank.
%	A single percent sign by itself indicates that exactly one alphanumeric or national character can occupy that position.
%%	One to eight percent signs can be specified in each qualifier.

Mask scenarios

Consider a scenario in which you have the following dataset groups:

- MYGROUP.DATA.SET1
- MYGROUP.DATA.SET2
- MYGROUP.DATA.SET30
- VSAM.MYGROUP.DATA.SET

In the **Sample masks** table, you can find some mask examples and the information that each mask returns.

Sample masks		
Mask	**Information returned**	**Information excluded**
MYGROUP.DATA.*	MYGROUP.DATA.SET1 MYGROUP.DATA.SET2 MYGROUP.DATA.SET30	VSAM.MYGROUP.DATA.SET

Mask	Information returned	Information excluded
.DATA.	MYGROUP.DATA.SET1, MYGROUP.DATA.SET2 MYGROUP.DATA.SET30, and VSAM.MYGROUP.DATA.SET	-
MYGROUP.DATA.SET%	MYGROUP.DATA.SET1 and MYGROUP.DATA.SET2	MYGROUP.DATA.SET30 and VSAM.MYGROUP.DATA.SET
MYGROUP.DATA.SET%%	MYGROUP.DATA.SET30	MYGROUP.DATA.SET1, MYGROUP.DATA.SET2, and VSAM.MYGROUP.DATA.SET

Using dataset groups

The Dataset Group feature gives you a powerful facility to examine and monitor the space and performance attributes of a collection of datasets.

By defining a series of dataset groups and masks, any collection of datasets can be displayed as a report within the portal, and Situations can be written to monitor the space or performance attributes of those datasets.

Here you can find some examples and ideas of how this function can add value for storage management and administration.

Think of Dataset Groups as of a more powerful version of the ISPF DSLIST utility (that is, 3.4). In addition to creating lists of datasets, with OMEGAMON® Dataset Group feature, you can name and save those lists in groups and provides much more powerful name masking capabilities.

All the attributes associated with the datasets in the generated lists can also be used in reports, alerts, and Storage Toolkit actions. The OMEGAMON® Dataset Group feature is not intended as a direct substitute for ISPF 3.4.

ISPF 3.4 is geared toward producing interactive, user-generated ad hoc lists of datasets and running simple actions against those datasets, one at a time. OMEGAMON® Dataset Groups are more static in nature and are better used for viewing large sets of related files and their attributes (individually and as a group), generating alerts based on those attributes, and taking actions.

> **Note:** The Dataset Group feature differs from the User DASD Group and Application functions in the following way. Dataset Groups start at the system catalog by using a name mask to create a list of matching datasets, while User DASD Groups start at the volume or storage group level, and Applications start at the job level, to create their volume/dataset lists.
> The important difference here is that for using Dataset Groups, you do not need to know where a dataset is located or which application accesses it. Using this dataset orientation can be more useful for end-users and storage administrators who prefer to perform certain tasks from this perspective.

The following examples describe how to use OMEGAMON® for Storage to monitor events that might occur in your system. They are:

- Obtaining an overview of space usage
- Identifying datasets with excess space
- Create an alert when response time exceeds specific thresholds
- Identify VSAM datasets with excessive CA/CI splits
- Cataloged and not-allocated datasets
- Notify the Production Scheduling group when a critical dataset exceeds a specified number of extents

- Automatically migrate infrequently used datasets

Obtaining an overview of space usage

Follow this procedure if you want to determine the amount of space that is used by your applications. Initially, you can create a group of your critical datasets, possibly dividing them into development, test, and production datasets.

1. Use the processes described in "Adding a dataset group" on page 196 to add a new dataset group. In the case of the example shown in Figure 2, there are four groups, for development, test, production, and history datasets.

Identifying datasets with excess space

Use dataset groups to identify datasets with extra allocated space.

PREREQUISITE

If DFSMS Data and Storage Class ACS routines are not configured for space restrictions or for releasing unused primary space, the total space might exceed the specified threshold.

To monitor this condition, perform the following steps:

- Create a group of datasets.
- Create a threshold to monitor the Largest Free Space for those datasets.
- After you have isolated a group of datasets, it is possible to analyze individual datasets.

> **Note:** You can also use the Dataset Attribute Summary workspace to monitor space usage on a system-wide basis if the group orientation does not provide sufficient scope for your storage management activities.

To set up this threshold, perform the following steps:

1. Right-click anywhere within the Dataset Group Summary table.
2. Click **Properties** to open the Properties editor.
3. Click the **Thresholds** tab.
4. Select the **Largest Space Free** attribute.
5. Change the operator in the first row to Greater than (>) and the value to 1000 (or any reasonable value for your system).
6. Check that the indicator is set to critical.
7. Repeat the process on the second and third line to set a threshold that indicates an acceptable warning and informational level.
8. Click **Apply** and **OK**.
   Whenever the Largest Free Space value for the dataset group exceeds the specified thresholds, that group is highlighted in the table view. You can analyze the resulting list of datasets and take appropriate action.

Create an alert when response time exceeds specific thresholds

Use the Application Monitoring feature of OMEGAMON to create alerts by using the name of the JOB, STC, or TSO user whose datasets need to be monitored.

PREREQUISITE

Note: There are other cases, where you need to monitor a group of datasets, no matter how or where they are accessed.

For this scenario, do the following:

- Create a group called for example Critical, encompassing the critical datasets.
- Define both Group Summary thresholds and situations to generate an alert when response time (MSR) or the number of I/O events exceeds the expected limits.

You can create a situation by using the following procedure.

1. Right-click the Dataset Group Summary node in the navigation pane.
2. Click **Situations** in the pop-up menu.
   The Situations for - Dataset Group Summary dialog box is displayed.
3. Click the **Create new Situation icon** to access the Create Situation dialog box.
4. Enter the name and description of your new situation.
5. Click **OK** to access the Select condition dialog box.
6. Select **S3_Dataset_Group_Details** in the Attributes Group column.
7. While pressing the **Ctrl** key, select both **Dataset Group Name** and **Maximum MSR** in the **Attribute Item** column.
8. Click **OK** to return to the Situations for - Dataset Group Summary dialog box.
9. In the **Formula** tab, enter appropriate values in the fields. For example, enter an operator of **Equal (==)** and a value of "Production" for the Dataset Group Name, and an operator of **Greater than (>)** and a value of 20 for the Maximum MSR, see Figure 4.
10. Select the **Distribution** tab.
11. Add the appropriate managed systems to the Assigned systems list for this situation.
12. Click **Apply** and **OK** to complete the definition your new situation and close the dialog box.
    Whenever a response time exceeds specific thresholds, this situation is triggered and an alert is sent.

Identify VSAM datasets with excessive CA/CI splits

Learn how to identify VSAM datasets with excessive CA/CI splits and efficiently fix those issues.

PREREQUISITE

VSAM datasets undergo CI and CA splits as they are updated and records are added or inserted. This behavior is not always a bad thing, however for some dataset update patterns it can cause excessive processing and I/O.

Note: You can also use the CA/CI Split Summaries from the Dataset Attribute Summary workspace to see the datasets with the largest number of splits across the entire enterprise, rather than limited to defined groups.
To monitor for this condition, perform the following steps:

- Define a group encompassing the VSAM datasets that you know do not experience high split rates.
- Create a Situation to generate a warning alert when any dataset in that group exceeds the threshold.
- Using Action or Workflow you can also automatically issue a command that would be trapped by, for example, AOC to cause the dataset in question to be REPROed and reDEFINEd to eliminate the splits.

Assuming that you have identified and grouped the VSAM datasets that you want to monitor, create a situation as follows:

1. Right-click the **Dataset Group Summary** node in the navigation pane.
2. Click **Situations** in the pop-up menu. The Situations for - Dataset Group Summary dialog box is displayed.

3. Click **Create new Situation** icon to access the Create Situation dialog box.
4. Enter the name and description of your new situation.
5. Click **OK** to access the Select condition dialog box.
6. Select **S3_Dataset_Group_Details** in the Attributes Group column.
7. While pressing the **Ctrl** key, select both **CA Splits** and **CI Splits** in the Attribute Item column.
8. Click **OK** to return to the Situations for - Dataset Group Summary dialog box.
9. In the **Formula** tab, enter appropriate values in the fields. Change the operators to Greater than (>) and enter the value 8 for CA Splits and 4 for CI Splits.
10. Click **Apply**.
11. To view the formula, click the formula icon in the right side of the Situation editor panel.
12. Select the **Distribution** tab.
13. Add the appropriate managed systems to the Assigned systems list for this situation.
14. Click **Apply** and **OK** to complete the definition of your new situation and close the dialog box. Whenever excessive CA/CI splits exist, this situation is triggered, and the correction action can occur automatically.

Cataloged and not-allocated datasets

Learn how to identify cataloged and not-allocated datasets with the Dataset Group feature.

PREREQUISITE

Catalog all datasets in a DFSMS environment. Although you follow this policy, there can still be cases where a dataset exists in the catalog but has been moved or deleted from the specified volume.

> **Note:** As a converse to this scenario, you can use the Uncataloged Dataset Summary report of the Dataset Attribute Summary to identify datasets that are allocated, but uncataloged.
> To monitor for this condition, perform the following steps:
> - Create a dataset group.
> - Create a situation by using the Allocation Flag attribute from the Dataset Group Details Volume attribute group in the Situation editor.

Create a situation as follows:

1. Right-click the **Dataset Group Summary** node in the navigation pane.
2. Click **Situations** in the pop-up menu. The Situations for - Dataset Group Summary dialog box is displayed.
3. Click the **Create new Situation** icon to access the Create Situation dialog box.
4. Enter the name and description of your new situation.
5. Click **OK** to access the Select condition dialog box.
6. Select **S3_Dataset_Group_Details_Volume** in the Attribute Group column.
7. Select **Allocation Flag** in the Attribute Item column.
8. Click **OK** to return to the Situations for - Dataset Group Summary dialog box.
9. In the **Formula** tab, enter values for Allocation Flag as follows:
    a. On the first line of the Allocation flag, open the drop-down list and select **CANDIDATE VOLUME/ CATALOG ENTRY ONLY**.
    b. On the second line of the Allocation flag, open the drop-down list and select **OVERFLOW VOLUME/ CATALOG ENTRY ONLY**.
10. To view the formula, click the **Formula** icon in the right side of the **Situation editor** panel.

11. Select the **Distribution** tab.
12. Add the appropriate managed systems to the Assigned systems list for this situation.
13. Click **Apply** and **OK** to complete the definition your new situation and close the dialog box.

Notify the Production Scheduling group when a critical dataset exceeds a specified number of extents

Learn how to notify the Production Scheduling group when a critical dataset exceeds a specified number of extents.

PREREQUISITE

Overnight batch jobs have a small window to complete, leaving less room for ABENDs and reruns. If a production dataset exceeds the available space and the batch update job terminates with an x37 ABEND, online applications might not start in time, which might affect business.

To monitor for this condition, perform the following steps:

- Create a Dataset Group containing the critical production datasets.
- Define a Situation to raise an alert if any of those datasets approaches the 16-extent limit.
- Add an action to the Situation to send a pager message to the relevant person.

> **Note:** You can use the name of the dataset and the number of its extents as substitution variable in the action.

Create the following situation:

1. Right-click the **Dataset Group Summary** node in the navigation pane.
2. Click **Situations** in the pop-up menu. The Situations for - Dataset Group Summary dialog box is displayed.
3. Click the **Create new Situation** icon to access the Create Situation dialog box.
4. Enter the name and description of your new situation.
5. Click **OK** to access the Select condition dialog box.
6. Select **S3_Dataset_Group_Details_Volume** in the Attribute Group column.
7. Select **Extents** in the Attribute Item column.
8. Click **OK** to return to the Situations for - Dataset Group Summary dialog box.
9. In the **Formula** tab, enter values for Extents. Change the operator to Greater than (>) and enter the value 12.
10. Select the **Distribution** tab.
11. Add the appropriate managed systems to the Assigned systems list for this situation.
12. Select the **Action** tab.
13. Click **Universal message** and enter `send message &dsname`.
14. Click **Apply** and **OK** to complete the definition of your new situation and close the dialog box. Whenever the number of extents for your critical dataset exceeds your threshold, this situation is triggered and notification is sent.

    The paging command that you use depends on the underlying product in the Windows or z/OS environment.

Automatically migrate infrequently used datasets

Learn how to configure automatic migration of infrequently used datasets with the Dataset Group feature.

PREREQUISITE

During the course of a day a lot of permanent datasets are created, possibly accessed once, and then never referenced again. If those datasets are not automatically migrated by DFHSM from disk to near-line storage, they waste valuable space.

To monitor for this condition, perform the following steps:

- Create a Situation by using the Dataset Group Details attribute group.
- When this situation is true, the **HMIGRATE** command is issued at the host to cause DFHSM to migrate the dataset in question.

> **Note:** This process does not mean to replace Management class and DFSMShsm processing options, but rather a way of augmenting those features by using OMEGAMON intelligent functions.

Create a Dataset Group for permanent datasets that you know are created, accessed once, and should then be migrated. Then create the following situation:

1. Right-click the **Dataset Group Summary** node in the navigation pane.
2. Click **Situations** in the pop-up menu.
   The Situations for - Dataset Group Summary dialog box is displayed.
3. Click the **Create new Situation** icon to access the Create Situation dialog box.
4. Enter the name and description of your new situation.
5. Click **OK** to access the Select condition dialog box.
6. Select **S3_Dataset_Group_Details** in the Attribute Group column.
7. Select **Primary Volume** in the Attribute Item column.
8. Click **OK** to return to the Situations for - Dataset Group Summary dialog box.
9. In the **Formula** tab, enter values for Primary Volume. Change the operator to Not equal (!=) and enter the value `MIGRAT`.
10. Click **Add Conditions** to access the Select condition dialog box.
11. Choose **Days Unreferenced** in the Attribute Item column.
12. Click **OK** to return to the Situations for - Dataset Group Summary dialog box.
13. In the **Formula** tab, enter values for Days Unreferenced. Change the operator to Greater than (>) and enter the value `1`.
14. Click **Add Conditions** to access the Select condition dialog box.
15. Choose **Dataset Group Name** in the Attribute Item column.
16. Click **OK** to return to the Situations for - Dataset Group Summary dialog box.
17. In the **Formula** tab, enter values for Dataset Group Name. Change the operator to Equal (==) and enter the value `groupname`. For example, the group name could be PAYROLL_TEMP. This group can be a specific set of datasets that you know are created and then referenced only once.
18. Select the **Distribution** tab.
19. Add the appropriate managed systems to the Assigned systems list for this situation.
20. Select the **Action** tab.
21. Enter the command you want to use in the System command entry box. For example, you could enter **`HMIGRATE &S3_Dataset_Group_Details:Dataset_Name`**. To add the variable, click **Attribute Substitution** and select **Dataset name**.
22. Select the **Formula** tab and click the `fx` button.
23. Click **Apply** and **OK** to complete the definition of your new situation and close the dialog box.

Dataset Attributes Database

With the Dataset Attributes Database feature, you can conveniently manage your datasets.

IBM® Z OMEGAMON® AI for Storage® provides Dataset Attributes Database to help you manage datasets. This database maintains attributes related to dataset space, DCB, and cataloging. The principal workspace links to related workspaces that track data for common aspects of dataset management and administration.

For instructions on how to configure the feature, see "Configuring the Dataset Attribute Database feature" on page 207.

Benefits

With Dataset Attributes Database, you can get summary reports on general trends, identify resources that require attention, initiate Storage Toolkit action commands for selected datasets, and more.

The database collects data regarding all datasets on all volumes in the z/OS® environment (except for volumes that you have excluded), so that you can perform the following operations:

- View comprehensive summary reports regarding general trends or specific low-level details.
- Identify exception conditions regarding datasets throughout the environment.
- Select datasets from a workspace view and initiate Storage Toolkit action commands against them.
- Obtain an installation-wide view of dataset space utilization, exceptional conditions, and summary statistics.
- Quickly identify resources that require attention, such as datasets that have excessive unused space, extents, or CA/CI splits.
- Avoid the use of utility programs that have limited ability to aggregate and display data, such as IEHLIST.

Types of Dataset Attribute collection

There are three types of Dataset Attribute collection: Full collection, Partial Sweep collection, and Incremental collection.

Full collection

This collection type is used for the initial collection after the OMEGAMON for Storage Agent's startup.

Full collection performs a complete scan of all volumes for VTOC data and collects information about datasets. And then if the Catalog data is requested, it additionally retrieves the Catalog data.

Partial Sweep collection

Partial Sweep collection goes through the VTOC entries of datasets that have changed instead of performing a full VTOC scan of all volumes.

This allows the Agent to update all datasets from the changed volumes (compared to the previous Full collection or Partial Sweep collection). As a result, the Dataset Attribute Database gets information about the latest datasets and consumes less CPU time than the Full collection.

Incremental collection

Incremental collection is also called Delta collection.

If OMEGAMON for Storage is started on all LPARs in the PLEX, after having an initial Full collection, the system becomes a reporter of data. Incremental collection retrieves dataset changes (deltas) from all other OMEGAMON for Storage Agents (from other LPARs in the PLEX) and applies them every specified interval (15min, 30min, 60min, 3h, 6h).

This allows the Agent to keep the Dataset Attribute Database up to date without high CPU consumption.

Optimization

To optimize performance, configure the Dataset Attributes Database feature.

Otherwise, the amount of data that you collect might consume excessive system resources.

For more information, see "Configuring the Dataset Attribute Database feature" on page 207 and also "Optimizing data collection for the Dataset Attributes Database feature" on page 209.

For example, you can optimize configuration of the following behaviors:

- Exclusion of specific SMS Storage Groups and Volumes
- Collection interval
- Number of items listed in the Summary reports
- Inclusion of catalog data

Data collected for the Dataset Attributes Database is maintained in the in-memory data space for high-performance access. The data is not saved to disk, except for those Summary tables that have been configured for historical collection.

For this reason, after a restart of the Tivoli® Enterprise Monitoring Server address space, you must initiate a collection cycle manually or at a scheduled time to rebuild the database.

> **Tip:** It is recommended to schedule collection once per week (for example, during the weekend) and Incremental collection every interval.
> The more often Incremental collection runs, the more CPU will be used. Most of the CPU is used to update the summarization workspaces and update saved Dataset Attribute Groups. You can choose between having more actual data and reducing CPU consumption.

Attribute derivation and definition

This section clarifies the concepts and values associated with specific attributes.

Percent used and free

- Percent Used and Free for PDS, PDSE, and HFS datasets is calculated based on the allocated and used tracks from the VTOC, and does not take into account internal usage or wastage within the dataset.
- Tracks Used is set equal to Tracks Allocated for PDSE and HFS datasets.
- For VSAM datasets Percent Used and Free is based on the volume-level High Used and High Allocated RBAs (relative byte addresses).

Space values

The following details govern the collection and display of space values in the product:

- Within the Dataset Attributes Database reports and queries, all space values are displayed as Tracks, MBs or GBs. By default, the space unit for all delivered Views is Tracks. The default space value columns can be changed as follows:
  a. Modify the **Filters** tab of the Properties dialog box for the workspace.
  b. Save the workspace under a new name.

  > **Note:** Refer to the *IBM® Tivoli® Monitoring: User's Guide* to learn how to display alternate columns in a predefined workspace.

- MBs are defined as `1024*1024`. GBs are defined as `1024*1024*1024`.
- All calculated MB and GB values are calculated from the underlying byte totals, and are rounded up.

- Space values in summary reports for multivolume datasets are for the portion of the dataset located on the indicated volume.
- The Dataset Detail workspace lists all individual pieces of a multivolume dataset (assuming that all volumes on which it is located were processed during collection and selected by the Query).
- Tracks allocated/used
  - **Tracks Allocated** is always the sum of all tracks in all extents on the volume from the DSCB.
  - **Tracks used for non-VSAM datasets** comes from DS1LSTAR in the DSCB.
  - **Tracks used for VSAM datasets** is calculated as follows:

    ```
    High Used RBA on the volumes/(VSAM physical record size) *
    (VSAM physical records per track)
    ```

    That value is then used to calculate Percent Used and Percent Free and to calculate the total MBs and GBs used.

    > **Note:** Other utilities such as DCOLLECT might assume that VSAM datasets are either 100% Used or 100% Free. If you are comparing aggregate results between utilities (for example, with enterprise-wide space values) be aware that the OMEGAMON calculation is intended to give a more accurate picture of space usage within VSAM datasets.

- The SYS1.HASPACE dataset (also called the JES2 spool) is always assumed to be 100% Used (`Space Used=Space Allocated`).

Catalog Values

The following details pertain to catalog values in the product:

- The VTOC Index (`SYS1.VTOCIX.V`*nnnnnn*) and VVDS (`SYS1.VVDS.`*nnnnnn*) are not included in the Uncataloged Summary.
- Because the Dataset Attribute database is built starting at the Volume level, there might be entries in a Catalog that are not represented in the Catalog Detail display. The entries that are not displayed can include the following:
  - datasets that are in the catalog but do not physically exist on any volume
  - Tape datasets
  - VSAM Clusters that do not physically exist on disk
  - Datasets that are cataloged in more than one catalog

Use cases

The Dataset Attributes Database feature enables scenarios for resource management.

The scenarios, or use cases, are the following:

- Identify datasets in `Unmovable` status or with excessive CA/CI splits that might affect application performance
- Find large datasets that have a small `Blocksize` value, and reblock them to reduce elapsed time. The referencing application must support reblocking.
- Define a situation that raises an alert when the number of extents used by a catalog exceeds a specified value.
- Use the Tivoli® Enterprise Portal to immediately fix a dataset problem, as follows:

Using the Tivoli Enterprise Portal to manage datasets	
Phase	**Action item**
Phase One: Become aware of a problem and decide how to respond	1. Create a situation to monitor available free space in your storage groups. 2. When the situation fires, find it and take action.
Phase Two: Identify targets and take corrective action*	1. Access any row within any TotalStorage™ workspace. 2. Click on the Unused column to resort the table by most unused space. 3. Select the candidate datasets, right click, and issue the **Release Space** action from the **Dataset Actions** popup menu item. This action initiates a **DFDSS RELEASE** command for each selected dataset.

* The Dataset Attributes Database and Storage Toolkit features that you use in Phase Two replace procedures that require more effort:

- If you don't have the Dataset Attributes Database feature, look at a workspace for each volume in the storage group, one at a time. Also, compile and sort the list of datasets for each volume to determine which ones have most unallocated space.
- If you don't have the Storage Toolkit feature, issue DFDSS RELEASE for each dataset individually. For more information, see "Storage Toolkit" on page 143.

Configuring the Dataset Attribute Database feature

With the Dataset Attribute Collection Control dialog box, you can configure data collection for the Dataset Attributes Database feature.

The Dataset Attributes Database feature is turned off by default.

You can configure data collection for the Dataset Attributes Database feature by using the Dataset Attribute Collection Control dialog box.

> **Tip:** It is recommended that you first run collection on a subset of volumes to test how the feature affects performance. And then with the results of the first test collection, you can estimate the impact of collection on the address space of the Tivoli® Enterprise Monitoring Server and on the overall z/OS® system.
> Thus, you can choose optimal configuration settings, such as the volumes and storage groups to exclude from collection, what time of the day to start data collection, and what deadline (if any) to impose for collection completion.

To access the Dataset Attribute Collection Control dialog box, right-click the row of the Dataset Attribute Summary workspace and select **Dataset Attribute Collection Control**. The dialog box will open.

The **Dataset Attribute Collection Control dialog box** table describes the settings that you can configure.

Dataset Attribute Collection Control dialog box	
Section in the dialog box	**Comments**
Data Collection settings:	
• Start data collection at (*hh*:*mm*)	When you select this checkbox, most of the other fields in the dialog box become active for you to configure. *hh* is an integer from 1 to 12 and *mm* is an integer from 0 to 60.

Section in the dialog box	Comments
• Stop data collection at (*hh*:*mm*)	Use this setting to prevent data collection from continuing into a time period when critical jobs run that require significant amounts of system resources. *hh* is an integer from 1 to 12 and *mm* is an integer from 0 to 60. If you specify a stop time and collection is halted because that time is reached, attributes and reports for every dataset processed up to the point when collection was halted are available for viewing from the Tivoli® Enterprise Portal.
• Collect catalog data	When you clear this checkbox, no catalog information is obtained for any dataset. **Note:** By deselecting this option, you reduce resource consumption and elapsed time. However, several summary reports are empty when you do not collect this data, for example, CA/CI Split Summaries and Catalog/Uncataloged Summaries.
• Number of data rows to collect for Top *N* reports	The default is 20 rows. The range is 1 to 100 rows. Many of the reports for the Dataset Attribute Database feature involve Top *N* data. You can lower or raise this setting to adjust the amount of impact that this feature has on system resources versus the number of rows returned in each Summary report. For datasets that have the same key value as the smallest key in the table (for example, lowest number of allocated tracks in the Top *N* Space Allocated table), the first dataset encountered with that value is the one that is displayed in the report. Defaults for Top *N* data collection are as follows: • An inefficient block size is defined as `blksize < 128`. • A mature dataset is defined as `days since last reference > 366`.
Exclusions:	
• Volumes • Storage Groups	You can specify VOLSER or storage group names to exclude volumes from data collection. A volume that matches either exclusion specifications is excluded. You can specify up to 32 exclusions of each type. You can specify a partial name pattern and place the wildcard character (*) at the end. Volumes and Storage Groups whose names match the pattern are excluded.
Data Collection Status:	
• Status message • **Force Start** button or **Force Stop** button • **Apply Updates Now** button	Status messages include *Collection not active*, *Collection in progress*, *Collection waiting*, *Collection failed*, *Collection halted due to deadline*, and *Collection halted by user or operator*. Click **Force Start** to manually force initial Full collection or Partial Sweep collection to start. The **Force Stop** button is available while data collection is running to manually force data collection to stop. Results are displayed, including all datasets processed up to the point when the collection was stopped. Click **Apply Updates Now** to manually force Incremental collection to start.
Collection Interval:	

Section in the dialog box	Comments
• Apply Incremental updates every	You can specify the interval how often Incremental collection will be started to apply updates to the Dataset Attribute collection. The collection must be in sync and OMEGAMON for Storage Agent must be started on all LPARs in the PLEX. Available values: 15 minutes, 30 minutes, 1 hour, 3 hours, 6 hours

Optimizing data collection for the Dataset Attributes Database feature

Follow these guidelines to optimize the use of the Dataset Attributes Database feature.

- In a multi-system, shared hard-disk-drive environment, only enable Dataset Attribute collection on one of the connected systems, ideally a non-production system. This approach avoids collection of the same dataset space information on multiple systems, and places the work for collection on the most lightly loaded system.
- Use product-provided queries and workspaces rather than user-created, as they have been crafted and tuned for maximum performance.
  User-defined queries can incur significant overhead if a very large number of datasets needs to be examined to resolve the predicates.
- User-created situations written against the Dataset Detail table run against the complete dataset database. Do not set a Situation Sampling Interval smaller than the Dataset Attribute Database collection cycle, because attributes in the database are not updated more often than that.
- In IBM® Z OMEGAMON® AI for Storage® version 6.1.0, you can define a dataset attribute group to logically group a related set of datasets.
  When you craft a group, define the properties and attributes to limit the number of datasets that make up the group. If a large number of datasets in the Dataset Attribute Database match the criteria for your group, or a large number of datasets must to be examined to determine whether they match the criteria, this activity can require excessive processing, network and space resources.
- VSAM datasets are typically reported under their cluster name in the Dataset Attribute Database workspaces. Although the datasets contain a cluster component, a data component and an index component, these components are counted as a single dataset in the Dataset Attribute Database workspaces.
  In addition, only the cluster component name displays in the table views to represent the dataset. In counters displaying the "Total Datasets" in a group, a single VSAM dataset is counted only once.

  > **Note:** After a VSAM dataset has been migrated, the cluster, data and index components are treated as individual datasets by the Dataset Attribute Database workspaces. Each component appears on a separate line in table views. In counters displaying the "Total Datasets" in a group, each component is counted separately, increasing the dataset count (in contrast to the count prior to migration). This discrepancy occurs because the catalog entry for a VSAM dataset is changed after the dataset is migrated. The catalog entry change causes an inaccurate count. In fact, the Catalog Search Index reports these component names as type 'A' (Non-VSAM datasets).

Command-line option

If you have the authority to issue z/OS® operator commands, you can use the command-line interface to manage data collection for the Dataset Attributes Database feature.

To manage data collection for the Dataset Attributes Database feature, enter the following command and adjust the *temsname* and *command_function* according to the purpose:

```
F temsname,S3DA command_function
```

where *temsname* is the name of the Tivoli® Enterprise Monitoring Server started task and *command_function* is one of the following:

Command function	Purpose
START or COLLECT	Immediately starts a collection cycle. Previously defined volume and storage group exclusion lists are used, if any exist.
STATUS	Displays the current status of a Dataset Attributes Database collection cycle that is running.
SYSPLEX[,TYPE=X]	Displays the synchronization status between the current Agent and other Agents on various LPARs within the PLEX. TYPE=X also shows information about excluded systems.
HALT or STOP	Immediately stops a collection cycle. Attributes and reports for every dataset are processed up to the point when the STOP command was issued and are available for viewing from the Tivoli Enterprise Portal.
TERM	Terminates all processing for the Dataset Attributes Database feature.

You can find messages and output from any of these commands in the RKLVLOG output from the Tivoli Enterprise Monitoring Server.

Custom queries

When you create queries for the Dataset Attributes Database feature, consider the guidelines listed in this section.

- Ensure that the queries you create for the Dataset Details table (a query that involves the entire dataset attribute database) have at least one predicate specified.
  If you use only one attribute for selection, you cannot use the attribute in conjunction with STR or COUNT.
- Use uppercase characters for the character string values in your queries, such as dataset name or volume serial number. Otherwise, the data is not displayed in the Tivoli® Enterprise Portal.
- A single custom query can return a maximum of 16000 rows. This behavior prevents storage shortages and potential abends in the Tivoli® Enterprise Monitoring Server.
  A query (for example, DSORG=PO) might match a large number of dataset detail rows. This processing load can negatively affect all other IBM® Tivoli® Monitoring applications.
- When you specify the predicate for a custom query, position the most specific attribute as the first column, so that it functions as the primary index. The primary index limits the search for the additional attributes and, if specified with sufficient granularity, can greatly reduce response time.
  For example, if you want to see all datasets that are greater than 1000 cylinders on volume PRI000, specify the VOLUME attribute as primary index, so that the query initially searches for datasets from volume PRI000 only.

  If you specify the number of cylinders (1000) as the primary index, the query finds every dataset in the z/OS® environment that is larger than 1000 cylinders and then selects those that are located on PRI000. This query would be much more time consuming.
- The string 00:00:00 functions as a wildcard specification in queries that use Creation Date, Last Reference Date or Expiration date.
  Specify the Time value of such queries to be 00:00:00 because that value does not have a time associated with it and the returned converted times must not contradict the query specification.
- The following examples show the potential value of queries:
  - To see all components associated with a specific VSAM Cluster, enter the cluster name in the Associated Dataset attribute column.
  - Find VSAM datasets that have REPLICATE or IMBED.
  - Locate Uncataloged VSAM datasets.
  - Identify catalogs that have exceeded a threshold number of extents.

Working with dataset attribute groups

A dataset attribute group is a logical grouping of datasets that you can define and configure.

Datasets within a dataset attribute group are typically related according to specific criteria. You can define a dataset attribute group based on many of the attributes associated with datasets such as dataset name, volume, SMS class, Percent Free, and Last Backup Date Time.

After you define a dataset attribute group, you can use the Dataset Attributes Group Summary workspace to view aggregated metrics for the group. From this workspace, you can link to the Dataset Attributes Group Details to display metrics for datasets in the group.

To learn how to add and configure a dataset attribute group, see "Adding a dataset attribute group" on page 211.

Adding a dataset attribute group

You can add a dataset attribute group by using the Add User Dataset Group dialog box.

Group tab

Use this tab to specify your dataset attribute group name and description.

Name

Enter a name for the dataset attribute group that you are creating. Names may consist of any combination of characters and must not exceed 30 characters in length. Names are case sensitive.

Description

Enter an optional description for the dataset attribute group. Descriptions may consist of any combination of characters and must not exceed 50 characters in length.

Persist

Click this check box to indicate that the dataset attribute group should persist (or be saved) across recycles of the Tivoli Enterprise Monitoring Server (monitoring server or TEMS). Dataset attribute groups are discarded during a recycle of the monitoring server if **Persist** is not selected.

Disable Collection

Click this check box to indicate that the dataset attribute group collection should be disabled for the current group. Dataset attribute group collection will be skipped for those groups with **Disable Collection** selected.

Paste from clipboard

Click this button and paste copied group definition in a special format to fill the group's fields automatically.

Properties tab

Use this tab to specify properties for your dataset attribute group. Datasets that match any of the properties that you specify will be included in your group, unless further constrained by criteria in the Attributes tab.

Group Properties

Enter one or more dataset name, catalog name, and associated name to include those datasets in your dataset attribute group. Dataset names support the masking rules described in "Dataset name masking rules" on page 213.

SMS

Enter one or more SMS data class, storage class, management class, or storage group to include datasets belonging to those classes or groups in your dataset attribute group.

Owner

Enter one or more owner to include VSAM datasets with those owners in your dataset attribute group.

Volser

Enter one or more volume names to include datasets on those volumes in the dataset attribute group. You can specify a partial name pattern and place the wildcard character (*) at the end. For example, a

mask of PRI* will include datasets in your dataset attribute group that are on all volumes whose names begin with "PRI".

Attributes tab

Use this tab to specify attribute constraints for your dataset attribute group. Datasets that match any of the properties and all of the attributes that you specify will be included in your group.

Use the drop-down list in the Attribute column to select an attribute name. Use the drop-down list in the Operator column to select an operator. Use either the drop-down list in the Value column to select an enumerated value or enter a value in the field.

For more information, see "How properties and attributes are matched" on page 212.

Editing a dataset attribute group

Use this dialog box to edit a dataset attribute group.

For descriptions of the fields in the Group, Properties, and Attributes tabs see "Adding a dataset attribute group" on page 211.

> **Note:** The group name cannot be modified.

Deleting a dataset attribute group

In the Dataset Attributes System Summary workspace, you can delete a dataset attribute group.

To delete a dynamically defined dataset attribute group, complete the following steps:

1. Navigate to the Dataset Attributes System Summary workspace.
2. Right-click the Dataset Attributes Group Summary table's row for the group that you want to delete.
3. From the menu that appears, select **Delete Group**.

The selected dataset attribute group is deleted.

Additional dataset attribute group considerations

When you work with dataset attribute groups, be aware of the following additional requirements that apply.

> **Note:** See "How properties and attributes are matched" on page 212 for more information.

How properties and attributes are matched

Datasets that match any of the properties and all of the attributes that you specify will be included in your group.

For properties, logical OR is used. For attributes, logical AND is used. Logical AND is used between properties and attributes. As a result, the dataset attribute group definition looks like the following:

```
(Property1 OR Property2 OR ... OR PropertyN) AND (Attrbibute1 AND Attrbibute2 AND ... AND AttrbibuteM)
```

For example, if you want to overview large datasets that reference more than 180 days for specific storage groups, you could create a dataset attribute group with several storage groups defined in the **Properties** tab, and then define the **Allocated Tracks with Days Since Reference** attribute in the **Attributes** tab.

Dataset name masking rules

The dataset name property supports the use of mask characters. Asterisks (*) and percent signs (%) can be used to represent qualifiers or parts of qualifiers in dataset names.

"Mask characters for dataset names" on page 213 explains the use of asterisks and percent signs in dataset name masks.

Mask characters for dataset names

Mask Characters	Description
*	A single asterisk by itself indicates that either one qualifier or one or more characters within a qualifier can occupy that position. An asterisk can precede or follow a set of characters.
**	A double asterisk indicates that zero or more qualifiers can occupy that position. A double asterisk cannot precede or follow any characters; it must be preceded or followed by either a period or a blank.
%	A single percent sign by itself indicates that exactly one alphanumeric or national character can occupy that position.
%%	One to eight percent signs can be specified in each qualifier.

Dataset attribute groups and the Dataset Attribute Database

Dataset attribute groups access the information in the complete Dataset Attribute Database to determine the datasets that match the criteria specified in the groups.

When a Dataset Attribute Database collection cycle begins, the information associated with each group is cleared and is updated upon completion of the collection cycle.

Optimizing dataset attribute groups

You can optimize attribute order to improve processing speed, reduce dataset checks and resource usage.

In a large enterprise, the Dataset Attributes Database can be tracking millions of datasets.

When you craft a dataset attribute group, define the properties and attributes so that they limit the number of datasets that make up the group. If a large number of datasets in the Dataset Attribute Database match the criteria for your group, or a large number of datasets need to be examined to determine whether they match the criteria, this activity can require excessive processing, network, and space resources.

You can modify the order of the attributes to reduce cost of data collection for dataset attribute group.

By default, the first property (or the first attribute) is used as primary index for initial filtering. All other attributes are checked against the resulting datasets after the initial filtering.

For example, you have 10000 datasets. You specify two attributes: the first attribute filters 9000 datasets and the second attribute filters only 100. The second filter will be applied against the 9000 datasets from the first filter.

However, if you change the order of these two filters, you will have only 100 checks after the first filter, optimizing for faster results and less CPU consumption.

After any dataset attribute collection with any type, all groups are updated. Large amount of dataset attribute groups could cause high CPU consumption, and every incremental collection might run every 15, 30 or 60 minutes.

Dataset Attributes Group Details dataset limit

Although a dataset attribute group can aggregate the metrics for more than 5,000 datasets, the Dataset Attributes Group Details workspace returns the details for a maximum of 5,000 datasets.

> **Tip:** You can **Extract Attributes** in CSV format without the 5,000 limitation to add datasets that exceed the limitation from Dataset Attributes Group Details.

Dataset attribute groups VSAM criteria

If a dataset attribute group includes an attribute associated exclusively with VSAM datasets, then the dataset attribute group contains only VSAM datasets.

For example, if you specify a value for CI_Size (control interval size), then the dataset attribute group consists of VSAM datasets that meet this and all other attribute criteria specified for the group.

The following attributes are associated exclusively with VSAM datasets:

- Associated_Entry_Name
- Average_LRECL
- CI_Size
- Data_Component
- Erase
- Imbed
- Maximum_Record_Size
- Maximum_Buffer_Size
- Compress_User_Data_Size
- Index_Component
- Speed
- Recovery
- Reuse
- Replicate
- Spanned
- Page_Space
- Swap_Space
- Percent_Free_CIs_per_CA
- Percent_Bytes_Free_in_CI
- Share_Option
- Owner
- CIs_per_CA
- Free_CIs_per_CA
- Free_Bytes_per_CI
- Logical_Records
- Deleted_Records
- Inserted_Records
- Updated_Records
- Retrieved_Records
- Component_Free_Space
- Number_of_CI_Splits

- Number_of_CA_Splits
- Number_of_EXCPs
- High_Allocated_RBA
- High_Used_RBA
- Unique
- VSAM_Dataset_Organization

Application Launch and TotalStorage Productivity Center

Learn about an Application Launch feature provided by IBM® Tivoli® Monitoring and how OMEGAMON® for Storage uses this feature in selected workspaces to provide launching of the IBM TotalStorage Productivity Center.

IBM TotalStorage Productivity Center

The IBM TotalStorage Productivity Center is a software tool that helps you manage complex, heterogeneous open storage environments to increase storage capacity utilization and to improve administrative efficiency. Typically, this software is integrated with a disk subsystem, such as a DS8000 storage server.

OMEGAMON for Storage launch feature

This feature provides convenient access to the IBM TotalStorage Productivity Center. The feature enables the following example scenario:

Phase	Stage
Phase One: Become aware of a problem and decide how to respond.	• You suspect that a DS8000 storage facility has slow response time. • You decide to investigate hardware issues in IBM TotalStorage Productivity Center.
Phase Two: Identify targets and take appropriate corrective action. **Note:** If you did not have the launch feature, you must complete the following manual process: 1. Log in to the mainframe computer 2. Find the affected resource. 3. Find the HTML page for IBM TotalStorage Productivity Center and launch it.	• You right-click the workspace to access the context menu. • You select **Launch** and immediately access the sign-on screen for the instance of IBM TotalStorage Productivity Center for the storage resource that you want to investigate. • You discover a need to allocate space to the extent pool that generated the alert. • You allocate some unallocated storage space to the extent pool.

Launching the IBM TotalStorage Productivity Center

Learn how to launch the IBM TotalStorage Productivity Center.

PREREQUISITE

The IBM TotalStorage Productivity Center is a browser-based software environment. In OMEGAMON® for Storage, use the URL of the browser page to make a connection.

Run the **KS3LaunchConfig** command to configure the launch feature of OMEGAMON® for Storage. The command identifies the location of the browser software and the URL of the IBM TotalStorage Productivity Center.

For more information about configuring the feature, see “Configuring the launch feature” on page 217.

You can use the features of the IBM TotalStorage Productivity Center to manage the hardware resource that you select in a workspace.

1. Right-click a resource that is displayed in a workspace. See “Workspaces that provide the launch feature” on page 217.

2. Select **Launch** from the menu.
   The Create or Edit Launch Definitions dialog box is displayed.
3. Select the instance of the IBM TotalStorage Productivity Center that you want to launch.
4. Click **Launch**.
   The sign-on screen for the IBM TotalStorage Productivity Center is displayed in a Web browser page.

Workspaces that provide the launch feature

The following OMEGAMON® for Storage workspaces provide the launch feature.

You can access these linked (secondary) workspaces through links in either of two primary workspaces: Cache CU Performance workspace or Cache CU Status workspace. The secondary workspace links are available when you have selected a DS8000 resource in a workspace view.

- TotalStorage Configuration
  - TotalStorage Ranks
    - TotalStorage Array Configuration
    - TotalStorage Rank Trend
  - TotalStorage Extent Pool Volumes
    - TotalStorage Volume Trend
  - TotalStorage Extent Pool Trend
    - TotalStorage Rank History
    - TotalStorage Volume History
- TotalStorage CU Volumes
  - TotalStorage Volume Trend

Configuring the launch feature

You use the command-line interface to maintain the launch points (insert, update, and remove operations). The available launch points are predefined by OMEGAMON® for Storage.

You can configure the ability to launch IBM TotalStorage Productivity Center in OMEGAMON® for Storage.

1. Access a command-line interface.
2. Navigate to the directory where the command is stored. By default, the command is stored in the following path:
   `C:\IBM\ITM\CNP\`
3. Run the KS3LaunchConfig command to insert, change, or remove target points for the launch feature.

   Note: Before you run the KS3LaunchConfig command you must navigate in the command line interface to the directory where the command is stored. By default the command is stored in the following path:
   `C:\IBM\ITM\CNP\`

   - To add a new launch target use the following flags:
     KS3LaunchConfig -insert -loc=*LOCATION* -url=*URL*

     where:
     - *LOCATION* is the local path of the browser application that you use to access the IBM TotalStorage Productivity Center. For example, the entry for Microsoft Internet Explorer might be "`C:\Program Files\Internet Explorer\iexplore.exe`".
     - *URL* is the URL of the browser page for IBM TotalStorage Productivity Center.

> **Note:** If the **LOCATION** parameter or the **URL** parameter contains embedded blank spaces, you must enclose the parameter in quotation marks.

- To make a change or add a launch target use the following flags:
  KS3LaunchConfig -change -loc=*LOCATION* -url=*URL*

  where:

  - *LOCATION* is the local path of the browser application that you use to access the IBM TotalStorage Productivity Center. For example, the entry Microsoft Internet Explorer might be "`C:\Program Files\Internet Explorer\iexplore.exe`".
  - *URL* is the URL of the browser page for IBM TotalStorage Productivity Center.

    > **Note:** If the **LOCATION** parameter or the **URL** parameter contains embedded blank spaces, you must enclose the parameter in quotation marks.

- To remove all of the launch target points that were created by this command use the following flags:
  KS3LaunchConfig -remove

The following command illustrates the use of the command:

> **Note:** Before you run the KS3LaunchConfig command you must navigate in the command line interface to the directory where the command is stored. By default the command is stored in the following path:
> `C:\IBM\ITM\CNP\`

```
KS3LaunchConfig -insert -loc="c:\Program Files\Internet Explorer\iexplore.exe"
-url=http://your.TPC.url
```

where: *your.TPC.url* represents the URL of the browser page for IBM TotalStorage Productivity Center.

DFSMSrmm Status feature

OMEGAMON® for Storage provides a DFSMSrmm monitoring and analysis function to help you manage the critical DFSMSrmm subsystem.

The DFSMSrmm Status workspace provides a high-level overview of DFSMSrmm statistics and links to related workspaces that track data for common aspects of DFSMSrmm management and administration.

For related information, see:

- "RMM Collection Control dialog box" on page 223.
  You use this dialog box to configure the data collection settings for the DFSMSrmm feature.
- "Monitoring DFSMSrmm status" on page 86.
  Learn about the workspaces that monitor the status of DFSMSrmm resources. You access these workspaces through the Tape Management Status node in the Tivoli Enterprise Portal.

Benefits

You can learn about the benefits of the DFSMSrmm Status feature.

The DFSMSrmm Status feature analyzes and reports on entries in the DFSMSrmm Control Dataset (CDS) regarding all datasets on all tape volumes that DFSMSrmm is managing. With this feature, you can perform the following operations:

- View comprehensive summary reports regarding general trends or specific low-level details:
  - Locations Owners
  - Accounting codes
  - Creating Jobs or Programs
  - Vital Record Specifications (VRS)
  - Volumes and Datasets
- Identify exception conditions regarding DFSMSrmm CDS space and backup, scratch tape counts, configuration state, and so on.
- Select tape volumes, datasets, owners, or VRSs from a workspace view and initiate Storage Toolkit action commands against them.
- Obtain an installation-wide view of DFSMSrmm CDS space utilization, exceptional conditions, and summary statistics.
- Identify resources that require attention, such as Pending Action or Pending Move volumes, and initiate that action by using Storage Toolkit.
- Avoid the use of the standard DFSMSrmm reporting utility program that requires user customization.

Optimization

Learn how to optimize performance of the DFSMSrmm Status feature.

You should configure the DFSMSrmm Status feature to optimize performance. Otherwise, the amount of data that you collect can use excessive system resources.

For example, you can optimize configuration of the following behaviors:

- Collection interval
- Inclusion of dataset-level attribute data

Also see "RMM Collection Control dialog box" on page 223 and also "Optimizing data collection for the DFSMSrmm Status feature" on page 221.

Note: Data collected for the DFSMSrmm Status feature is maintained in an in-memory data space for high-performance access. The data is not saved to disk, except for those Summary tables that have been configured for historical collection.
For this reason, after you restart the Tivoli Enterprise Monitoring Server address space, the data space must be rebuilt by initiation of a collection cycle, either manually or at a scheduled time.

Attribute derivation and definition

You find clarifications of the concepts and values associated with specific attributes of the DFSMSrmm Status feature.

- Only volumes with a type of PHYSICAL and a status of MASTER or USER are included when calculating the following statistics:
  - Capacity
  - Utilization
  - Compression
- Compression, which is meant to represent the percentage reduction of data written from the host that is actually recorded on the tape volume, is calculated as follows:
  `1-((Volume Capacity*Volume % Full)/(App Data written to volume))`
- The following attributes include PHYSICAL, LOGICAL, and STACKED volume types:
  - Total volumes
  - Total datasets
  - Average volume use count
  - Average datasets per volume
  - Percent of volumes read/written previous day
- The following dates are shown as NONEXPIRING in the Days To/Days From attributes that are calculated by OMEGAMON® :
  - 99000
  - 99365
  - 99366
  - 98000
  - 98365
- The Days Until Expiration attribute displays a positive value to represent a date in the future, while a negative the number represents the days since expiration. Similarly, a positive **Days for Retention** value signifies the remaining number of days that the dataset or volume is retained, while a negative value indicates the number of days since the retention date.
- History can be enabled for the **DFSMSrmm Summary**, **Configuration**, and **CDS** tables only. CDS and Configuration data is collected at the interval specified in the History Collection Configuration dialog box of the Tivoli® Enterprise Portal. Summary attributes are written to the history files only once during any collection cycle, because the summary data does not change until the next collection cycle is run.
  - Locations Owners
  - Accounting codes
  - Creating Jobs or Programs
  - Vital Record Specifications (VRS)
  - Volumes and Datasets

Scenarios for using the DFSMSrmm feature

The DFSMSrmm feature enables scenarios for powerful tape environment monitoring and management.

Examples of such scenarios are the following:

- Using the existing OMEGAMON® Tape Devices workspace (under the navigation tree Tape Group node), you can link to the new DFSMSrmm Volume Details workspace to see all DFSMSrmm attributes for a mounted and in-use tape volume.
- When an individual leaves the organization, you can use the Owner Summary workspace and link to the Owner Realtime Volumes workspace to see all volumes owned by that individual. Using the Storage Toolkit Change Volume command, you can change the owner name to a current employee.

Using the Tivoli Enterprise Portal to manage DFSMSrmm tape volumes and datasets	
Phase	**Action item**
One: Become aware of a problem and decide how to respond.	Create a copy of the product-provided situation KS3_RMM_Scratch_Tape_Warning and set the low and high scratch volume counts to values appropriate for your installation. Start the situation.
Two: Identify targets and take corrective action.	When the situation is triggered, go to the DFSMSrmm Pending Actions workspace and use the Storage Toolkit Confirm Actions for all Volumes command for any volumes that are in pending scratch state. You could also use the Storage Toolkit Add Volume command to add new volumes to DFSMSrmm.

Optimizing data collection for the DFSMSrmm Status feature

Learn how to optimize data collection for the DFSMSrmm Status feature.

Follow these guidelines to optimize the usage of the DFSMSrmm Status feature:

- In a multisystem, shared hard-disk-drive environment where the DFSMSrmm subsystem is running on multiple systems and sharing a single CDS, only enable OMEGAMON® DFSMSrmm collection on one of the connected systems, ideally a non-production system.
  This behavior avoids collection of the same DFSMSrmm information on multiple systems and places the work for collection on the system that is most lightly loaded.
- Use product-provided queries and workspaces rather than user-created because they have been crafted and tuned for maximum performance.
  User-defined queries can incur significant overhead when a very large number of datasets or volumes must to be examined to resolve the predicates.
- User-created situations written against the **DFSMSrmm Summary** table should not have their **Sampling Interval** smaller than the DFSMSrmm collection cycle because attributes are not updated more often than that.

Command-line option

If you have the authority to issue z/OS® operator commands, you can use the command-line interface to manage data collection for the DFSMSrmm Status feature.

To manage data collection for the DFSMSrmm Status feature, use the following command:

F *temsname*,S3RM *command_function*

where *temsname* is the name of the Tivoli® Enterprise Monitoring Server started task and *command_function* is one of the following:

- START or COLLECT: Immediately start a collection cycle.
- STATUS: Display the current status of a running DFSMSrmm collection cycle.

- HALT or STOP: Immediately stop a collection cycle. Attributes and reports for every volume and dataset are processed up to the point that the STOP command was issued are available for viewing from the Tivoli® Enterprise Portal.
- TERM: Terminate all processing for the DFSMSrmm Status feature.

You can find messages and output from any of these commands in the RKLVLOG output from the Tivoli® Enterprise Monitoring Server.

Custom queries

You can use the following tips and suggestions when creating queries for the DFSMSrmm Status feature.

You must use uppercase characters for the character string values in your queries, such as dataset name or volume serial number; otherwise, the data is not displayed in the Tivoli® Enterprise Portal.

A single DFSMSrmm table query can return a maximum of 4000 rows. This behavior prevents storage shortages and potential abends in the Tivoli® Enterprise Monitoring Server.

A query (for example, using the predicate JOB=$*) might match an extremely large number of dataset detail rows. This processing load can negatively affect all other IBM® Tivoli® Monitoring applications.

Dialog boxes

Dialog boxes

This section includes dialog boxes that allow you to configure data collection settings for various features, such as DFSMSrmm Status, CA1, Copy Services, Record Level Sharing, and ADABAS, as well as to set data collection intervals, define groups of datasets, and copy the group.

RMM Collection Control dialog box

You use this dialog box to configure the data collection settings for the DFSMSrmm feature.

The DFSMSrmm Status feature is turned off by default. It is recommended that you first run a collection cycle off-hours to test how the feature may affect performance. Use the information from this test session to estimate the impact of collection on the address space of the Tivoli® Enterprise Monitoring Server and on the overall z/OS system. This estimate can help you choose optimal configuration settings, such as what time of day to start collection of data, what deadline (if any) to impose for completion of collection, and whether to include dataset-level attributes in the collection.

Details regarding the Collection Control dialog box

Use the DFSMSrmm Collection Control dialog box to configure data collection for the DFSMSrmm feature. Right-click in a row of the DFSMSrmm Status workspace and select **RMM Collection Control** to launch the dialog box.

RMM Collection Control dialog box	
Section/option	**Comments**
Collection Scheduling	
Start collection at (*hh*:*mm*)	Select this check box to activate most of the other fields in the dialog box. *hh* is an integer from 1 to 12, and *mm* is an integer from 0 to 60. By default, collection is set to run daily. You deselect specific days of the week in the **Every** area to prevent collection on specific days. **Start collection at** causes the backend controller to schedule Summary Collection to begin at the specified time on the specified days of the week. If unchecked, Summary collection does not automatically run.
Stop collection at (*hh*:*mm*)	Use this setting to prevent data collection from continuing into a time period when critical jobs run that require significant amounts of system resources. You can also use this setting to prevent collection when the DFSMSrmm subsystem itself is in heavy use and requires exclusive access to its CDS. *hh* is an integer from 1 to 12, and *mm* is an integer from 0 to 60. If you specify a stop time and collection is halted because that time is reached, attributes and reports for every volume and dataset processed up to the point that collection was halted are available for viewing from the Tivoli® Enterprise Portal. **Stop collection at** causes the backend controller to halt the RMM Summary Collection process at the indicated time of day. This response can be useful in environments where the RMM CDS is extremely large and the summary process interferes with normal RMM processing, or where a limit should be placed on the amount of elapsed time/resources consumed by the summary collection process.
Collection Options	

Section/option	Comments
Summarize RMM dataset attributes	When you deselect this check box, no DFSMSrmm information is obtained for any tape dataset. You continue to see a total dataset count, however, because this metric is obtained by accumulating the total number of datasets on each processed volume. **Note:** By deselecting this option, you reduce resource consumption and elapsed time. However, several summary report workspaces are empty when you do not collect this data, specifically, Job, Program, and VRS summary. In addition, several attributes that appear on various summaries can also not be calculated: • Datasets created previous day • Data created previous day • Datasets read/written previous day
Collection Status	
Status message: **Force Start** button or **Force Stop** button	Status messages include, `Collection not active`, `Collection in progress`, `Collection waiting`, `Collection failed`, `Collection completed`, `Collection halted due to deadline`, and `Collection halted by user or operator`. Click **Force Start** to manually force data collection to start. A **Force Stop** button is available while data collection is running to manually force data collection to stop. Results are displayed that include all volumes and datasets processed up to the point that collection was halted. And the status shown is `Collection completed`. If you stop the collection by clicking on **Force Stop** button, the state in the Collection Status area is `Collection halted by user or operator`. The status is not `Stopped`.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

CA 1 Collection Control dialog box

You use this dialog box to configure the data collection settings for the CA 1 feature.

Note: The requirements specified in this topic only apply if you want to monitor the CA 1 Tape Management System.

Account for the following issues when you install IBM® Z OMEGAMON® AI for Storage® to ensure that the CA 1 feature operates correctly:

- You must delete, reallocate, and initialize the RKS3DSAx persistent data stores. Otherwise, the new CA 1 tables cannot be correctly defined for historical collection.

The CA 1 Status feature is turned off by default. It is recommended that you first run a collection cycle off-hours to test how the feature may affect performance. Use the information from this test session to estimate the impact of collection on the address space of the Tivoli® Enterprise Monitoring Server and on the overall z/OS system. This estimate can help you choose optimal configuration settings, such as what time of day to start collection of data and what deadline (if any) to impose for completion of collection. Please refer to the next section for sizing and provisioning requirements.

Sizing and provisioning requirements

The collector module will create the following VSAM datasets during a collection cycle:

- A primary data cluster with 3 alternate indices. This is a non-transient cluster which holds the output from the collection cycle.

- A secondary data cluster with 3 alternate indices. This is an optional copy of the primary cluster and is used by reporting modules to provide data while a collection cycle is in progress. If you disable "allow reporting during collection" then this cluster will not be created.
- A scratch data cluster. This is a transient cluster which only exists during the collection cycle.

Sizing

Use the following formulas to determine the required primary and secondary allocations of each VSAM cluster or index.

TMC values

- V (volume): total volume count (active + scratch, not including deleted volumes)
- D (dataset): active volume count + active DSNB count
- T (total objects): V + D + (D * 0.10) + (V * 0.01)

VSAM values

- VP (primary allocation in MB): ((T * 612) + (T * 612 * 0.10)) / (1024 * 1024)
- VS (secondary allocation in MB): VP * 0.20
- VT (total allocation in MB): VP + VS

Total space

If "allow reporting during collection" is enabled:

- Total Primary Allocation in MB: VP * 9
- Total Secondary Allocation in MB: VS * 9
- Total Allocation in MB: VT * 9

(1 primary data cluster with 3 alternate indices, 1 secondary data cluster and 3 alternate indices and 1 scratch cluster)
If "allow reporting during collection" is disabled:

- Total Primary Allocation in MB: VP * 5
- Total Secondary Allocation in MB: VS * 5
- Total Allocation in MB: VT * 5

(1 primary data cluster with 3 alternate indices and 1 scratch cluster)

SMS data class

An SMS data class must be provided to govern the VSAM datasets maintained by the collector. The SMS data class should have the following attributes:

- Extended Addressability is required for VSAM datasets greater than 4GB in size.
- Multi-volume is recommended.

Details regarding the Collection Control dialog box

Use the CA 1 Collection Control dialog box to configure data collection for the CA 1 feature. Right-click in a row of the CA 1 Status workspace and select **CA 1 Collection Control** to launch the dialog box.

CA 1 Collection Control dialog box	
Section/option	**Comments**
Collection Scheduling	

Section/option	Comments
Start collection at (*hh*:*mm*)	Select this check box to activate most of the other fields in the dialog box. *hh* is an integer from 1 to 12, and *mm* is an integer from 0 to 60. By default, collection is set to run daily. You deselect specific days of the week in the **Every** area to prevent collection on specific days. **Start collection at** causes the backend controller to schedule Summary Collection to begin at the specified time on the specified days of the week. If unchecked, Summary collection does not automatically run.
Stop collection at (*hh*:*mm*)	Use this setting to prevent data collection from continuing into a time period when critical jobs run that require significant amounts of system resources. You can also use this setting to prevent collection when the CA 1 subsystem itself is in heavy use and requires exclusive access to its TMC. *hh* is an integer from 1 to 12, and *mm* is an integer from 0 to 60. If you specify a stop time and collection is halted because that time is reached, attributes and reports for every volume and dataset processed up to the point that collection was halted are available for viewing from the Tivoli® Enterprise Portal. **Stop collection at** causes the backend controller to halt the CA 1 Summary Collection process at the indicated time of day. This response can be useful in environments where the CA 1 TMS is extremely large and the summary process interferes with normal CA 1 processing, or where a limit should be placed on the amount of elapsed time/resources consumed by the summary collection process.
Collection Options	
Allow reporting during collection	When you select this check box, a second VSAM cluster will remain in existence so that the user may request data during a collection cycle. Data from the previous collection cycle will be shown. By deselecting this check box, only a single VSAM cluster will be maintained and access to data will not be possible during a collection cycle. Please refer to the previous Sizing and provisioning requirements section for details on how this will affect your DASD requirements.
SMS Data Class	An explicit SMS data class must be provided to the govern the VSAM datasets used by the collector module. Please refer to the previous Sizing and provisioning requirements section for details on how this will affect your DASD requirements.
Collection Status	
Status message: **Force Start** button or **Force Stop** button	Status messages include, `Collection not active`, `Collection in progress`, `Collection waiting`, `Collection failed`, `Collection completed`, `Collection halted due to deadline`, and `Collection halted by user or operator`. Click **Force Start** to manually force data collection to start. A **Force Stop** button is available while data collection is running to manually force data collection to stop. Results are displayed that include all volumes and datasets processed up to the point that collection was halted. And the status shown is `Collection completed`. If you stop the collection by clicking on **Force Stop** button, the state in the Collection Status area is `Collection halted by user or operator`. The status is not `Stopped`.

Copy Group dialog box

Use this dialog box to copy the group.

To copy a group:

1. Right-click groups (either Dataset Attributes or User DASD), and then click **Copy Group**.
2. In the **Copy Group** dialog box, use the Available Nodes table to select the destination of copying, or use the Copy button to copy group definition in the clipboard.
3. When you finish, click **OK** to copy.

The **Copy Group** dialog box has the following fields:

Available Nodes

The list of copying destination storage nodes.

Copy button

The button to copy group definition in the clipboard.

> **Note:** This button is disabled if several groups are selected for copying.

Copy Services Collection Control dialog box

You use this dialog box to configure the data collection settings for the Copy Services feature.

The Copy Services Status feature is turned off by default. It is recommended that you first run a collection cycle off-hours to test how the feature may affect performance. Use the information from this test session to estimate the impact of collection on the address space of the Tivoli® Enterprise Monitoring Server and on the overall z/OS system. This estimate can help you choose optimal configuration settings, such as what time of day to start collection of data, what deadline (if any) to impose for completion of collection, and whether to include dataset-level attributes in the collection.

Details regarding the Collection Control dialog box

Use the Copy Services Collection Control dialog box to configure data collection for the Copy Services feature. Right-click in a row of the Copy Services Status workspace and select **Copy Services Collection Control** to launch the dialog box.

Copy Services Collection Control dialog box	
Section/option	**Comments**
Collection Scheduling	
Start collection at (*hh*:*mm*)	Select this check box to activate most of the other fields in the dialog box. *hh* is an integer from 1 to 12, and *mm* is an integer from 0 to 60. By default, collection is set to run daily. You deselect specific days of the week in the **Every** area to prevent collection on specific days. **Start collection at** causes the backend controller to schedule Summary Collection to begin at the specified time on the specified days of the week. If unchecked, Summary collection does not automatically run.
Stop collection at (*hh*:*mm*)	Use this setting to prevent data collection from continuing into a time period when critical jobs run that require significant amounts of system resources. You can also use this setting to prevent collection when the Copy Services subsystem itself is in heavy use and requires exclusive access to its CDS. *hh* is an integer from 1 to 12, and *mm* is an integer from 0 to 60. If you specify a stop time and collection is halted because that time is reached, attributes and reports for every volume and dataset processed up to the point that collection was halted are available for viewing from the Tivoli® Enterprise Portal. **Stop collection at** causes the backend controller to halt the Copy Services Summary Collection process at the indicated time of day. This response can be useful in environments where the Copy Services CDS is extremely large and the summary process interferes with normal Copy Services processing, or where a limit should be placed on the amount of elapsed time/resources consumed by the summary collection process.
Collection Options	

Section/option	Comments
Summarize Copy Services dataset attributes	When you deselect this check box, no Copy Services information is obtained for any tape dataset. You continue to see a total dataset count, however, because this metric is obtained by accumulating the total number of datasets on each processed volume. **Note:** By deselecting this option, you reduce resource consumption and elapsed time. However, several summary report workspaces are empty when you do not collect this data, specifically, Job and Program summary.
Collection Status	
Status message: **Force Start** button or **Force Stop** button	Status messages include, `Collection not active`, `Collection in progress`, `Collection waiting`, `Collection failed`, `Collection completed`, `Collection halted due to deadline`, and `Collection halted by user or operator`. Click **Force Start** to manually force data collection to start. A **Force Stop** button is available while data collection is running to manually force data collection to stop. Results are displayed that include all volumes and datasets processed up to the point that collection was halted. And the status shown is `Collection completed`. If you stop the collection by clicking on **Force Stop** button, the state in the Collection Status area is `Collection halted by user or operator`. The status is not `Stopped`.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

RLS Collection Control dialog box

You use this dialog box to configure the data collection settings for the Record Level Sharing feature.

See "DFSMSrmm Status feature" on page 219 for additional information.

Details regarding the Collection Control dialog box

Use the RLS Collection Control dialog box to configure data collection for the Record Level Sharing feature. Right-click in a row of the Record Level Sharing Status workspace and select **RLS Collection Control** to launch the dialog box.

RLS Collection Control dialog box	
Section/option	**Comments**
RLS Collection Interval	
RMF Interval	Select this option to use the RMF Interval as the RLS collection interval
On Demand	Select this option to collect RLS data on demand.
Mins (5 - 360)	Select this option to set a time interval for RLS collection. Specify a time in minutes between 5 and 360.

To learn about the standard capabilities of a dialog box, see "Standard tabs in dialog boxes of Storage Toolkit" on page 157. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

ADABAS Collection Control dialog box

You use this dialog box to configure the data collection settings for the ADABAS feature.

Details regarding the ADABAS Collection Control dialog box

Use the ADABAS Collection Control dialog box to configure data collection for the ADABAS feature. To launch the dialog box, right-click on a row of the ADABAS Summary workspace and select **Collection Interval** .

If you change the interval to a new value that is less than time than the previous collection, the new collection will be started. Otherwise, the next collection will be started after the last collection at the specified time.

> **Warning:** Spool issues may occur when a low collection interval is specified.

ADABAS Collection Control dialog box	
Section/option	**Comments**
ADABAS Collection Interval	
Mins (0 - 2880)	Select this option to set a time interval between data collections. Specify a time in minutes between 0 and 2880. 0 - disable ADABAS collections.

Dataset Group Properties dialog box

You use this dialog box to define a group of datasets based on filtering rules (masks) and related properties.

You access this dialog box through the Dataset Group Summary workspace, which is described in "Monitoring dataset groups summary data" on page 118.

The settings to modify are as follows:

- **Enable Detail History Collection**. Select this checkbox to enable the collection of detailed historical data for each dataset member of the dataset group. (Only applicable to v420 IF4 and above.)
- **RLS Group**. Select this checkbox to enable collection of RLS data for each member of the group.
- **Mask**. You type a mask definition following the rules used by DFDSS dataset name masking.
  For detailed information, see the IBM® Z OMEGAMON® AI for Storage®: User's Guide.
- **Active**. Select this checkbox to cause data collection from the catalog to be performed for the datasets that fall under a mask.
- **Space Data**. Select this checkbox to collect space data for the datasets that fall under a mask.
- **RLS Mask**. Select this checkbox to collect RLS data for the datasets that fall under this dataset mask.

Dataset Group Data Collection dialog box

You use this dialog box to set the interval at which data collection occurs.

You access this dialog box through the Dataset Group Summary workspace. The settings to modify are as follows:

- **RMF interval**. Data is collected at the predefined time interval of the resource measurement facility (RMF).
- **On Demand**. Data is collected when the user requests collection.
- ***nnn* Mins (5 - 360)**. Data is collected at the interval of minutes that you define, where *nnn* is an integer from 5 to 360.

Attributes

In this section, you can find information about the attributes that you can use to build situations for monitoring the performance of the required managed systems.

Tivoli® Enterprise Portal gathers data from remote agents located within your systems. The portal stores this data for IBM® Z OMEGAMON® AI for Storage® in elements called attributes. You can use attributes to build situations to monitor the performance of the managed systems you are concerned with.

Attributes correspond to the column names contained on the table views in associated IBM® Z OMEGAMON® AI for Storage® workspaces. Related attributes are grouped into attribute groups. A given table view contains information provided by a single attribute group.

To use attributes effectively, you must first understand the structure of an attribute.

- An attribute is made up of an attribute group and an attribute item.
- An attribute group associates attributes that refer to some unique characteristic of the data that is being stored.
- An attribute item stores data for a particular property of an attribute group.
  For example, in the attribute TOTALSTORAGE_RANKS.Total_Bytes_Read
- Total Bytes Read is an attribute item that stores a value representing the total number of bytes read from the rank over the interval in the following attribute group.

You use attributes to create situations that monitor the state of your storage environment. A situation describes a condition you want to test. When you start a situation, the monitoring server compares the values that you have assigned for the situation's attributes with the values collected by IBM® Z OMEGAMON® AI for Storage® and registers a situation event if the condition is met. You are alerted to situation events by indicator icons that appear in the Navigator.

Each attribute is a characteristic of an object. For example, the All Channels Busy Percent attribute in the Logical Control Unit attribute group displays the percentage of the time that all channel paths to an LCU were busy.

Attribute groups

The IBM® Z OMEGAMON® AI for Storage® attributes are organized into groups of related items. These groups of attributes comprise the attribute tables for the product.

> **Note:** In the Tivoli® Enterprise Portal, the names of attribute groups are preceded by “S3,” which is the product code for IBM® Z OMEGAMON® AI for Storage®.

Attributes and workspaces

You can view the data reported for the attribute in the reports and charts displayed in the IBM® Z OMEGAMON® AI for Storage® workspaces. The workspaces contain the table views you use to obtain information about the various aspects of your environment. The table views and graphs within each workspace report the storage attribute information you are monitoring. You can use the collected data to analyze and monitor performance.

Attributes and situations

Attributes are used in the predefined situations for the product. You can also create custom situations for to monitor the performance of your z/OS® systems. These situations can monitor one aspect of your enterprise or analyze multiple conditions to alert you to problems that can occur when attribute values exceed their thresholds.

List of attribute groups and attributes

The following attribute groups are provided with the IBM® Z OMEGAMON® AI for Storage® product:

“ADABAS Databases” on page 233
“ADABAS Files” on page 234
“ADABAS Volumes” on page 236
“Application Datasets” on page 236
“Application Dataset Details” on page 236

“Application Dataset Monitor” on page 237
“Application Monitoring” on page 237
“Application RLS Datasets” on page 238
“Application Volume Performance” on page 238
“Application Volume Space” on page 240
“Cache Control Unit” on page 241
“Cache Devices” on page 242
“Cache Raid Rank” on page 245
“Catalog Dataset Attributes” on page 245
“CA1 Configuration” on page 246
“CA1 Controls” on page 246
“CA1 Dataset” on page 247
“CA1 Job Summary” on page 248
“CA1 Permanent Error Summary” on page 248
“CA1 Program Summary” on page 248
“CA1 Summary” on page 249
“CA1 Temporary Error Summary” on page 250
“CA1 Volume” on page 250
“Channel Path” on page 253
“Cross System Volume Summary” on page 255
“Cross System Volumes” on page 256
“DASD Group Vol Perform” on page 257
“DASD Group Vol Space” on page 258
“DASD Physical Group” on page 260
“DASD Volume Performance” on page 262
“DASD Volume Space” on page 264
“Dataset Attributes Database” on page 204
“Dataset Attributes Blocksize Summary” on page 265
“Dataset Attributes CA Split Summary” on page 266
“Dataset Attributes Catalog Summary” on page 266
“Dataset Attributes CI Split Summary” on page 267
“Dataset Attributes Creation Date Summary” on page 267
“Dataset Attributes Defs” on page 268
“Dataset Attributes Dataset Detail” on page 268
“Dataset Attributes DSORG Detail” on page 272
“Dataset Attributes DSORG Summary” on page 273
“Dataset Attributes Extent Summary” on page 274
“Dataset Attributes High Level Qualifier Summary” on page 274
“Dataset Attributes Inefficient Blocksize Summary” on page 275
“Dataset Attributes Mature Dataset Summary” on page 276
“Dataset Attributes Never Referenced Dataset Summary” on page 276
“Dataset Attributes Newborn Dataset Summary” on page 277
“Dataset Attributes Reference Date Summary” on page 277
“Dataset Attributes SMS Class Detail” on page 278
“Dataset Attributes SMS Class Summary” on page 279
“Dataset Attributes Space Allocated Summary” on page 274
“Dataset Attributes Space Unused Summary” on page 280
“Dataset Attributes System Summary” on page 280
“Dataset Attributes Uncataloged Detail” on page 282
“Dataset Attributes Uncataloged Summary” on page 283
“DSN Attr Group Detail” on page 283
“DSN Attr Group Summary” on page 286
“Dataset Group Defs” on page 288
“Dataset Group Details” on page 288
“Dataset Group Detail Attrs” on page 291
“Dataset Group Details Volume” on page 291
“Dataset Group Summary” on page 293
“Dataset Performance” on page 294
“Dataset SMS” on page 295
“Dataset Space” on page 295
“Dataset Users” on page 297
“FICON Director Control Units” on page 297
“FICON Director Ports” on page 297
“FICON Director Summary” on page 298
“S3 App Defs” on page 298
“S3 Dag Defs” on page 299
“S3 DASD Defs” on page 299
“S3 Group Defs” on page 300
S3 Device Class Details
S3 Device Summary
“S3 Storage Group Configuration” on page 301
“S3 User Storage Groups” on page 301
“S3 Volume Event” on page 303
“S3 Volume Model” on page 303
“Hitachi Data Systems 2107 Array Configuration” on page 304
“Hitachi Data Systems 2107 Configuration” on page 304
“Hitachi Data Systems 2107 Parity Group” on page 305
“HSM CDS” on page 305
“HSM Common Storage” on page 306
“HSM Cross System CRQ Hosts” on page 306

“HSM Cross System CRQplex” on page 307
“HSM CRQ Hosts” on page 308
“HSM CRQ Requests” on page 308
“HSM CRQplex” on page 309
“HSM Dataset Recover Version” on page 310
“HSM Function Details” on page 310
“HSM Function Statistics” on page 311
“HSM Function Summary” on page 312
“HSM Host Function Status” on page 312
“HSM Host Request Summary” on page 313
“HSM Host Status” on page 314
“HSM Private Storage” on page 315
“HSM Requests” on page 315
“HSM Status” on page 316
“Logical Control Unit” on page 317
“Logical Control Unit Channel Paths” on page 318
“LCU Control Unit” on page 319
“RLS Appl Dataset Details” on page 319
“RLS Buffer LSU Summary” on page 322
“RLS Buffer Pool Detail” on page 323
“RLS Collection Interval” on page 323
“RLS Dataset Details” on page 324
“Dataset Collection Interval” on page 326
“RLS Dataset Group Details” on page 326
“RLS Dataset Group Summary” on page 328
“RLS Lock Structure” on page 329
“RLS Lock Structure CF Detail” on page 331
“RLS Performance Overview” on page 331
“RLS Storage Class” on page 332
“RMM Configuration” on page 334
“RMM Control” on page 336
“RMM Control Dataset” on page 336
“RMM Dataset Details” on page 337
“RMM Location Details” on page 338
“RMM Owner Details” on page 339
“RMM Pending Actions” on page 339
“RMM System Summary” on page 339
“RMM Volume Details” on page 342
“RMM VRS Details” on page 347
“S3 PPRC Connection” on page 348
“S3 PPRC Volume” on page 349
“S3 RLS Cache Structure Summary” on page 349
“S3 XRC Session” on page 350
“S3 XRC Volume Pair” on page 351
“S3 Flash Copy Relation” on page 351
“S3 Flash Copy Volume” on page 352
“S3 Link Statistics” on page 353
“S3 VTS Cluster Worst Statistics Summary” on page 355
“SMS Cache Sets” on page 356
“SMS Configuration” on page 356
“SMS Data Class” on page 356
“SMS Management Class” on page 357
“SMS Storage Class” on page 359
“SMS Storage Group” on page 359
“SMS Storage Group Status” on page 361
“SMS Systems” on page 361
“SMSplex Navigation” on page 361
“SMSplex Summary” on page 362
“Storage Toolkit Requests” on page 363
“Storage Toolkit Result Summary” on page 364
“Storage Toolkit Result Detail” on page 365
“Storage Toolkit Result Resources” on page 366
“Symmetrix Configuration” on page 366
“Symmetrix Devices” on page 366
“Symmetrix Director Summary” on page 367
“Symmetrix Directors” on page 367
“Symmetrix SSIDs” on page 369
“System Summary” on page 369
“Tape Device” on page 371
“Tape Group” on page 371
“TotalStorageDS Array Configuration” on page 372
“TotalStorageDS Configuration” on page 373
“TotalStorageDS Extent Pool” on page 373
“TotalStorageDS Rank” on page 374
“TotalStorageDS SSIDs” on page 375
“UDG Cache Devices” on page 375
“Volume Group Summary” on page 377
“Volume Status” on page 380
“Volume Users” on page 380

"VTS Cache" on page 381
"VTS Capacity" on page 382
"VTS Composite Library" on page 383
"VTS Overview" on page 384
"VTS Physical Devices" on page 384
"VTS TS7700-series Cache Container" on page 385
"VTS TS7700-series Cache Partition Container" on page 386
"VTS TS7700-series Cache Preference Group" on page 387
"VTS TS7700-series Cluster" on page 388
"VTS TS7700-series Cluster Replication" on page 385
"VTS Virtual Devices" on page 389
"VTS Virtual Devices Summary" on page 390
"VTSS Channels" on page 391
"VTSS Configuration" on page 391
"VTSS RTD Queue" on page 392
"VTSS RTD Status" on page 392
"VTSS Status" on page 393
"VTSS Virtual Volumes" on page 393

ADABAS Databases

The ADABAS Databases attributes provide information on the databases related to ADABAS.

Managed System The SMS name of this system.

Collection Time: Time when data was collected.

Adabas Version: Version of Adabas instance.

Database Name: Name assigned to database.

Database Identification Number: Database Identification Number.

SVC number: SVC Number.

Max Files: Maximum number of files permitted for the database.

Log Tape Number: Number of the most recent data protection log for the database.

DB Load Time: Date and Time when the database was initially defined.

High File Number: Number of files loaded to database.

Bytes in RABN: Length of the blocks in the database.

DSST Alloc Extents: The extents allocated for Data Storage Space Table in Blocks.

DATA Allocated Bytes: Number of bytes allocated for DATA component.

DATA Allocated Blocks: Number of blocks allocated for DATA component.

DATA Allocated Cylinders: Number of cylinders allocated for DATA component.

ASSO Allocated Bytes: Number of bytes allocated for ASSO component.

ASSO Allocated Blocks: Number of blocks allocated for ASSO component.

ASSO Allocated Cylinders: Number of cylinders allocated for ASSO component.

DATA Unused Bytes: Number of bytes unused by DATA component.

DATA Unused Blocks: Number of blocks unused by DATA component.

DATA Unused Cylinders: Number of cylinders unused by DATA component.

ASSO Unused Bytes: Number of bytes unused by ASSO component.

ASSO Unused Blocks: Number of blocks unused by ASSO component.

ASSO Unused Cylinders: Number of cylinders unused by ASSO component.

DATA Used Bytes: Number of bytes used by DATA component.

DATA Used Blocks: Number of blocks used by DATA component.

DATA Used Cylinders: Number of cylinders used by DATA component.

ASSO Used Bytes: Number of bytes used by ASSO component.

ASSO Used Blocks: Number of blocks used by ASSO component.

ASSO Used Cylinders: Number of cylinders used by ASSO component.

DATA Percent Used: Percent of the storage, used by DATA component.

ASSO Percent Used: Percent of the storage, used by ASSO component.

Checkpoint File Number: Checkpoint file number.

Security File Number: Security file number.

Trigger File Number: Trigger file number.

Host: Host name.

Sysplex Name: Sysplex name.

ADABAS Files

The ADABAS DB File attributes provide summary information on ADABAS Database files.

Managed System: The SMS name of this system.

Collection Time: Time when data was collected.

Host: Host name.

SVC Number: SVC Number.

Database Identification Number: Database ID number.

File Number: Database file number.

File Name: Database file name.

Data Allocated Bytes: Storage allocated for data component for current file, in bytes.

Data Used Bytes: The used storage of the data component for the current file, in bytes.

Address Converter Allocated Bytes: Storage allocated for the Address Converter component for the current file, in bytes.

Address Converter Used Bytes: Storage used for the Address Converter component for the current file, in bytes.

Upper Index Allocated Bytes: Storage allocated for Upper Index for the current file, in bytes.

Upper Index Used Bytes: Storage used for Upper Index for the current file, in bytes.

Normal Index Allocated Bytes: Storage allocated for Normal Index for the current file, in bytes.

Normal Index Used Bytes: Storage used for Normal Index for the current file, in bytes.

Total Allocated Bytes: Total storage allocated for data, Address Converter, Upper Index and Normal Index components for current file, in bytes.

Total Used Bytes: Total storage used for data, Address Converter, Upper Index and Normal Index components for current file, in bytes.

Database Name: Name assigned to database.

Locked Update File is locked update.

Replicated: File is replicated.

Refresh Permitted: File can be refreshed using the E1 command.

User ISN: File was loaded with the USERISN option.

File Updated: File was updated.

No AC Extension: File permits increasing MAXISN setting.

Privileged Use: File is in privileged usage mode.

Reuse allowed: ISN reuse is allowed.

First Unused ISN: Highest ISN currently used in the file + 1.

Top Expect ISN: Highest ISN that can be assigned to a record in the file.

Highest Index Level: Highest Index Level.

Load Time: Date and Time when the file was loaded.

Last Update: Date and Time of the last file update.

Start ISN Number: Minimum ISN number.

ISN Size: Whether the file contains 3-byte or 4-byte ISNs.

Max Compressed Record Length: Maximum compressed record length permitted for the file.

Number Of Updates: Number of file updates.

Number Of Records: Number of records loaded to file.

Max DS Extentions In Blocks: Maximum number of blocks which may be allocated for each Data Storage secondary extent.

Max UI Extentions In Blocks: Maximum number of blocks which may be allocated for each Upper Index secondary extent.

Max NI Extentions In Blocks: Maximum number of blocks which may be allocated for each Normal Index secondary extent.

DATA Percent Used: Padding factor for DATA component.

ASSO Percent Used: Padding factor for ASSO component.

Sysplex Name: Sysplex name.

DATA Unused Bytes: The unused storage of DATA component for current file, in bytes.

Address Converter Unused Bytes: The unused storage of Address Converter for current file, in bytes.

Upper Index Unused Bytes: The unused storage of Upper Index for current file, in bytes.

Normal Index Unused Bytes: The unused storage of Normal Index for current file, in bytes.

Total Unused Bytes: The unused storage for DATA, Address Converter, Upper Index and Normal Index components for current file, in bytes.

Data Allocated Cyl: Storage allocated for DATA component for current file, in cylinders.

Data Used Cyl: The used storage of DATA component for current file, in cylinders.

Data Unused Cyl: The unused storage of DATA component for current file, in cylinders.

Address Converter Allocated Cyl: Storage allocated for Address Converter component for current file, in cylinders.

Address Converter Used Cyl: The used storage of Address Converter for current file, in cylinders.

Address Converter Unused Cyl: The unused storage of Address Converter for current file, in cylinders.

Upper Index Allocated Cyl: Storage allocated for Upper Index for current file, in cylinders.

Upper Index Used Cyl: The used storage of Upper Index for current file, in cylinders.

Upper Index Unused Cyl: The unused storage of Upper Index for current file, in cylinders.

Normal Index Allocated Cyl: Storage allocated for Normal Index for current file, in cylinders.

Normal Index Used Cyl: The used storage of Normal Index for current file, in cylinders.

Normal Index Unused Cyl: The unused storage of Normal Index for current file, in cylinders.

Total Allocated Cyl: Total storage allocated for DATA, Address Converter, Upper Index and Normal Index components for current file, in cylinders.

Total Used Cyl: The used storage for DATA, Address Converter, Upper Index and Normal Index components for current file, in cylinders.

Total Unused Cyl: The unused storage for DATA, Address Converter, Upper Index and Normal Index components for current file, in cylinders.

Data Storage Extents: Number of logical extents currently assigned to Data Storage.

Address Converter Extents: Number of logical extents currently assigned to Address Converter.

Upper Index Extents: Number of logical extents currently assigned to Upper Index.

Normal Index Extents: Number of logical extents currently assigned to Normal Index.

File UQ DE: File has options "DE,UQ".

ADABAS Volumes

The ADABAS Volume attributes provide summary information on ADABAS Volumes.

Managed System: The SMS name of this system.

Collection Time: Time when data was collected.

Allocated Bytes: Number of bytes allocated for system dataset.

Unused Bytes: Number of bytes unused by system dataset.

Last Reference Date: Last reference date for the dataset.

Host: Host name.

Dataset Name: The MVS™ dataset name.

Volume Serial: Represents the six-character identifier associated with this logical volume.

Database Identification Number: The database ID.

SVC Number: SVC Number.

Volume Sequence Number: Volume sequence number for multi-volume datasets.

Application Datasets

The Application Datasets attributes provide space information on the datasets allocated to a specific application.

Application: The name of the batch job, TSO userid or started task associated with this application.

ASID: The address space ID (in decimal) associated with this application.

Dataset Count: The total number of datasets allocated to this application.

Dataset Name: The name of a dataset in use by this application.

Managed System: The SMS name of this system.

RLS Processing: Indicates that a vsam cluster is open for RLS processing. The value will be no unless the dataset is a vsam component open for RLS processing.

SMF ID: Identifies, by SMF ID the system on which this application is running.

Volume Serial: The volume serial number of a disk device currently in use by this application.

Volume Count: The total number of volumes allocated to this application.

Application Dataset Details

The Application Dataset Details attributes provide space information on the datasets allocated to a specific application.

Block Size: Represents the physical length attribute of the records in this dataset. For a DASD dataset, block size is maintained in the DSCB. Because large blocks transfer more data, large block sizes result in higher connect time. Larger blocks with fewer buffers might be appropriate for sequential datasets. Consider using DASD cache control for heavily used, read-only sequential datasets. Random access files need smaller block sizes with a larger number of PDS libraries, block sizes should usually reflect average member size.

Dataset Name: The MVS™ dataset name.

Dataset Type: Represents the organization of this dataset. This attribute can take the following decimal values:

Direct_Access=11
Direct_Access_Unmovable=12
Extended_Seq_Unmovable=4
Extended_Sequential=3
Hierarchical_File=1

ISAM=7
ISAM_Unmovable=8
Partitioned=13
Partitioned_Unmovable=14
PDS_Extended=2
Physical_Seq_Unmovable=10
Physical_Sequential=9
Unknown=0
Unknown=15
VSAM=5
VSAM_Unmovable=6

Job Name: The MVS™ JOBNAME, TSOUSERID, or STARTED TASK NAME associated with this application.

Logical Record Length: The length attribute associated with the logical records in this dataset.

Number of Extents: The number of DASD extents allocated to this dataset. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Managed System: The SMS name of this system.

Tracks Allocated: Represents the number of DASD tracks occupied by this dataset. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Tracks Used: Represents the number of DASD tracks actually used by this dataset. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Tracks Used Percent: Represents the value of Tracks Used divided by Tracks Allocated for this dataset. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Volume: Represents the six-character identifier associated with this logical volume.

Application Dataset Monitor

The Application Dataset Monitor attributes provide information to other attribute groups. Users do not use these attributes. The attributes are documented for the purpose of background reference only.

Managed System: The SMS name of this system.

Application Monitoring

The Application Monitoring attributes provide summary performance information for an application and control the monitoring of applications with the Monitor Status and I/O Monitor Status columns.

Application: The name of the batch job, TSO user ID, or started task associated with this application.

ASID Hex: The address space ID (in hex) associated with this application.

ASID: Indicates the number of the address space ID (in decimal) associated with this application.

Dataset Count: Represents the total number of datasets allocated to this application in this address space. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Dataset Name: Specifies the MVS™ dataset name.

Dataset with High MSR: The MVS™ dataset name with the longest response time in milliseconds associated with this application.

High Dataset MSR: Represents the longest response time in milliseconds for this application to access any dataset. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

High Volume MSR: Represents the longest response time in milliseconds for any volume in use by this application. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

I/O Monitor Status: Used to toggle the status of I/O monitoring at the dataset level for the application. This value is used for input only when defining a situation to start or stop monitoring of an application. This value is meaningless in a report, because it is always reported as an empty field. When specified as Y (Yes), it triggers MSR (millisecond response time) monitoring at the dataset level for all datasets in use by the application.

Y = Turn ON dataset-level I/O monitoring for this application.
N = Turn OFF dataset-level I/O monitoring for this application.

I/O Second: Specifies the average number of I/O instructions per second issued by this application. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Managed System: The Storage Management Subsystem name of this system.

Monitor Status: Represents the current monitor status of this application. The values are: Y=The application is being monitored. N=The application is not being monitored.

RLS Dataset Open: An indicator as to whether the address space has VSAM datasets open and using Record Level Sharing (RLS) processing.

SMF ID: The SYSTEM ID of the MVS™ system which is running this application.

Volume Count: Represents the total number of volumes allocated to this application in this address space. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Volume with High MSR Dataset: The volume label of the logical volume containing the dataset with the longest response time in milliseconds for this application.

Volume with High MSR: The six-character identifier of the logical device with the longest response time in milliseconds in use by this application.

Volume: The six-character identifier associated with this logical volume.

Application RLS Datasets

The Application RLS Datasets attributes return the cluster names for an application that are opened for RLS processing.

ASID: The address space ID (in decimal) associated with this application.

Dataset Count: The total number of VSAM datasets open for RLS processing allocated to this application.

Dataset Name: The name of a VSAM cluster in use by this application opened for RLS processing.

Application: The name of the batch job, TSO userid or started task associated with this application.

ORIGINNODE: An internal node name used to identify the system.

RLS Processing: Indicates that the VSAM cluster is open for RLS processing. The value will be no unless the dataset is a VSAM component open for RLS processing.

SMF ID: Identifies, by SMF ID the system on which this application is running.

Volume: The volume serial number of a disk device currently in use by this application.

Volume Count: The total number of volumes allocated to this application on which VSAM datasets open for RLS processing reside.

Application Volume Performance

The Application Volume Performance attributes provide performance information for the logical volumes allocated to an application.

Average Command Response Delay: The average number of milliseconds that a successfully initiated start or resume function needs until the first command is indicated as accepted by the device over the RMF™ interval.

Average HyperPAV Alias Count: The average number of aliased devices assigned to a base volume in a HyperPAV environment over the RMF™ interval.

Busy Percent: The ratio, expressed as a percentage, of the total busy time this interval divided by the interval length.

Bypass Cache Percent: The percentage of I/O that goes directly to a DASD device regardless of whether the data might be or currently is cached.

Cache Status: Indicates whether caching is enabled for this device. Values are Timeout, Active, Inactive, Pending Active, Pending Inactive, and n/a.

Cache to DASD Total Tracks: The number of tracks written from cache to physical disk during the cache interval. This value is captured asynchronously with respect to the transfer of data from the channel to the cache.

Cache to DASD Tracks per Sec: The rate per second of tracks written from cache to physical disk during the cache interval. This value is captured asynchronously with respect to the transfer of data from the channel to the cache.

CFW Read Percent: The percentage of I/O requests read from cache that were in cache due to a cache fast write.

CFW Write Percent: The percentage of I/Os that were cache fast write requests. Cache Fast Write (CFW) is a function of 3990 controllers that writes data to cache at channel speed. The data is later written to DASD. Cache fast write differs from DASD fast write in that the data is not written to NVS (nonvolatile storage).

Connect Time: The part of I/O during which data is located on the track and transferred between the DASD device and central storage.

Current PAV Exposures: The number of exposures for the device when the data was gathered.

DASD to Cache Non Promotes Second: The rate per second of operations that did not promote (copy) a track to cache because free cache segments were not available. Non promotes are tracks that the system does not copy to cache because free cache segments were not available.

DASD to Cache Non Promotes: The number of operations that did not promote (copy) a track to cache because free cache segments were not available. Non promotes are tracks that the system does not copy to cache because free cache segments were not available.

DASD to Cache Total Tracks Non-sequential: The number of tracks read from physical disk to cache over the cache interval for non-sequential I/O operations.

DASD to Cache Total Tracks Sequential: The number of tracks read from physical disk to cache over the cache interval for sequential I/O operations.

DASD to Cache Total Tracks: The number of tracks read from physical disk to cache over the cache interval for I/O operations.

DASD to Cache Tracks per Sec Non-sequential: The rate per second of tracks read from physical disk to cache over the cache interval for non-sequential I/O operations.

DASD to Cache Tracks per Sec Sequential: The rate per second of tracks read from physical disk to cache over the cache interval for sequential I/O operations.

DASD to Cache Tracks per Sec: The rate per second of tracks read from physical disk to cache over the cache interval for I/O operations.

DCBs Open: The count of open DCBs for datasets.

Device Address: The MVS™ device number, assigned by HCD, associated with this device.

Device MPL: The device Multi-Processing Level, which is the I/O density for a device. This metric is the product of the MSR (millisecond response time) and I/O per second for a device.

DFW Hit Percent: The ratio of I/O requests where data is written to both cache and nonvolatile storage that was satisfied by the storage control cache without requiring access to the hard disk.

DFW Retry Percent: The percentage of DASD fast write operations that did not complete due to NVS (nonvolatile storage) constraint.

DFW Status: Indicates whether DASD Fast Write is enabled on this device.

Disconnect Time: Time that I/O spends, freed from the channel, searching for or requesting data. This time includes moving the device to the requested cylinder and track, waiting for the record to rotate under the head (rotational delay), and waiting for the channel to become available.

Interrupt Delay Time: The average interrupt delay time for an I/O. This is the elapsed time from the completion of an I/O until z/OS issues test subchannel to retrieve the results.

I/O Count: The total number if I/O operations performed on this device during this interval.

I/O Per Second: The average number of input/output operations, per second, directed at this volume during this interval.

Inhibit Cache Percent: The percentage of I/O that used the MVS™ I/O mechanism to prevent a request from being cached, even if it is eligible. If the data is already in cache, the request is satisfied by cache.

IOSQ Delay: The average time that an I/O waits because the device is already busy (signified by the UCBBUSY bit for the device being on).

Job Name: The MVS™ JOBNAME, TSO user ID, or started task name associated with this application.

Managed System: The SMS name of this system.

Maximum PAV Exposures: The maximum number of exposures for a device over the current RMF™ interval.

MSR Connect Time Percent: The ratio of average connect time divided by average total response time.

PAV Exposure Changed: An indicator that the number of PAV exposures for the volume has changed over the interval.

Pend Time: The time that I/O is delayed in the path to the device. Pending time can be attributed to the channel or control unit path being busy.

Physical Device: The physical device type.

Read Hit Percent: The ratio of read requests that were satisfied by the storage control unit cache without requiring access to DASD.

Read I/O Count: The number of read requests in the current collection interval.

Read I/O Percent: The percentage of I/O requests that were read requests.

Reserved Percent: The percentage of time during the interval when a shared device was reserved by this system.

Response Time: The total hardware service time and the front end software queuing time in milliseconds involved for the average I/O request to the device.

Solid State Device: Indicates whether the logical volume is backed by a solid state physical device.

Storage Group Name: Indicates whether the number of PAV exposures for this device has changed during this interval.

Subsystem ID: The subsystem ID.

Time Stamp: The date and time this data was reported.

Volume: Represents the six-character identifier associated with this logical volume.

Write Hit Percent: The ratio of write hit requests that were satisfied by the storage control unit cache without requiring access to DASD.

Write I/O Count: The number of write requests in the current collection interval.

Application Volume Space

The Application Volume Space attributes provide space information for the logical volumes allocated to an application.

Device Address: The MVS™ device number, assigned by HCD, associated with this device.

Device Type: The MVS™ device type, for example 3380, 3390, or 2105.

Fragmentation Index: The qualitative measure of the scattered free space on a volume. The higher the number, the more severe the fragmentation.

Free Space MB: Represents the total amount of free space in MB currently available on this device.

Job Name: The JOBNAME, TSO ID, or started task name associated with this application.

Largest Free Extent MB: The largest contiguous free area available in MB.

Managed System: The SMS name for this system.

Percent Free Space: The ratio of free space divided by capacity, expressed as a percentage.

SMS Conversion Status: The SMS conversion status of this volume. The values are: non-SMS, initial, converted, or excluded.

SMS Status: The SMS status of this volume. Values are enabled, quiesced_all, quiesced_new, disabled_all, or disabled_new.

Solid State Device: Indicates whether the logical volume is backed by a solid state physical device.

Storage Group Name: The 32-character SMS storage group name to which this volume belongs.

Total Capacity MB: Represents the total data capacity of this device in MB.

Volume: The six-character identifier associated with this logical device.

VTOC Index Status: The status of the VTOC index on this volume. Values are: Enabled, Disabled, Error, OS-VTOC, or Exclude List.

Cache Control Unit

The Cache Control Unit attributes provide performance and status information for cache control units.

Active Volumes: The number of volumes with ACTIVE caching status on the cache control unit.

Bypass Cache Percent: The percentage of I/Os that bypassed cache during this interval. The I/O request is sent directly to DASD. Tracks are not read from cache or promoted to cache. Also, the entries of tracks in the LRU list are not updated.

Cache MB Available: The total of cache storage in MB available on this Cache Control Unit.

Cache MB Configured: The total storage in MB of cache storage configured on this Cache Control Unit.

Cache Status: The caching status for this subsystem. Values are as follows: Active (in other words, caching for this subsystem is active), Inactive (in other words, caching for this subsystem is inactive), Pending Active, Pending Inactive, and Timeout.

CFW Read Percent: The percentage of all I/Os that were Read operations for data in cache that originated from a cache fast write during this interval. Cache fast write (CFW) is a form of fast write where the data is written directly to cache without using nonvolatile storage and is available for later destaging.

CFW Write Percent: The percentage or all I/Os that were Cache Fast Writes during this interval. Cache fast write (CFW) is a form of fast write where the data is written directly to cache without using nonvolatile storage and is available for later destaging.

Control Unit Type: The type of this cache control unit. This field is only populated for the physical devices that are supported with extended information. Examples of these devices include EMC Symmetrix, 2105, and 2107.

Deactivated Volumes: The number of volumes on this Cache Control for which the cache has been deactivated.

DFW Hit Percent: The percentage of all DASD I/Os to this control unit that were DASD Fast Writes during this interval. DASD Fast Write (DFW) is form of fast write to cache where the data is written concurrently to cache and nonvolatile storage and automatically scheduled for destaging to DASD.

DFW Retry Percent: The percentage of DASD I/O operations which required retry during this interval.

Inhibit Cache Percent: The percentage of I/O operations which inhibited the loading of data into cache.

Managed System: The SMS name of this system.

NVS KB Configured: The KB of cache storage configured on this subsystem. NVS (nonvolatile storage) facilitates the DASD Fast Write feature which allows data to be written at electronic speeds directly into this nonvolatile memory in the controller, thus avoiding the many delays associated with accessing the actual disk.

NVS KB Pinned: The amount of storage that is unavailable because DASD failure is preventing the subsystem from destaging data. The data is pinned in cache.

NVS Status: The overall status of the nonvolatile storage (NVS). Values are:

Timeout=0 Status is timed out.
Active=1 NVS is online and usable.
Inactive=2 NVS is not usable.
Pending_Active=3 Status is switching from inactive to active.
Pending_Inactive=4 Status is switching from active to inactive.
Timeout=5 Status is timed out.
Failed=6 An internal subsystem error occurred that terminated non-volatile storage availability. All data has not been successfully destaged or discarded.
Compromised=7 NVS is pending in status.

Read Hit Percent: The percentage of read hits for each cache control unit.

Read I/O Percent: The percentage of I/O requests to the cache control unit that were read requests.

Serial Number Index: The serial number of the control unit. Used as an index into hardware specific displays.

Storage Facility ID: The Storage Facility Identifier for the device (Type-Model-Serial Number).

Storage Facility Series: The series of control unit, which defines the physical geometry of the subsystem. Can be one of: HDS, Symmetrix, Iceberg, RAMAC®, shark, DS6000™, or DS8000™.

Subsystem ID Fullword: The subsystem ID (for Model 3990, 2105, and 2107 control units).

Subsystem ID Hex: The subsystem ID (for Model 3990, 2105, and 2107 control units).

Subsystem ID: The subsystem ID (for Model 3990, 2105, and 2107 control units).

Time Stamp: The date and time this data was reported. Primarily for use in historical reporting.

Track Destaging Rate: The number of tracks per second being removed from cache to physical disks over the cache interval.

Track Staging Rate: The number of tracks per second being brought into cache from the physical disks over the cache interval.

Write Hit Percent: The percentage of I/O operations that completed without accessing the hard disk drive.

Cache Devices

The Cache Devices attributes provide both cache status and cache performance information for a logical volume.

Avg Logical KB Read: The average number of KB read from the logical volume per second over the cache interval.

Avg Logical KB Written: The average number of KB written to the logical volume per second over the cache interval.

Avg Logical Read Time: The average time in milliseconds spent reading from the logical volume per I/O over the cache interval.

Avg Logical Write Time: The average time in milliseconds spent writing to the logical volume per I/O over the cache interval.

Avg Phys Read Response Time: The average time to read from the physical disk per operation in milliseconds over the cache interval.

Avg Phys Write Response Time: The average time in milliseconds to write to the physical disk per operation over the cache interval.

Bypass Cache Percent: The percentage of I/Os that bypassed cache during this interval. The I/O request is sent directly to DASD. Tracks are not read from cache and are not promoted to cache, and their entries in the LRU list are not updated.

Cache Status: The caching status for the subsystem to which this device is attached. Active means caching for this subsystem is active. Inactive means caching for this subsystem is disabled.

Cache to DASD Total Tracks: The number of tracks written from cache to physical disk during the cache interval. This value is captured asynchronously with respect to the transfer of data from the channel to the cache.

Cache to DASD Tracks per Second: The rate per second of tracks written from cache to physical disk during the cache interval. This value is captured asynchronously with respect to the transfer of data from the channel to the cache.

CFW Read Percent: The percentage of READ requests which were satisfied from cache due to a previous cache fast write.

CFW Write Percent: The percentage of all I/Os that were CACHE FAST WRITES during this interval. CACHE FAST WRITE allows data to be written to cache only. The data is later written to DASD. CFW differs from DASD Fast Write in that the data is not written to nonvolatile storage (NVS).

Channel Index: The identifier for channel path that identifies which channel path this device is attached to.

Channel Path ID 1: The channel path identifier (in decimal) of one channel path connecting this device.

Channel Path ID 2: The channel path identifier (in decimal) of one channel path connecting this device.

Channel Path ID 3: The channel path identifier (in decimal) of one channel path connecting this device.

Channel Path ID 4: The channel path identifier (in decimal) of one channel path connecting this device.

Channel Path ID 5: The channel path identifier (in decimal) of one channel path connecting this device.

Channel Path ID 6: The channel path identifier (in decimal) of one channel path connecting this device.

Channel Path ID 7: The channel path identifier (in decimal) of one channel path connecting this device.

Channel Path ID 8: The channel path identifier (in decimal) of one channel path connecting this device.

Control Unit Model: The model number of the control unit to which this device is attached.

DASD to Cache Non-promotes per Second: The rate per second of operations that did not promote (copy) a track to cache because free cache segments were not available.

DASD to Cache Non-promotes: The number of operations that did not promote (copy) a track to cache because free cache segments were not available.

DASD to Cache Total Tracks Non-sequential: The number of tracks read from physical disk to cache over the cache interval for non-sequential I/O operations.

DASD to Cache Total Tracks Sequential: The number of tracks read from physical disk to cache over the cache interval for sequential I/O operations.

DASD to Cache Total Tracks: The number of tracks read from physical disk to cache over the cache interval for I/O operations.

DASD to Cache Tracks per Sec: The rate per second of tracks read from physical disk to cache over the cache interval for I/O operations.

DASD to Cache Tracks per Second Non-sequential: The rate per second of tracks read from physical disk to cache over the cache interval for non-sequential I/O operations.

DASD to Cache Tracks per Second Sequential: The rate per second of tracks read from physical disk to cache over the cache interval for sequential I/O operations.

Device Address: The four hexadecimal digits that define the physical address of this device. Also known as unit or unit address.

Device Hex Address: The four-digit hexadecimal address of this device.

DFW Hit Percentage: The percentage of DASD Fast Writes that were satisfied from cache without requiring access to DASD. DASD Fast Writes are written to both cache and nonvolatile storage (NVS).

DFW Retry Percentage: The percentage of DASD Fast Write operations that did not complete due to nonvolatile Storage (NVS) constraints.

DFW Status: The status of DASD FAST WRITE function for this volume. Values are Active, Inactive, Unknown, or Timeout.

Extent Pool ID: The Extent Pool Identifier.

I/O Count: The number of cache eligible I/Os to the device over the cache collection interval.

Inhibit Cache Percentage: The percentage of I/O requests that set the Inhibit Cache Load mode in the channel program to prevent the data being cached. If the data is already in cache, the request is satisfied from cache.

LCU Number Hex: The logical control unit number for this device (hexadecimal). An LCU is the logical representation of either a single control unit, with or without attached devices, or a group of control units that share devices.

Logical Control Unit Index: The number of LCU to which this device is attached. An LCU is the logical representation of either a single control unit, with or without attached devices, or a group of control units that share devices.

Logical KB Read: The total number of KB read from the logical volume over the cache interval.

Logical KB Written Total: The total number of KB written to the logical volume over the cache interval.

Logical Read Time: The amount of time in milliseconds spent reading from the logical volume over the cache interval.

Logical Write Time: The amount of time in milliseconds spent writing to the logical volume over the cache interval.

Managed System: The SMS name of this system.

Physical Device: The identifier of the physical device type of this device.

Physical KB Read Rate: The KB read from the physical disk per second over the cache interval.

Physical KB Read Total: The total KB read from the physical disk over the cache interval.

Physical KB Write Rate: The KB written to the physical disk per second over the cache interval.

Physical KB Write Total: The total KB written to the physical disk over the cache interval.

Physical Read Rate: The rate of I/O read operations from the physical disk per second over the cache interval.

Physical Read Total: The total number of I/O read operations from the physical disk over the cache interval.

Physical Write Rate: The rate of I/O write operations to the physical disk per second over the cache interval.

Physical Write Total: The total number of I/O write operations to the physical disk over the cache interval.

RAID Flag: Indicates the type of RAID device.

Read Hit Percent: The percentage of READ requests that were satisfied from cache.

Read I/O Count: The number of read requests for this device.

Read I/O Percent: The percentage of I/O requests that are read requests.

Solid State Device: Indicates whether the logical volume is backed by a solid state physical device.

SSID Index: The fullword representation of the SSID used for indexing the product.

Storage Facility ID: The storage facility identifier for the device (Type-Model-Serial Number).

Storage Group Index: The index of the storage group to which this volume belongs. The storage group is an SMS construct used to group DASD devices in order to separate the physical requirements of storing data from the logical requirements.

Storage Group Name: The name of the storage group to which this volume belongs. The storage group is an SMS construct used to group DASD devices in order to separate the physical requirements of storing data from the logical requirements.

Subsystem ID Hex: The subsystem ID associated with this device.

Subsystem ID Selector: An index to limit the range of physical volumes for which data is requested.

Subsystem ID: The two-byte unique value used to identify a 3990 or equivalent storage control unit.

Time Stamp: Represents the date and time this data was collected.

Total Phys Read Response Time: The total time to read from the physical disk in milliseconds over the cache interval.

Total Phys Write Response Time: The total time to write to the physical disk in milliseconds over the cache interval.

User Group Index: A user DASD group is a facility that enables you to define groups of DASD volumes by VOLSER and device address. You can also specify generic VOLSERs and devices (patterns) by using the wildcard character.

Volume: The volume serial number of the disk device.

Write Hit Percent: The percentage of I/O WRITE operations that were satisfied without having to access the physical drive. This occurs as the result of DASD Fast Write or Cache Fast Write features. DASD Fast Write is a default in controllers that support it. Cache Fast Write can be requested in the Define Extent command of the channel program.

Write I/O Count: The number of write requests for this device.

zHPF Prefetch Hit Percent: The percentage in tenths of zHPF prefetch requests that resulted in data being returned without accessing the physical disk device.

zHPF Prefetch Hits: The number of zHPF prefetch I/O requests that were satisfied without accessing the physical device media.

zHPF Prefetch Requests: The number of zHPF prefetch requests for this device over the current monitoring interval.

zHPF Read I/O per Second: The rate of zHPF read I/O requests per second in tenths of a request.

zHPF Read I/O Requests: The number of zHPF read I/O requests for this device over the current monitoring interval.

zHPF Request Percent: The percent in tenths of cache I/O requests that utilized zHPF.

zHPF Write I/O per Second: The rate of zHPF write I/O requests per second in tenths of a request.

zHPF Write I/O Requests: The number of zHPF write I/O requests for this device over the current monitoring interval.

Cache Raid Rank

The Cache Raid Rank attributes provide performance and usage information for the ranks in an IBM® 2105 storage facility.

Device Adapter ID: The two hexadecimal digit of the device adapter associated with this rank. The device adapter is the hardware component that provides the path control and data transfer functions.

Hard Disk Drives in Rank: The number of disk drives contained in this rank.

Managed System: The SMS name of this system.

Raid Rank ID: The four hexadecimal digit ID associated with this raid rank.

Raid Rank ID String: The raid rank ID in string format.

Read MB per Request: The average number of MB of data transferred per read request.

Read MB per Second: The average number of MB of data transferred per second for read requests.

Read Requests per Second: The average number of read operations per second.

Read Response Time MS: The average response time in milliseconds for read requests.

Read Sectors per Request: The average number of hard disk sections read per read request.

Sector Size: The size in bytes of the sectors on the hard disks associated with this rank.

Sectors Read per Second: The average number of hard disk sectors read per second for read requests.

Subsystem ID Fullword: The subsystem ID of the logical control unit to which this rank is attached.

Subsystem ID: The subsystem ID of the logical control unit to which this rank is attached.

Time Stamp: The date and time associated with this row of data.

Volume: The volume serial number of the first logical volume on this rank.

Write MB Per Request: The average number of MB written per request for this rank.

Write MB Per Second: The average number of MB of data written per second for this rank.

Write Requests Per Second: The average number of write operations per second for this rank.

Write Response Time MS: The average time in milliseconds required to complete a write operation.

Write Sectors per Request: The average number of physical sectors written per write request.

Write Sectors per Second: The average number of physical sectors written per second.

Catalog Dataset Attributes

The Catalog Dataset Attributes attribute group provides dataset attributes for catalogs. This data is required for efficient catalog management and setting situations to get alerts.

Aliases per GB: The number of Aliases per GB.

Allocated Cyls: The cylinders allocated for the catalog's dataset.

Allocated MB: The MB allocated for the catalog's dataset.

Allocated Tracks: The tracks allocated for the catalog's dataset.

Allocation Type: The unit type used to allocate DASD space for the catalog.

Cache Type: The type of cache.

Catalog Name: The catalog name.

CA Splits: The number of control area splits.

CIs per CA: The number of control intervals in the control area for the catalog.

CI Size: The control interval size, in bytes, for the catalog.

CI Splits: The number of control interval splits.

Collection Time: The time when data was collected.

Extents: The number of DASD extents used by the catalog.

Is Extendable: The flag if the catalog can be extended on the volume.

High Allocated RBA: The high-allocated RBA (Relative Byte Addresses) for the catalog.

High Used RBA: The high-used RBA (Relative Byte Addresses) for the catalog.

Largest Free Extent In Alloc Type: The largest Volume Free Extent calculated in Allocation Type.

Locked: The catalog locked status.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Max Possible Extensions: The number of maximum possible extends in the largest volume free extend on which the catalog resides.

Number of Aliases: The number of Aliases in the catalog.

Primary Allocation: The primary space allocation of DASD space for the catalog.

Secondary Allocation: The secondary space allocation of DASD space for the catalog.

Shared: The catalog shared status.

SMS Managed: The SMS-managed status.

Used CAs: The number of control areas within the HURBA for the catalog.

Volume: The volume where the catalog resides.

% Used CAs: The percentage of the HARBA for the catalog.

% 4GB Limit: The HURBA expressed as a percentage of 4 GB.

CA1 Configuration

The CA1 Configuration attributes provide configuration and control information for the CA1 tape management system.

Managed System: The SMS name of this system.

Number of Formatted DSNB Records: Total number of formatted (ready to use) DSNB records.

Number of Active DSNB Records: Total number of DSNB records currently in use.

First DSNB Offset: Relative address of the first DSNB record.

Next Free DSNB Offset: Relative address of the next free DSNB record.

First Volume Definition Offset: Relative address of the first volume definition record.

Number of Volume Definition Records: Number of volume definition records.

Number of Audit Records Generated: Number of audit records generated.

Next Audit Offset: Relative address of the next available audit record.

First Audit Record Since Last Backup: First audit record since last backup.

Last TMC Backup: Time of last TMC backup.

TMC Backup VOLSER: VOLSER of last backup tape.

TMC Backup DSN: DSN of last backup tape.

TMC Backup Indicator: TMC backup indicator.

TMSEXTEND Indicator: TMS EXTEND audit flag.

New TMC VOLSER: VOLSER of the new backup tape.

CA1 Controls

The CA1 Controls attributes provide agent collection parameter information for the CA1 tape management system.

Managed System: The SMS name of this system.

Message: Message.

Controls: Controls.

Timestamp: Time table row was created in CT timestamp format.

CA1 Dataset

The CA1 Dataset attributes provide dataset information for the CA1 tape management system.

Managed System: The SMS name of this system.

Dataset Name: Dataset name.

DSNB Active Flag: DSNB record is active.

DSNB Flag1: DSNB flag #1.

Temp Dataset: First dataset is temporary.

Dataset Recreated: First dataset has been recreated.

Dataset is on catalog: Secondary dataset is on OS catalog.

Dataset was on catalog: Dataset was on OS catalog.

DSNB Offset: DSNB record offset.

Next DSNB Offset: Next DSNB record offset.

Previous DSNB Offset: Previous DSNB record offset.

Volser: Dataset starts on this volume.

DSNB Chain Volser: Volume associated with the beginning of the DSNB chain.

Dataset Sequence Number: Dataset sequence number in the volume set.

Record Format: Record format.

Logical Record Length: Logical record length.

Block Size: Block size.

Block Count: Block count.

SMS Management Class: SMS management class.

Compression Ratio: Compression ratio.

Catalog Count: Catalog count.

Encryption Key Index: Encryption key index.

Volume Percent: Percentage of volume used by this dataset.

Creation Date: The date when the dataset was first written to tape.

Days Since Creation: Days since creation.

Creating Program Name: The name of the program that created the dataset.

Creating Job Name: The name of the job that created the dataset.

Creating Step Name: The name of the step that created the dataset.

Creating DD Name: The DD statement that created the dataset.

Creating Unit Address: The device number of the drive on which the volume was mounted.

Expiration Date: The expiration date of the dataset.

Expiration Value: Contains additional retention information derived from expiration date.

Days Until Expiration: Days until expiration. This will be zero if retention policy is specified in Expiration Value attribute.

Default Expiration used at open: Default expiration used at open.

CATALOG Retention Period: Retain for specified number of days and then retain while dataset is cataloged to the operating system.

Expired by Catalog Control: Expired by catalog control.

Expired by TMS: Expired by TMS interface.

Closed by ABEND processing: Dataset closed by ABEND.

Eligible for RDS Override: Dataset is eligible for RDS override.

Audit Last CA1: ID of the last CA1 program to update the record.

Audit Last Intercept: ID of last intercept to update the record.

Audit Update: Date of last update to the record.

Audit User: ID of last user to update the record.

Audit CPU: ID of last CPU used to update the record.

Audit Code: Audit code.

Audit Flag: Audit flag.

Volume Sequence Number: Volume sequence number for multi-volume datasets.

Volume Set ID: First volume in the volume set.

Dataset Mask: DSN mask used by search.

CA1 Job Summary

The CA1 Job Summary attributes provide job summary information for the CA1 tape management system.

Managed System: The SMS name of this system.

Job Name: Creating job name.

Total Datasets: The total number of datasets on all volumes created by the job.

Average Dataset Size: Average dataset size expressed in MB to 1 decimal place.

CA1 Permanent Error Summary

The CA1 Permanent Error Summary attributes provide permanent error summary information for the CA1 tape management system.

Managed System: The SMS name of this system.

Volume: Volume serial number.

Perm Read Errors Since Init: Number of permanent read errors since initialization.

Perm Write Errors Since Init: Number of permanent write errors since initialization.

Perm Read Errors Since Clean: Number of permanent read errors since clean.

Perm Write Errors Since Clean: Number of permanent write errors since clean.

Total Perm Read Errors: Total number of permanent read errors (since init or clean).

Total Perm Write Errors: Total number of permanent write errors (since init or clean).

Total Perm Errors: Total number of permanent errors.

Total Errors: Total number of errors.

CA1 Program Summary

The CA1 Program Summary attributes provide program summary information for the CA1 tape management system.

Managed System: The SMS name of this system.

Program Name: Creating program name.

Total Datasets: The total number of datasets on all volumes created by the program.

Average Dataset Size: Average dataset size expressed in MB to 1 decimal place.

CA1 Summary

The CA1 Summary attributes provide performance information for the CA1 tape management system.

Average Compression Ratio: Average compression ratio.

Average Datasets Per Volume: Average number of datasets per volume.

Average Dataset Size: Average dataset size expressed in MB to 1 decimal place.

Average Dataset Usage Percent: Average dataset usage percent.

Average Volume Open Count: Average number of times tape opened since birth date.

Average Volume Usage Percent: Average volume usage percent.

Collection End Time: End time of data collection.

Collection Start Time: Start time of data collection.

Collection Status: Status of the data collector.

Managed System: The SMS name of this system.

Total Volumes: Total number of active volumes.

Total New Scratch Volumes: Total number of new (uninitialized) scratch volumes.

Total Scratch Volumes: Total number of scratch volumes.

Temp Read Errors Since Init: Number of temporary read errors since initialization.

Vol Temp Read Errors Since Init: Number of volumes with temporary read errors since initialization.

Temp Write Errors Since Init: Number of temporary write errors since initialization.

Vol Temp Write Errs Since Init: Number of volumes with temporary write errors since initialization.

Temp Read Errors Since Clean: Number of temporary read errors since clean.

Vol Temp Read Errs Since Clean: Number of volumes with temporary read errors since clean.

Temp Write Errors Since Clean: Number of temporary write errors since clean.

Vol Temp Write Errors Since Clean: Number of volumes with temporary write errors since clean.

Perm Read Errors Since Init: Number of permanent read errors since initialization.

Vol Perm Read Errs Since Init: Number of volumes with permanent read errors since initialization.

Perm Write Errors Since Init: Number of permanent write errors since initialization.

Vol Perm Write Errors Since Init: Number of volumes with permanent write errors since initialization.

Perm Read Errors Since Clean: Number of permanent read errors since clean.

Vol Perm Read Errors Since Clean: Number of volumes with permanent read errors since clean.

Perm Write Errors Since Clean: Number of permanent write errors since clean.

Vol Perm Write Errors Since Clean: Number of volumes with permanent write errors since clean.

Total Temp Read Errors: Total number of temporary read errors (since init or clean).

Total Vol Temp Read Errors: Total number of volumes with temporary read errors (since init or clean).

Total Temp Write Errors: Total number of temporary write errors (since init or clean).

Total Vol Temp Write Errors: Total number of volumes with temporary write errors (since init or clean).

Total Perm Read Errors: Total number of permanent read errors (since init or clean).

Total Vol Perm Read Errors: Total number of volumes with permanent read errors (since init or clean).

Total Perm Write Errors: Total number of permanent write errors (since init or clean).

Total Vol Perm Write Errors: Total number of volumes with permanent write errors (since init or clean).

Total Temp Errors: Total number of temporary errors (since init or clean).

Total Vol Temp Errors: Total number of volumes with temporary errors (since init or clean).

Total Perm Errors: Total number of permanent errors (since init or clean).

Total Vol Perm Errors: Total number of volumes with permanent errors.

Total Errors: Total number of errors.

Total Volume Errors: Total number of volumes with errors.

Total RDS Override Volumes: Total volumes eligible for RDS Override.

Total Volumes in Delete Status: Total volumes in delete status.

Total Volumes To Be Cleaned: Total volumes listed to be cleaned.

Total Volumes Closed by ABEND: Total volumes closed by ABEND.

Total Volumes Closed Normally: Total volumes closed normally.

Total Datasets: Total number of datasets.

Total Usage: The amount of uncompressed data written to the tape, expressed as GB.

Volume 0% Utilized: Number of volumes not being utilized.

Volume 1-10% Utilized: Number of volumes being utilized 1-10%.

Volume 11-20% Utilized: Number of volumes being utilized 11-20%.

Volume 21-30% Utilized: Number of volumes being utilized 21-30%.

Volume 31-40% Utilized: Number of volumes being utilized 31-40%.

Volume 41-50% Utilized: Number of volumes being utilized 41-50%.

Volume 51-60% Utilized: Number of volumes being utilized 51-60%.

Volume 61-70% Utilized: Number of volumes being utilized 61-70%.

Volume 71-80% Utilized: Number of volumes being utilized 71-80%.

Volume 81-90% Utilized: Number of volumes being utilized 81-90%.

Volume 91-100% Utilized: Number of volumes being utilized 91-100%.

CA1 Temporary Error Summary

The CA1 Temporary Error Summary attributes provide temporary error summary information for the CA1 tape management system.

Managed System: The SMS name of this system.

Volume: Volume serial number.

Temp Read Errors Since Init: Number of temporary read errors since initialization.

Temp Write Errors Since Init: Number of temporary write errors since initialization.

Temp Read Errors Since Clean: Number of temporary read errors since clean.

Temp Write Errors Since Clean: Number of temporary write errors since clean.

Total Temp Read Errors: Total number of temporary read errors (since init or clean).

Total Temp Write Errors: Total number of temporary write errors (since init or clean).

Total Temp Errors: Total number of temporary errors.

Total Errors: Total number of errors.

CA1 Volume

The CA1 Volume attributes provide volume information for the CA1 tape management system.

Actual VOLSER: VOLSER provided by offsite backup facility.

Actual VOLSER In Use: See "Actual VOLSER" attribute for value.

Additional Files In Vol Set: Additional files exist in the volume set.

Address of First DSNB: Relative address of first DSNB associated with the volume set.

Address of Last DSNB: Relative address of last DSNB associated with this volume.

Audit Code: The audit code.

Audit CPU: ID of last CPU used to update the record.

Audit Date: Date of last update to the record.

Audit Flag: The audit flag.

Audit Last CA1: ID of the last CA1 program to update the record.

Audit Last Intercept: ID of last intercept to update the record.

Birthdate: Date when tape was first initialized.

CA-9 R Plus Indicated Bad Tape: CA-9 R Plus indicated bad tape.

Catalog Count: The catalog count.

CATALOG Retention Period: Retain for specified number of days and then retain while dataset is cataloged to the operating system.

Closed by ABEND: Volume has been closed by ABEND.

Closed Normally: Volume has been closed normally.

Compression Ratio: The compression ratio.

Controlled by CA-DYNAM_T: Controlled by CA-DYNAM_T.

Created by CA-1 COPYCAT: Created by CA-1 COPYCAT.

Creating DD Name: The DD statement that created the volume or set.

Creating Job Name: The name of the job that created the volume or set.

Creating Program Name: The name of the program that created the volume or set.

Creating Step Name: The name of the step that created the volume or set.

Creating Unit Address: The device number of the drive on which the volume was mounted.

Creation Date: The date when the volume changed status, either to or from scratch.

CYCLE Retention Period: Retain for specified number of cycles.

Dataset Count: Number of datasets stored on this volume.

Dataset Erase Required: Dataset erase required.

Dataset Name: First dataset on the volume or the volume set.

Dataset Recreated: Dataset recreated.

Dataset was on MVS Catalog: Dataset was on MVS catalog.

Days Since Creation: The number of days since creation.

Days Since First Use: Number of days since first use.

Days Since Last Used: The number of days since last used.

Days Until Expiration: Days until expiration. This will be zero if retention policy is specified in Expiration Value attribute.

Default Expiration Date Used: Default expiration date used.

Delete Status: The delete status.

Density: Recording density.

DSN Pad: Padding because DSNTAIL changed length.

DSN Tail: Last 17 bytes of first dataset DSN on this volume.

Eligible for RDS Override: Volume is eligible for RDS override.

Encryption Key Index: The encryption key index.

Expiration Date: The latest expiration date across the volume set.

Expiration Value: Contains additional retention information derived from expiration date.

Expired by CA-1: Volume is listed to be cleaned.

Expired from Catalog Control: Volume has been expired by Catalog control.

Expired from CYCLE Control: Volume has been expired by CYCLE control.

Expired from LDATE Control: Volume has been expired by LDATE control.

External Data Manager ID: External data manager ID.

External Data Manager In Use: Controlled by an external data manager.

File On OS Catalog: File on OS catalog.

Hold Status: Reason code indicating why the tape is being held.

Internal Field Changed by User: Internal field changed by user.

Label Type: Label type.

Last Clean Date: The last clean date.

Last Cleaning Use Count: Use count since last cleaning.

Last Used Date: Date when dataset was last read from or written to the volume or set.

Last Used Job Name: The name of the job to last use the volume or set.

Last Used Program Name: The name of the program to last use the volume or set.

Last Used Unit Address: The device number of the drive which last mounted the tape.

LDATE Retention Period: Retain for specified number of days after date on which tape was last used.

Listed to be Cleaned: Volume is listed to be cleaned.

Location: Offsite vault name.

Managed System: The SMS name of this system.

Media Recording: Media recording technique.

Next VOLSER of Dataset: Next volume for multi-volume dataset.

No Further Stacking Allowed: No further stacking is allowed.

Non-resident tape: Non-resident tape.

Number of DSN Blocks: Number of dataset name blocks.

Num of Tape Cleanings: Number of tape cleanings.

Num Opens Since First Use: The number of times the tape has been opened since first use.

Out of Area Date: Date when the tape was sent offsite.

Out of Area ID: Offsite vault name.

Permanent Retention: Retain the dataset permanently.

Perm Read Errors Since Clean: Number of permanent read errors since clean.

Perm Read Errors Since Init: Number of permanent read errors since initialization.

Perm Write Errors Since Clean: Number of permanent write errors since clean.

Perm Write Errors Since Init: Number of permanent write errors since initialization.

Previous VOLSER of Dataset: Previous volume for multi-volume dataset.

Robotic Device Indicator: Robotic device indicator.

Robotic Tape Type: Robotic tape type.

Scratch Status: The scratch status.

SMS Management Class: The SMS Management class.

Tape Expired By SMS Retention: Tape Expired By SMS MAX Retention Rules.

Tape Has Been Degaussed: Tape has been degaussed.

Tape Has Been Used By RTS: Tape has been used by RTS.

Tape Is A 3592 WORM Cartridge: Tape is a 3592 WORM cartridge.

Tape Is In Use For RTS: Tape is in use for RTS.

Tape Released By Vault Manager: Tape released by external vault manager.

Temp Dataset: Temp dataset.

Temp Read Errors Since Clean: Number of temporary read errors since clean.

Temp Read Errors Since Init: Number of temporary read errors since initialization.

Temp Write Errors Since Clean: Number of temporary write errors since clean.

Temp Write Errors Since Init: Number of temporary write errors since initialization.

TMC Flag6: TMC Flag6.

Total Errors: Total number of errors.

Total Perm Errors: Total number of permanent errors.

Total Perm Read Errors: Total number of permanent read errors (since init or clean).

Total Perm Write Errors: Total number of permanent write errors (since init or clean).

Total Read Errors: Total number of read errors.

Total Temp Errors: Total number of temporary errors.

Total Temp Read Errors: Total number of temporary read errors (since init or clean).

Total Temp Write Errors: Total number of temporary write errors (since init or clean).

Total Write Errors: Total number of write errors.

User ID: ID of last user to update the record.

User Job Accounting Area: The user job accounting area.

Vault Slot Number: Offsite vault slot number.

Vault Specific Request: Vault specific request.

Vendor: Tape vendor.

Virtual Vol On Physical Vol: Virtual volume residing on physical volume.

Volume Is Owned By Agent: Volume is owned by agent.

Volume Mask: VOLSER mask used by search.

Volume Opened for Output: Volume opened for output.

Volume Percent Full: Percentage of this volume that contains data.

Volume Record Updated By User: Volume record updated by user.

Volume Sequence Number: Volume sequence number for multi-volume datasets.

Volume Serial Number: Volume serial number.

Volume Set Count: Number of volumes for multi-volume datasets.

Volume Set ID: First volume in the volume set.

WORM World Wide ID: World Wide ID for 3592 WORM cartridges.

WORM Writer Mount Count: Write mount count for 3592 WORM cartridges.

WORM WWID And WMC Are Present: WORM WWID and WMC are present.

Channel Path

The Channel Path attributes provide channel configuration and performance information. Extended information is provided for FICON® channels.

Bus Percent Utilization: For FICON® channels only, represents the percentage of the maximum available bus cycles being used in this collection interval.

Complex MB Second Read: For FICON® channels only, represents the number of MB of data transferred for read requests by this channel path for all LPARs in this complex.

Complex MB Second Written: For FICON® channels only, represents the average number of MB written by this channel path per second during this interval for the complex.

Complex Percent Read Units: For FICON® channels only, represents the ratio of total read units for this channel, for the complex, divided by the maximum read units for this channel.

Complex Percent Utilized: The ratio of total channel busy time divided by elapsed time for this interval.

Complex Percent Write Units: For FICON® channels only, represents the ratio of the number of write units for the complex divided by the maximum write units for this channel.

CPMF Group Number: The number of the Dynamic Channel Path Management Group to which this channel path is assigned.

DCM Status: The Dynamic Channel Path Management Status. DCM allows you to dynamically assign channels for controlling units to respond to peaks in demand.

Description: The type of channel connection. The values are as follows:

Unknown=0X00
Parallel_Block_Multiplex=0X01
Parallel_Byte_Multiplex=0X02
ESCON_Point_to_Point=0X03
ESCON_Switched_or_Point_to_Point=0X04
ESCON_Switched_Point_to_Point=0X05
ESCON_Path_to_a_Block_Converter=0X06
Native_Interface=0X07
CTC_Point_to_Point=0X08
CTC_Switched_Point_to_Point=0X09
CTC_Switched_or_Point_to_Point=0X0A
Coupling_Facility_Sender=0X0B
Coupling_Facility_Receiver=0X0C
Unknown=0X0D Unknown=0X0E
ESCON_Path_to_a_Byte_Converter=0X0F
OSA_EXPRESS=0X10
OSA_DIRECT_EXPRESS=0X11
Open_Systems_Adapter=0X12
Internal_System_Device=0X13
Unknown=0X14
Unknown=0X15
Cluster_Bus_Sender=0X16
Cluster_Bus_Receiver=0X17
Internal_Coupling_Sender=0X18
Internal_Coupling_Receiver=0X19
FICON_Point_to_Point=0X1A
FICON_Switched=0X1B
FICON_to_ESCON_Bridge=0X1C
FICON_Incomplete=0X1D
Direct_System_Device=0X1E
Emulated_I/O=0X1F
Unknown=0X20
Integrated_Cluster_Bus_Peer=0X21
Coupling_Facility_Peer=0X22
Internal_Coupling_Peer=0X23
Internal_Queued_Direct_Comm=0X24
FCP_Channel=0X25
Coupling_Over_Infiniband=0X26
OSA_ZBX_Data=0X30
OSA_ZBX_Management=0X31
High_Performance_FICON=0X99

LPAR MB Second Read: For FICON® channels only, represents the number of MB of data transferred for read requests by this channel path for this LPAR.

LPAR MB Second Written: For FICON® channels only, is the average number of MB written by this channel path per second during this interval for this LPAR.

LPAR Percent Read Units: For FICON® channels only, is the ratio of total read units for this LPAR, divided by the maximum read units for this channel.

LPAR Percent Utilized: The ratio of channel path busy time for this LPAR divided by elapsed time for this interval.

LPAR Percent Write Units: For FICON® channels only, is the ratio of total write units for this channel, for this LPAR, divided by the maximum write units for this channel.

Managed System: The SMS name of this system.

NonSMS Volumes: Represents the number of DASD volumes accessible through this channel path that are not managed by SMS.

Path ID Hex: The two-digit hexadecimal number assigned by HCD identifying this channel path to the system.

Path ID: The two-digit hexadecimal identifier for this channel path assigned by HCD.

Shared Path: Denotes whether this path is shared with other LPARs (yes), or is dedicated to this LPAR (no).

SMS Volumes: The number of SMS-controlled DASD volumes accessible through this channel path.

Status: Denotes the availability state of this channel path.

Time Stamp: The date and time this data was reported. Primarily for use in historical reporting.

Total Volumes: Represents the total number of DASD volumes accessible by this channel path.

Type: Indicates the path type, including Unknown, Block Channel, Byte Channel, ESCON Director Port, ESCON Channel, ESCON Converter, RMF Undetermined Type, and Channel-To-Channel Adaptor.

zHPF Requests Deferred per Second: The rate of zHPF requests attempted by the channel per second that were deferred due to unavailable resources. This will be zero when zHPF is not available.

zHPF Requests per Second: The rate of zHPF requests attempted by the channel per second. This will be zero when zHPF is not available.

Cross System Volume Summary

The Cross System Volume attributes provide an overview of the performance of a logical device across multiple z/OS ® images, that is, it helps to identify problems with shared DASD.

Control Unit Model: The model number associated with this control unit.

Cumulative I/O Rate: Represents the sum of I/O rates for all sharing systems.

Device Contention Index: The percentage of time that I/Os to the device are delayed (pending), averaged for all sharing systems. This is a measure of the time that each system spends waiting for the device.

Device Model: The model of DASD. Device geometry for the most common models or their equivalents are as follows:

Model	Tracks	Cylinders	Paths	Density
3390-1	16685	1113	4	TRIPLE
3390-2	33390	2226	4	TRIPLE
3390-3	50055	3339	4	TRIPLE
3380-K	39825	2655	4	TRIPLE
3380-J	13275	885	4	SINGLE
3380-E	26550	1770	2	DUAL
3380-A	13275	885	2	SINGLE
3375	11508	959	2	SINGLE
3350	16650	555	2	SINGLE

High Disconnect Time System: The SMFID of the sharing system reporting the highest disconnect time for this volume.

High Disconnect Time Value: The highest disconnect time reported by any sharing system.

High MSR System: The SMFID of the sharing system reporting the highest MSR (millisecond response time) value for this volume.

High MSR Value: Represents the highest MSR (millisecond response time) time reported by any sharing system.

Managed System: The SMS name of this system.

Storage Group Name: The 32-character name of the SMS storage group to which this volume belongs.

True % Busy: The percentage of time a device is busy as a result of activity from all sharing systems. This is a measure of the time that the device spends actually doing work.

User DASD Grp Name: The name of the user DASD group to which this volume belongs. User DASD group is a facility that enables you to define groups of DASD volumes by VOLSER and device address. You can also specify generic VOLSERs and devices (patterns) by using the wildcard character.

This feature gives you the flexibility of reporting the performance and status of any combination of DASD volumes as a single entity. User DASD groups can include both SMS managed volumes and non-SMS managed volumes.

User_Group\Name: The user-assigned name for a group of DASD volumes.

Volume: The six-character identifier assigned to a DASD or tape volume. Also known as the volume serial number or volume ID.

Cross System Volumes

The Cross System Volumes attributes provide performance information for a logical volume on each z/OS® image on which it is shared.

Busy Percent: The percentage of channel path usage. For FICON® channels, this is the percentage of the maximum work units being used.

Connect Time: A part of I/O during which data is located on the track and transferred between the DASD device and central storage.

Control Unit Model: The model number of the control unit for this hard disk drive.

Device Address: The MVS™ device number.

Device Model: The model of the DASD device. Attributes of the most common models or their equivalents are as follows:

MODEL	TRACKS	CYLINDERS	PATHS	DENSITY
3390-1	16685	1113	4	TRIPLE
3390-2	33390	2226	4	TRIPLE
3390-3	50055	3339	4	TRIPLE
3390-9	982800	65520	4	TRIPLE
3390-A	3940320	262668	4	TRIPLE
3380-K	39825	2655	4	TRIPLE
3380-J	13275	885	4	SINGLE
3380-E	26550	1770	2	DUAL
3380-A	13275	885	2	SINGLE
3375	11508	959	2	SINGLE
3350	16650	555	2	SINGLE

Disconnect Time: The time that I/O spends, free from the channel, searching for or requesting data. This time includes moving the device to the requested cylinder, waiting for the record to rotate under the head (rotational delay), and waiting for the channel to become available.

I/O Per Second: The average number of I/O operations per second.

IOSQ Delay: The average time that an I/O waits because the device is already busy (signified by the UCBBUSY BIT of the device being on).

Managed System: The SMS name of this system.

Pend Time: The time that I/O is delayed in the path to the device. Pending time can be attributable to the channel or control unit path being busy.

Response Time: The amount of time in milliseconds required for an I/O request to complete to a given device. MSR (millisecond response time) is divided into four components: IOS Queue, Pending, Disconnect, and Connect.

Storage Group Name: The 32-character name of the SMS storage group to which this volume belongs.

System: The systems management facility ID of the system that reported this row of data.

Volume: The six-character name assigned to DASD or tape volume. Also known as volume serial number of volume ID.

DASD Group Vol Perform

The DASD Group Vol Perform attributes provide information about the volumes and associated metrics that are in the DASD Group Vol Perform workspace.

Average Command Response Delay: The average number of milliseconds that a successfully initiated start or resume function needs until the first command is indicated as accepted by the device over the RMF interval.

Average HyperPAV Alias Count: The average number of alias devices assigned to a base volume in a HyperPAV environment over the RMF interval.

Busy Percent: The percentage of time a resource is busy processing I/O requests. For Parallel Access Volumes, this value is normalized by dividing by the number of PAV exposures.

Channels: The Channel Path IDs (in hexadecimal) through which this device may be accessed.

CHPID SEL: The CHPID selector is used as an index to limit the range of physical volumes for which data is requested.

Connect Time: The average time required for seek and transfer. This is the time it takes for data to be located on the track and transferred between the DASD device and central storage.

Control Unit Busy Delay Time: The average amount of I/O delay (in milliseconds) to this device caused by a control unit busy condition. This is the time it takes for data to be located on the track and transferred between the DASD device and central storage.

CU Model: The model indicator for this cache control unit.

Current PAV Exposures: The current number of UCBs through which the device may be accessed.

DCBs Open: The number of open datasets on this device.

Device Address: Four hexadecimal digits that define the physical address of this device. Also known as unit or unit address.

Device Busy Delay Time: The average amount of I/O delay (in milliseconds) to this device caused by a device busy condition.

Device Hex Address: The device address (in hexadecimal).

Device Model: The model of the DASD device.

Device MPL: The device multi-processing level, which is the I/O density for a device. This metric is the product of the millisecond response time (MSR) and I/O rate per second for the device. For a PAV volume, this value is normalized by dividing by the number of exposures.

Device Type: The MVS device type (derived from UCBTBYT4).

Director Busy Delay Time: The average amount of I/O delay (in milliseconds) to this device caused by a director port busy condition.

Disconnect Time: The average time, in milliseconds that I/O waited because the device was not connected to the channel. This time includes moving the device to the requested cylinder, waiting for the record to rotate under the head (rotational delay), and waiting for the channel to become available.

Has Catalog: Indicates that this DASD volume contains a catalog.

Hardware serial number: Hardware serial number is composed of the manufacturing plant identifier and a sequence number.

Interrupt Delay Time: The average interrupt delay time for an I/O. This is the elapsed time from the completion of an I/O until z/OS issues test subchannel to retrieve the results.

I/O Count: The actual number of I/O requests issued for this device during this interval.

I/O Per Second: The average number of I/O operations per second.

IOSQ Delay: The average time that an I/O waits because the device is already busy. A device is determined to be busy if its UCBBUSY bit is on.

LCU Number: The logical control unit number.

LCU Number Hex: The logical control unit number for this device (hexadecimal).

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Maximum PAV Exposures: The maximum number of UCBs through which the device could be accessed.

MSR Connect Time Percent: The percentage of the device response time during which the device was connected. This value is obtained be dividing the device pending time by the total millisecond response time.

MVS Status: The status of this device in the local MVS system.

PAV Exposure Changed: Indicates whether the number of PAV exposures for the volume has changed over the interval.

Pend Time: The average time that I/O is delayed in the path to the device. Pending time may be attributable to the channel or control unit path being busy.

Physical Device: The type of physical device containing this logical device.

RAID Model: The RAID type indicator (Symmetrix, Iceberg, RAMAC OR 2105).

Reserved Percent: The percentage of time during the interval when a shared device was reserved by this system.

Response Time: The total hardware service time and the front-end software queuing time required for the average I/O request to the device.

SMFID: The SMF ID of the system from which this information was obtained.

SMS Conversion Status: The SMS conversion status of this volume (non-SMS, initial, converted, or excluded).

SMS Status: The SMS volume status for local system (enabled, quiesced_all, quiesced_new, disabled_all, or disabled_new).

Solid State Device: An indication of whether the logical volume is backed by a solid state physical device.

Storage Group Name: The name of the SMS storage group to which this volume belongs.

Subsystem ID: A two-byte unique value used to identify a 3990 or equivalent storage control unit.

Time Stamp: The date and time this data was reported. Primarily for use in historical reporting.

User Group Name: The user-assigned name for a group of DASD volumes.

Volume: The volume serial number of a disk device currently in use by this application.

VTOC Index Status: The type of VTOC present on this device.

DASD Group Vol Space

The DASD Group Vol Space attributes provide information about the volumes and associated metrics that are in the DASD Group Vol Space workspace.

Allocated Space MB: The amount of space in MB on this volume that is currently allocated to datasets.

Cu Model: The model indicator for this device's control unit.

Cylinder Managed Space: The space on the volume, beginning at cylinder address 65520, is managed only in multicylinder units.

Device Address: Four hexadecimal digits that define the physical address of this device. Also known as unit or unit address.

Device Model: The model indicator for this device.

Device Type: The device model. Possible values include 3380 and 3390.

Extended Address Volume: The volume has more than 65520 cylinders and is an Extended Address Volume.

Fragmentation Index: A measure of the degree to which free space on the volume is fragmented. The higher the number, the more severe the fragmentation.

Free DSCBs: The number of DSCBs in the VTOC of this volume that are unused.

Free Extents: The total number of free areas on this volume.

Free Space MB: Amount of free space on this volume in MB.

Free VTOC Index Records: The number of index records in the VTOC index that are available for use.

Has Catalog: Indicates that this DASD volume contains a catalog.

Hardware serial number: Hardware serial number is composed of the manufacturing plant identifier and a sequence number.

Largest Free Extent MB: The size in MB of the largest contiguous free area on the volume.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Maximum Free Cylinders: The largest contiguous area of free cylinders on this device.

Maximum Free Tracks: The number of additional free tracks in the extent that contains the Maximum Free Cylinders.

Percent Free Space: Portion of this volume that is currently unused, expressed as a percentage.

RAID Model: The RAID type indicator (Symmetrix, Iceberg, RAMAC OR 2105).

SMFID: The SMF ID of the system from which this information was obtained.

SMS Conversion Status: SMS conversion status of this volume (non-SMS, initial, converted, or excluded).

SMS Status: The SMS status of this volume (enabled, quiesced_all, quiesced_new, disabled_all, or disabled_new).

Solid State Device: An indication of whether the logical volume is backed by a solid state physical device.

Storage Group Name: The name of the SMS storage group to which this volume belongs.

Time Stamp: The name of the user-defined group of DASD volumes to which this volume belongs.

Total Capacity MB: The data capacity of this volume in MB.

Total Free Cylinders: The total number of unallocated whole cylinders on this volume.

Total Free Tracks: The total number of free tracks that are not contained the total number of free whole cylinders.

Track Managed Free Cylinders: The total number of unallocated whole cylinders in the track managed area on this volume.

Track Managed Free Extents: The total number of free extents in the track managed area on this volume.

Track Managed Free Space: Amount of free space in the track managed area on this volume in MB.

Track Managed Free Tracks: The total number of free tracks in the track managed area that are not contained in the total number of track managed area free whole cylinders.

Track Managed Fragmentation Index: A measure of the degree to which free space in the track managed area on the volume is fragmented. The higher the number, the more severe the track managed area fragmentation.

Track Managed Maximum Free Cylinders: The largest contiguous area of free cylinders in the track managed area on this device.

Track Managed Maximum Free Tracks: The number of additional free tracks in the extent that contains the maximum free cylinders.

Track Managed Largest Free Extent: The size in MB of the largest contiguous free area in the track managed area on the volume.

Track Managed Percent Free: The portion of the track managed area on this volume that is currently unused, expressed as a percentage.

Track Managed Total Capacity: Data capacity of the track managed area on this volume in MB.

User Group Name: The name of the user-defined group of DASD volumes to which this volume belongs.

Volume: The volume serial number of a disk device currently in use by this application.

Total Capacity GB: The data capacity of this volume in GB.

Free Space GB: Amount of free space on this volume in GB.

Total Used GB: Amount of space on this volume in GB, which is currently allocated to datasets.

DUPLEX_INF: PPRC State.

VTOC Index Status: The status of the VTOC index on this volume (enabled, disabled, error, OS-VTOC, or exclude).

DASD Physical Group

The DASD Physical Group attributes provide performance information for a physical volume group. A physical volume group is the set of logical volumes that are related by a hardware specific construct.

Average Command Response Delay: The average number of milliseconds that a successfully initiated start or resume function needs until the first command is indicated as accepted by the device over the RMF™ interval.

Average HyperPAV Alias Count: The average number of aliased devices assigned to a base volume in a HyperPAV environment over the RMF™ interval.

Busy Percent: The percentage of time a resource is busy processing I/O requests. For PAV volumes, this value is normalized by dividing the number of PAV exposures.

Bypass Cache Percentage: The percentage of I/Os that bypassed cache this interval. The I/O request is sent directly to DASD. Tracks are neither read from cache nor promoted to cache. Also, their entries in the LRU list are not updated.

Cache FW % Read: The percentage of all I/Os that were read operations for data in cache that originated from a cache fast write during this interval. Cache fast write is a form of fast write where the data is written directly to cache without using nonvolatile storage and is available for later destaging.

Cache FW % Write: The percentage of all I/Os that were cache fast writes during this interval. Cache fast write is a form of fast write where the data is written directly to cache without using nonvolatile storage and is available for later destaging.

Cache Status: The caching status for this volume. Active means caching for this volume is active and inactive means caching for this volume is inactive.

Cache to DASD Total Tracks: The number of tracks written from cache to physical disk during the cache interval. This value is captured asynchronously with respect to the transfer of data from the channel to the cache.

Cache to DASD Tracks per Sec: The rate per second of tracks written from cache to physical disk during the cache interval. This value is captured asynchronously with respect to the transfer of data from the channel to the cache.

Connect Time: Part of I/O during which data is located on the track and transferred between the DASD device and central storage.

Control Unit Model: The model number associated with the control unit to which this device is attached.

Current PAV Exposures: The current number of UCBs through which the device can be accessed.

DASD to Cache Non Promotes Second: The rate per second of operations that did not promote (copy) a track to cache because free cache segments were not available.

DASD to Cache Non-promotes: The number of operations that did not promote (copy) a track to cache because free cache segments were not available.

DASD to Cache Total Tracks Non-sequential: The number of tracks read from physical disk to cache over the cache interval for non-sequential I/O operations.

DASD to Cache Total Tracks Sequential: The number of tracks read from physical disk to cache over the cache interval for sequential I/O operations.

DASD to Cache Total Tracks: The number of tracks read from physical disk to cache over the cache interval for I/O operations.

DASD to Cache Tracks per Sec Non-sequential: The rate per second of tracks read from physical disk to cache over the cache interval for non-sequential I/O operations.

DASD to Cache Tracks per Sec Sequential: The rate per second of tracks read from physical disk to cache over the cache interval for sequential I/O operations.

DASD to Cache Tracks per Sec: The rate per second of tracks read from physical disk to cache over the cache interval for I/O operations.

DCBs Open: The count of open datasets.

Device Address: The MVS™ device number assigned by HCD and associated with this device.

Device MPL: The device multi-processing level, which is the I/O density for a device. This metric is the product of the MSR (millisecond response time) and I/O per second for a device.

DFW Hit Percent: The percentage of all DASD I/Os to this volume that were DASD fast writes during this interval. DASD fast write is a form of fast write to cache where the data is written concurrently to cache and nonvolatile storage and automatically scheduled for destaging to DASD.

DFW Retry Percent: The percentage of DASD I/O operations that required retry during this interval.

DFW Status: The status of the DASD fast write option for this volume. Value can be Active, Inactive, Unknown or Timeout.

Disconnect Time: The time that I/O spends, free from the channel searching for or requesting data.

Interrupt Delay Time: The average interrupt delay time for an I/O. This is the elapsed time from the completion of an I/O until z/OS issues test subchannel to retrieve the results.

I/O Count: The actual number of I/O requests issued for this device during this interval.

I/O per Second: The average number of I/O operations per second.

Inhibit Cache Percent: The percentage of I/O operations that inhibited the loading of data into cache. If the data is already in cache, the request is satisfied from cache.

IOSQ Delay: The average time that an I/O waits because the device is already busy. A device is determined to be busy if its UCBBUSY bit is on.

LCU Contention Rate: The average number of requests delayed due to LCU contention. LCU contention is the rate per second at which the SAP places delayed I/O requests on the CU-HDR for this LCU. This calculation is performed when all paths to the subchannel are busy and at least one path to the control unit is busy. For devices with only one path, or devices where multiple paths exist and the busy condition is immediately resolved, the IOP does not count the condition.

LCU Number: The three hexadecimal digit logical control unit number.

Managed System: The SMS name of this system.

Maximum PAV Exposures: The maximum number of UCBs through which the device can be accessed in this interval.

MSR Connect Time Percent: The percentage of MSR (millisecond response time) during which data is transferred (SEARCH + TRANSFER).

PAV Exposure Changed: Indicates that the number of PAV exposures to this device changed during this interval.

Pend Time: The time that I/O is delayed in the path to the device. Pending time can be attributable to the channel or control unit path being busy.

Read Hit Percent: The percentage of read hits for this volume.

Read I/O Count: The number of read operations directed to this device that completed during this interval.

Read I/O Percent: The percentage of I/O requests to this volume that were read requests.

Reserved Percent: The percentage of time during the interval when a shared device was reserved by this system.

Response Time: The amount of time in milliseconds it takes for an I/O request to complete to a given device. MSR (millisecond response time) is divided into four components: IOS Queue, Pending, Disconnect and Connect.

Selected Volser: The volume serial number associated with the physical volume selected for display.

SMS Status: The SMS volume status for the local system. Values are: Enabled, Quiesced_All, Quiesced_New, Disabled_All or Disabled_New.

Storage Group Name: The 32-character name of the SMS storage group to which this volume belongs.

Subsystem ID: The MVS™ identifier of the logical control unit to which this device is attached.

Volume: The VOLSER associated with this physical volume.

VTOC Index Status: The status of the VTOC index on this volume. Values are Enabled, Disabled, Error, OS_VTOC or Excluded.

Write Hit Percent: The percentage of I/O operations that completed without accessing the hard disk drive.

Write I/O Count: The number of write operations directed to this device that completed during this interval.

DASD Volume Performance

The DASD Volume Performance attributes provide status and performance information for logical volumes.

Average Command Response Delay: The average number of milliseconds that a successfully initiated start or resume function needs until the first command is indicated as accepted by the device over the RMF™ interval.

Average HyperPAV Alias Count: The average number of aliased devices assigned to a base volume in a HyperPAV environment over the RMF™ interval.

Busy Percent: The percentage of time a resource is busy processing I/O requests. For Parallel Access Volumes, this value is normalized by dividing by the number of PAV exposures.

Channel Path ID 1: The channel path identifier (in decimal) of one channel path connecting this device.

Channel Path ID 2: The channel path identifier (in decimal) of one channel path connecting this device.

Channel Path ID 3: The channel path identifier (in decimal) of one channel path connecting this device.

Channel Path ID 4: The channel path identifier (in decimal) of one channel path connecting this device.

Channel Path ID 5: The channel path identifier (in decimal) of one channel path connecting this device.

Channel Path ID 6: The channel path identifier (in decimal) of one channel path connecting this device.

Channel Path ID 7: The channel path identifier (in decimal) of one channel path connecting this device.

Channel Path ID 8: The channel path identifier (in decimal) of one channel path connecting this device.

Channels: The channel Path IDs (in hexadecimal) through which this device can be accessed.

CHPID SEL The CHPID selector is used as an index to limit the range of physical volumes for which data is requested.

Connect Time: The average time (in milliseconds) required for seek and transfer. This is the time it takes for data to be located on the track and transferred between the DASD device and central storage.

Control Unit Busy Delay Time: The average amount of I/O delay (in milliseconds) to this device caused by a control unit busy condition.

CU Model: The model indicator for this cache control unit.

Current PAV Exposures: The current number of UCBs through which the device can be accessed.

DCBs Open: The number of open datasets on this device.

Device Address: The MVS™ device number associated with this device.

Device Busy Delay Time: The average amount of I/O delay (in milliseconds) to this device caused by a device busy condition.

Device Hex Address: The device address in hexadecimal numbers.

Device Model: The MVS™ device model.

Device MPL: The device Multiprocessing Level, which is the I/O density for a device. This metric is the product of the MSR (millisecond response time) and I/O rate per second for the device. For a PAV volume, this value is normalized by dividing by the number of exposures.

Device Type: The MVS™ device type.

Director Busy Delay Time: The average amount of I/O delay (in milliseconds) to this device caused by a director port busy condition.

Disconnect Time: The average time in milliseconds that I/O waited because the device was not connected to the channel. This time includes moving the device to the requested cylinder, waiting for the record to rotate under the head (rotational delay), and waiting for the channel to become available.

Has Catalog: Indicates that this DASD volume contains a catalog.

Hardware serial number: Hardware serial number is composed of the manufacturing plant identifier and a sequence number.

Interrupt Delay Time: The average interrupt delay time for an I/O. This is the elapsed time from the completion of an I/O until z/OS issues test subchannel to retrieve the results.

I/O Count: The actual number of I/O requests issued to the device over the Resource Monitoring Facility (RMF™) interval.

I/O Per Second: The average number of I/O operations per second.

IOSQ Delay: The average time (in milliseconds) that an I/O waits because the device is already busy. A device is determined to be busy if its UCBBUSY bit is on.

LCU Number Hex: The hexadecimal LCU number for this device.

LCU Number: The hexadecimal logical control unit number.

Managed System: The SMS name of this system.

Maximum PAV Exposures: The maximum number of UCBs through which the device can be accessed in this interval.

MSR Connect Time Percent: The percentage of the device response time during which the device was connected. This value is obtained be dividing the device pending time (in milliseconds) by the total millisecond response time.

MVS™ Status: The status of this device in the local MVS™ system. Values are Online, Offline, Pending Offline, Boxed, Not Ready, and Long Busy.

PAV Exposure Changed: Indicates whether the number of PAV exposures for the volume has changed over the interval.

Pend Time: The average time (in milliseconds) that I/O is delayed in the path to the device. Pending time can be attributable to the channel or control unit path being busy.

Physical Device: The type of physical device containing this logical device (supported for the following devices only: Shark 2105, Symmetrix, TDS). Values are 2105, Symmetrix, and TDS.

Raid Model: The raid type indicator. Values are Symmetrix, Iceberg, Ramac, or 2105.

Reserved Percent: The percentage of time during the interval when a shared device is reserved by this system.

Response Time: The total hardware service time and the front end software queuing time (in milliseconds) involved for the average I/O request to the device.

Solid State Device: Indicates whether the logical volume is backed by a solid state physical device.

SMFID: The SMF ID of the system from which this information was obtained.

SMS Conversion Status: The SMS conversion status of this volume. Values are Non-SMS, Initial, Converted, or Excluded.

SMS Status: The SMS volume status for a local system. Values are Not Connected, Enabled, Quiesced_All, Quiesced_New, Disabled_All, or Disabled_New.

Storage Group Name: The SMS storage group name to which this volume belongs.

Subsystem ID Hex: The subsystem ID (for Model 3990 and 2150 control units), expressed in hexadecimal.

Subsystem ID Index: The two-byte unique value used to identify a 3990 or equivalent storage control unit.

Time Stamp: The date and time associated with this row. Primarily for use in historical reporting.

User Group Index: The user-assigned name for a group of DASD volumes.

Volume: The volume serial number of a disk device currently in use by this application.

VTOC Index Status: The type of VTOC present on this device. Values are Enabled, Disabled, Error, OS VTOC, and Excluded.

DASD Volume Space

The DASD Volume Space attributes provide status and space information for logical volumes.

Allocated Space MB: The amount of space in MB on this volume which is allocated to datasets.

Cylinder Managed Space: Space on the volume, beginning at cylinder address 65520, is managed only in multicylinder units.

Device Address: The MVS™ device number assigned by HCD, associated with this device.

Device Model: The actual device model. Each device model has unique space and geometry attributes.

Device Hex Address: The device address in hexadecimal.

Device Type: The MVS™ device type. For example: 3380, 3390, or 2105.

Extended Address Volume: Volume has more than 65520 cylinders and is an Extended Address Volume.

Fragmentation Index: A number between 0 and 1000 that indicates the degree of scattered free space on the volume. The higher the number, the higher the amount of fragmentation.

Free DSCBs: The number of DSCBs in the VTOC of this volume that are unused.

Free Extents: The amount of free areas on the volume.

Free Space MB: Represents the amount of free space on this volume in MB.

Free VTOC Index Records: The number of index records in the VTOC index that are available for use.

Hardware serial number: Hardware serial number is composed of the manufacturing plant identifier and a sequence number.

Has Catalog: Indicates that this DASD volume contains a catalog.

Largest Free Extent MB: The size in MB of the largest contiguous free area on the volume.

Managed System: The SMS name of this system.

Maximum Free Cylinders: The largest contiguous area of free cylinders on this device.

Maximum Free Tracks: The number of tracks in the largest free extent.

Percent Free Space: The ratio of free space divided by the total space for this volume expressed as a percentage.

RAID Model: Indicates the RAID type, including Symmetrix, Iceberg, RAMAC and 2105.

Solid State Device: Indicates whether the logical volume is backed by a solid state physical device.

SMS Conversion Status: The SMS conversion status of this volume. Values are: Non-SMS, Initial, Converted, or Excluded.

SMS Status: The SMS volume status for the local system. Values are: Enabled, Quiesced_All, Quiesced_New, Disabled_All, or Disabled_New.

Storage Group Name: The 32-character name of the SMS storage group to which this volume belongs.

Time Stamp: Represents the date and time associated with this row.

Total Capacity MB: The data capacity of this volume in MB.

Total Free Cylinders: Represents the total number of unallocated cylinders on this volume.

Total Free Tracks: The total number of separated areas of one track or larger on this volume that are available to allocate.

Track Managed Fragmentation Index: A measure of the degree to which free space in the track managed area on the volume is fragmented. The higher the number, the more severe the track managed area fragmentation.

Track Managed Free Cylinders: Total number of unallocated whole cylinders in the track managed area on this volume.

Track Managed Free Extents: Total number of free extents in the track managed area on this volume.

Track Managed Free Space: Amount of free space in the track managed area on this volume in MB.

Track Managed Free Tracks: The total number of free tracks in the track managed area which are not contained in the total number of track managed area free whole cylinders.

Track Managed Largest Free Extent: The size in MB of the largest contiguous free area in the track managed area on the volume.

Track Managed Maximum Free Cylinders: Largest contiguous area of free cylinders in the track managed area on this device.

Track Managed Maximum Free Tracks: Number of additional free tracks in the extent which contains the Maximum Free Cylinders.

Track Managed Percent Free: Portion of the track managed area on this volume which is currently unused, expressed as a percentage.

Track Managed Total Capacity: Data capacity of the track managed area on this volume in MB.

User Group Index: The name of the user-defined group of the DASD volume to which this volume belongs.

Volume: The six-character identifier associated with this logical device.

Total Capacity GB: The data capacity of this volume in GB.

Free Space GB: Amount of free space on this volume in GB.

Total Used GB: Amount of space on this volume in GB, which is currently allocated to datasets.

DUPLEX_INF: PPRC State.

VTOC Index Status: The status of the VTOC index on this volume. Values are: Enabled, Disabled, Error, OS-VTOC, or Excluded.

Dataset Attributes Blocksize Summary

The Dataset Attributes Blocksize Summary attributes provide the count of the number of datasets that exist in each product-defined unique DCB BLKSIZE range. See "Dataset Attributes Database" on page 204 for more information.

Allocated GB: Total allocated GB for all datasets in the range.

Allocated MB: Total allocated MB for all datasets in the range.

Allocated Tracks: Total allocated tracks for all datasets in the range.

High Blocksize: The highest value in the blocksize range.

Low Blocksize: The lowest value in the blocksize range.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Number of Datasets: The number of datasets in the range.

Percent of Allocated Space: The percentage of total allocated space that datasets in this blocksize range represent.

Percent of Total Datasets: The percentage of total processed datasets that datasets in this blocksize range represent.

Timestamp: The time that the row was created.

Unused GB: Total unused GB for all datasets in the range.

Unused MB: Total unused MB for all datasets in the range.

Unused Tracks: Total unused tracks for all datasets in the range.

Used GB: Total used GB for all datasets in the range.

Used MB: Total used MB for all datasets in the range.

Used Tracks: Total used tracks for all datasets in the range.

Wasted GB: Total wasted GB for all datasets in this range.

Wasted MB: Total wasted MB for all datasets in this range.

Dataset Attributes CA Split Summary

The Dataset Attributes CA Split Summary attributes provide data on the Top N VSAM datasets in terms of total number of CA Splits. You define the value of n in the Dataset Attributes Data Collection dialog box. See "Dataset Attributes Database" on page 204 for more information.

Address: Address of detail information in the data space of the dataset.

ALET: ALET (Address List Entry Token) of information in the data space of the dataset.

Allocated GB: Total allocated GB for the dataset.

Allocated MB: Total allocated MB for the dataset.

Allocated Tracks: Total allocated tracks for the dataset.

Allocated Volumes: The number of volumes on which space has been allocated for the dataset.

Associated DSN: The name of this component's associated dataset.

CA Splits: Total number of Control Area splits.

CI Splits: Total number of Control Interval splits.

Dataset Name: The name of the dataset.

Extents: The number of extents.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Timestamp: The time that the row was created.

Unused GB: Total unused GB for the dataset.

Unused MB: Total unused MB for the dataset.

Unused Tracks: Total unused tracks for the dataset.

Used GB: Total used GB for the dataset.

Used MB: Total used MB for the dataset.

Used Tracks: Total used tracks for the dataset.

Volser: The VOLSER on which the dataset is located.

Dataset Attributes Catalog Summary

The Dataset Attributes Catalog Summary attributes provide the number of datasets cataloged in each system catalog. This data shows the distribution of datasets across catalogs and can allow rebalancing to improve performance.

When the Collect Catalog Data option is enabled, the catalog data is retrieved for each dataset encountered in the VTOC of the volume. See "Dataset Attributes Database" on page 204 for more information.

Allocated GB: Total allocated GB represented by all of the datasets in the catalog.

Allocated MB: Total allocated MB represented by all of the datasets in the catalog.

Allocated Tracks: Total allocated tracks represented by all of the datasets in the catalog.

Catalog Name: Name of the catalog.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Timestamp: The time that the row was created.

Total Datasets: The total number of processed datasets contained in the catalog.

> **Note:** Datasets that are cataloged, but do not exist on any processed volume are not counted. The datasets are not counted because the dataset name from the VTOC is the starting point for the catalog lookup.

Unused GB: Total unused GB represented by all of the datasets in the catalog.

Unused MB: Total unused MB represented by all of the datasets in the catalog.

Unused Tracks: Total unused tracks represented by all of the datasets in the catalog.

Used GB: Total used GB represented by all of the datasets in the catalog.

Used MB: Total used MB represented by all of the datasets in the catalog.

Used Tracks: Total used tracks represented by all of the datasets in the catalog.

Wasted GB: Total wasted GB for all datasets in this range.

Wasted MB: Total wasted MB for all datasets in this range.

Dataset Attributes CI Split Summary

The Dataset Attributes CI Split Summary attributes provide data on the Top N VSAM datasets in terms of total number of CI Splits. You define the value of n in the Dataset Attributes Data Collection dialog box. See "Dataset Attributes Database" on page 204 for more information.

Address: Address of detail information in the data space of the dataset.

ALET: ALET (Address List Entry Token) of information in the data space of the dataset.

Allocated GB: Total allocated GB for the dataset.

Allocated MB: Total allocated MB for the dataset.

Allocated Tracks: Total allocated tracks for the dataset.

Allocated Volumes: The number of volumes on which space has been allocated for the dataset.

Associated DSN: The name of this component's associated dataset.

CA Splits: Total number of Control Area splits.

CI Splits: Total number of Control Interval splits.

Dataset Name: The name of the dataset.

Extents: The number of extents.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Timestamp: The time that the row was created.

Unused GB: Total unused GB for the dataset.

Unused MB: Total unused MB for the dataset.

Unused Tracks: Total unused tracks for the dataset.

Used GB: Total used GB for the dataset.

Used MB: Total used MB for the dataset.

Used Tracks: Total used tracks for the dataset.

Volser: The VOLSER on which the dataset is located.

Dataset Attributes Creation Date Summary

The Dataset Attributes Creation Date Summary attributes provide the count of the number of datasets that exist in each product-defined unique Creation Date (in other words, number of days since creation) range. See "Dataset Attributes Database" on page 204 for more information.

Allocated GB: Total allocated GB for all datasets in the range.

Allocated MB: Total allocated MB for all datasets in the range.

Allocated Tracks: Total allocated tracks for all datasets in the range.

High Days Since Creation: The highest value in the creation date range.

Low Days Since Creation: The lowest value in the creation date range.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Number of Datasets: The number of datasets in the range.

Percent of Total Allocated Space: The percentage of total allocated space that datasets in this creation date represent.

Percent of Total Datasets: The percentage of total processed datasets that datasets in this creation date range represent.

Timestamp: The time that the row was created.

Unused GB: Total unused GB for all datasets in the range.

Unused MB: Total unused MB for all datasets in the range.

Unused Tracks: Total unused tracks for all datasets in the range.

Used GB: Total used GB for all datasets in the range.

Used MB: Total used MB for all datasets in the range.

Used Tracks: Total used tracks for all datasets in the range.

Wasted GB: Total wasted GB for all datasets in this range.

Wasted MB: Total wasted MB for all datasets in this range.

Dataset Attributes Defs

The Dataset Attributes Defs attributes provide information to other attribute groups. Users do not use these attributes. The attributes are documented for the purpose of background reference only. See "Dataset Attributes Database" on page 204 for more information.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Message: Message.

Dataset Attribute Definitions: Dataset Attribute bulk data.

Dataset Attributes Dataset Detail

The Dataset Attributes Dataset Detail attributes provide dataset information from the VTOC and catalog. The attributes in this group include details regarding control intervals (CIs) and control areas (CAs). See "Dataset Attributes Database" on page 204 for more information.

> **Note:** You can create queries and situations that generate more specific data than that provided in the default workspaces. Be aware that in a large enterprise the dataset database can be tracking millions of datasets. Estimate the potential size of the responses to your queries and situations to prevent them from requiring excessive processing, network, and disk space resources. Also be aware that a maximum of 16384 rows are returned on any single query against the Dataset Attribute Detail table.

% Free Bytes per CI: The VSAM percentage of free bytes in all CIs.

% Free CIs per CA: The VSAM percentage of free CIs in all CAs.

Allocated GB: The GB allocated to the dataset on the volume.

Allocated MB: The MB allocated to the dataset on the volume.

Allocated Tracks: The tracks allocated to the dataset on the volume.

Allocated Volumes: The number of volumes on which space is allocated for the dataset.

Associated Entry Name: The VSAM associated dataset name.

Average LRECL: The VSAM average LRECL.

BLKSIZE: The block size.

Bytes Saved By CA-Reclaim The number of bytes saved by using the CA-Reclaim feature.

CA Reclaim: CA Reclaim.

CA Reclaimed: Number of Reclaimed CAs.

CA Reused: Number of Reused CAs.

CA Splits: The VSAM number of CA splits.

Catalog Entry Type: The catalog entry type. The values are as follows:

Alternate Index
Cluster
Data Component
Generation Data Group
Generation Dataset
Index Component
Non-VSAM Dataset
Unknown
USERCAT

Catalog Name: The name of the catalog in which dataset is cataloged.

Cataloged: The dataset is cataloged.

CI Size: The VSAM control interval size.

CI Splits: The VSAM number of CI splits.

CIs per CA: The VSAM number of CIs per CA.

Collection Timestamp: Date and time of data collection for this row.

Component Free Space: Bytes of free space in a VSAM dataset component.

Compressed Size: The VSAM compressed user data size in bytes.

Compressed Size MB: The VSAM compressed user data size in MB.

Compressed Size GB: The VSAM compressed user data size in GB.

Compressible: This is a compressible dataset.

Creation Date: The creation date.

DAEAddress: Address of detail information in the data space of the dataset.

DAEALET: ALET of detail information in the data space of the dataset.

Data Component: The dataset is a VSAM data component.

Dataset Mask: The dataset mask of the dataset.

Dataset Name: The name of the dataset.

Days Since Creation: The number of days since creation.

Days Since Last Backup: The days since last backup.

Days Since Reference: The number of days since last reference.

Days Since/Until Expiration: The numbers days until expiration. This value is negative if expiration has passed.

Device Type: The device type.

DSORG: The dataset organization type. The values are as follows:

VS: VSAM
IS: ISAM
PS: Physical Sequential
DA: Direct Access
PO: Partitioned
GS: Graphics
CX: BTAM OR QTAM Line Group
CQ: QTAM DIRECT ACC MSG QUE
MQ: QTAM PROB PROG MSG QUE
TX: TCAM Line Group
TQ: TCAM Message Queue
AM: VSAM Dataset/Data Space
TR: TCAM 3705

??: Unknown

Encrypted: The dataset is encrypted.

Encryption Algorithm: The dataset encryption algorithm.

Encryption Key Length: The dataset encryption key length.

Encryption Mode: The dataset encryption mode.

Extended Address Space Eligible: Dataset is eligible to use Extended Address Space of an Extended Address Volume.

Erase: The VSAM erase attribute.

EXCPS: The VSAM number of EXCPs.

Expiration Date: The expiration date.

Extended Format: This is an extended format dataset.

Extents: The number of extents.

Free Bytes per CI: The VSAM number of free bytes in all CIs.

Free CIs per CA: The VSAM number of free CIs per CA.

GDG Attributes: GDG attributes.

GDG Last Altered Date: GDG last altered date.

GDG Level: GDG level.

Generation Dataset: This is a generation dataset.

HFS: This is an HFS dataset.

High Allocated RBA: The high allocated RBA for the dataset on this volume where RBA is relative block address.

High Allocated RBA String: The high allocated RBA for the dataset on this volume as a string where RBA is relative block address.

High Used RBA: The high used RBA for the dataset on this volume where RBA is relative block address.

High Used RBA String: The high used RBA for the dataset on this volume as a string where RBA is relative block address.

Imbed: The VSAM imbed attribute.

Index Component: The dataset is a VSAM index component.

Is a Catalog: The dataset itself is a catalog.

Key Label: The Key Label for encrypted dataset.

Key Length: The DASD hardware key length or, for keyed VSAM datasets, the keylen value.

Last Backup Date: The last backup date.

Log Stream ID: Value of LOGSTREAMID parameter set by IDCAMS DEFINE/ALTER.

LRECL: The logical record length.

Managed System: Internal node name used to identify the system from which this row of data was obtained.

Maximum Buffer Size: The VSAM maximum buffer size.

Maximum Datasets in GDG: For generation datasets, the maximum number of datasets in the GDG.

Maximum Record Size: The VSAM maximum record size.

Non-VSAM Attribute: The Non-VSAM attribute information. The predefined values are as follows:

n/a: 0
Simple non-VSAM dataset: S
Active GDS: H
Deferred GDS: N
Rolled-off GDS: M
PDSE: L
POSIX dataset: P
Active GDS-PDSE: I

Deferred GDS-PDSE: J
Rolled-Off GDS-PDSE: K
Unknown: U

Owner: The VSAM dataset owner.

Page Space: The cluster is a page space.

PDSE: This is a persistent data store extended (PDSE) dataset.

Percent Free: The percent of allocated space free. For VSAM datasets, this value is calculated using: (High Allocated RBA - High Used RBA) \ High Allocated RBA.

Percent Used: The percent of allocated space used. For VSAM datasets, this is calculated using (High Used RBA \ High Allocated RBA).

Primary GB: The primary space allocation (in other words, the first extent) in GB.

Primary MB: The primary space allocation (in other words, the first extent) in MB.

Primary Tracks: The primary space allocation (in other words, the first extent) in tracks.

RACF®: This dataset is RACF® protected. This attribute is based on the VTOC RACF® indicator, not on whether the dataset might be RACF-protected by default.

RECFM: The record format. The values are as follows:

F: Fixed Length
V: Variable Length
U: Undefined Length
T: Track Overflow
B: Blocked
S: Spanned/Standard
A: ANSI Control Character
M: Machine Control Character

Records Deleted: The VSAM number of deleted records.

Records Inserted: The VSAM number of inserted records.

Records Retrieved: The VSAM number of retrieved records.

Records Total: The VSAM number of logical records.

Records Updated: The VSAM number of updated records.

Recovery: The VSAM recovery attribute.

Reference Date: The last reference date.

Replicate: The VSAM Attribute: Replicate.

Reuse: The VSAM reuse attribute.

Secondary Allocation Quantity: The secondary allocation amount.

Secondary Allocation Record Units: The secondary space record allocation units.

Secondary Allocation Record Value: The secondary space allocation value when specified in Records.

Secondary Allocation Units: The space allocation units. The values are as follows:

CYLINDER/CYL
TRACK/TRK
ABSTRACK/ABS
AVG BLOCK/AVB
AVG BLOCK ROUNDED/AVD
AVG RECORD/AVR

Share Option: The VSAM share options (cross-region, cross-system). The values are as follows:

(1,3), (1,4),
(2,3), (2,4),
(3,3), (3,4),
(4,3), (4,4)

SMS Data Class: The SMS data class.

SMS Managed: The dataset is SMS managed.

SMS Management Class: The SMS management class.

SMS Storage Class: The SMS storage class.

SMS Storage Group: The SMS storage group.

Spanned: The VSAM spanned attribute.

Speed: The VSAM speed attribute.

Stripe Count: The stripe count.

Swap Space: The cluster is a swap space.

System Determined Blocksize: The block size as determined by the system.

Temporary Dataset: This is a temporary dataset.

Timestamp: The time that the row was created.

Unique: The VSAM unique attribute.

Unmovable: The dataset is unmovable.

Unused GB: The GB not used by the dataset on the volume.

Unused MB: The MB not used by the dataset on the volume.

Unused Tracks: The tracks unused by the dataset on the volume.

Used GB: The GB used by the dataset on the volume.

Used MB: The MB used by the dataset on the volume.

Used Tracks: The tracks used by the dataset on the volume.

Using Cylinder Managed Space: Dataset is using Cylinder Managed Space on an Extended Address Volume.

Volser: The volume on which the dataset is located.

Volume Count: The number of volumes on which the dataset is located.

Volume Mask: The volume mask of the dataset.

Volume Sequence: The volume sequence number for this portion of the dataset.

VSAM Dataset Organization: The VSAM dataset organization types. The values are as follows:

KSDS
ESDS
RRDS
LDS
AIX®

Wasted GB: Wasted GB for the dataset.

Wasted MB: Wasted MB for the dataset.

ZFS: The dataset is ZFS.

Dataset Attributes DSORG Detail

The Dataset Attributes DSORG Detail attributes provide monitoring data for the Top N datasets in terms of allocated space for a specific DSORG. You define the value of *n* in the Dataset Attributes Data Collection dialog box. See "Dataset Attributes Database" on page 204 for more information.

Address: Address of detail information in the data space of the dataset.

ALET: ALET (Address List Entry Token) of information in the data space of the dataset.

Allocated GB: Allocated GB for the dataset.

Allocated MB: Allocated MB for the dataset.

Allocated Tracks: Allocated tracks for the dataset.

Allocated Volumes: The number of volumes on which space has been allocated for the dataset.

BLKSIZE: The block size of the dataset.

Creation Date: The date dataset was created.

Dataset Name: The name of the dataset.

DSORG: The dataset organization.

Extents: The number of extents.

LRECL: The logical record length.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

RECFM: The record format.

Reference Date: The date dataset was last referenced.

Timestamp: The time that the row was created.

Unused GB: Unused GB for the dataset.

Unused MB: Unused MB for the dataset.

Unused Tracks: Unused tracks for the dataset.

Used GB: Used GB for the dataset.

Used MB: Used MB for the dataset.

Used Tracks: Used tracks for the dataset.

Volser: The VOLSER on which the dataset is located.

Dataset Attributes DSORG Summary

The Dataset Attributes DSORG Summary attributes provide summary counts for the number of datasets in each encountered DSORG. See "Dataset Attributes Database" on page 204 for more information.

Allocated GB: The allocated GB for all datasets with this DSORG.

Allocated MB: The allocated MB for all datasets with this DSORG.

Allocated Tracks: The allocated tracks for all datasets with this DSORG.

DSORG String: The dataset organization type for this dataset in character string format.

DSORG: The dataset organization type. The values are as follows:

Unknown
Hierarchical_File
PDS_Extended
Extended_Sequential
Extended_Seq_Unmovable
VSAM
VSAM_Unmovable
ISAM
ISAM_Unmovable
Physical_Sequential
Physical_Seq_Unmovable
Direct_Access
Direct_Access_Unmovable
Partitioned
Partitioned_Unmovable
Unknown

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Number of Datasets: The number of processed datasets within the indicated DSORG.

Percent of Allocated Space: The percentage of total allocated space for all DSORGs that datasets with this DSORG represent.

Percent of Total Datasets: The percentage of total datasets for all DSORGs that datasets with this DSORG represent.

Timestamp: The time that the row was created.

Unused GB: Total unused GB for datasets with this DSORG.

Unused MB: Total unused MB for datasets with this DSORG.

Unused Tracks: Total unused tracks for datasets with this DSORG.

Used GB: Total used GB for datasets with this DSORG.

Used MB: Total used MB for datasets with this DSORG.

Used Tracks: Total used tracks for datasets with this DSORG.

Wasted GB: Total wasted GB for all datasets with this DSORG.

Wasted MB: Total wasted MB for all datasets with this DSORG.

Dataset Attributes Extent Summary

The Dataset Attributes Extent Summary attributes provide data on the Top N largest datasets in terms of total number of extents. You define the value of n in the Dataset Attributes Data Collection dialog box. See "Dataset Attributes Database" on page 204 for more information.

Address: Address of detail information in the data space of the dataset.

ALET: ALET (Address List Entry Token) of information in the data space of the dataset.

Allocated GB: Allocated GB for the dataset.

Allocated MB: Allocated MB for the dataset.

Allocated Tracks: Allocated tracks for the dataset.

Allocated Volumes: The number of volumes on which space has been allocated for the dataset.

Dataset Name: The name of the dataset.

DSORG: The dataset organization.

Extents: The number of extents.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Timestamp: The time that the row was created.

Unused GB: Unused GB for the dataset.

Unused MB: Unused MB for the dataset.

Unused Tracks: Unused tracks for the dataset.

Used GB: Used GB for the dataset.

Used MB: Used MB for the dataset.

Used Tracks: Used tracks for the dataset.

Volser: The VOLSER on which the dataset is located.

Dataset Attributes High Level Qualifier Summary

This topic details the columns in the Dataset Attributes High Level Qualifier Summary table.

Allocated Bytes: Total allocated bytes for all datasets with this HLQ.

Allocated GB: Total allocated GB for all datasets with this HLQ.

Allocated KB: Total allocated KB for all datasets with this HLQ.

Allocated MB: Total allocated MB for all datasets with this HLQ.

Allocated Tracks: Total allocated tracks for all datasets with this HLQ.

Catalog Name: The catalog name for Alias for this HLQ.

Cataloged Dataset Count: Cataloged dataset count under HLQ.

High-Level Qualifier: High-Level Qualifier.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Number of Datasets: Total number of processed datasets under this HLQ.

Percent of Datasets: Percentage of total datasets represented by datasets under this HLQ.

Percent of Allocated Space: Percentage of total allocated space represented by datasets under this HLQ.

Timestamp: The time that the row was created.

Uncataloged Dataset Count: Uncataloged dataset count under HLQ.

Unused Bytes: Total unused bytes for all datasets with this HLQ.

Unused GB: Total unused GB for all datasets with this HLQ.

Unused KB: Total unused KB for all datasets with this HLQ.

Unused MB: Total unused MB for all datasets with this HLQ.

Unused Tracks: Total unused tracks for all datasets with this HLQ.

Used Bytes: Total used bytes for all datasets with this HLQ.

Used GB: Total used GB for all datasets with this HLQ.

Used KB: Total used KB for all datasets with this HLQ.

Used MB: Total used MB for all datasets with this HLQ.

Used Tracks: Total used tracks for all datasets with this HLQ.

Wasted Bytes: Total wasted bytes for all datasets with this HLQ.

Wasted GB: Total wasted GB for all datasets with this HLQ.

Wasted KB: Total wasted KB for all datasets with this HLQ.

Wasted MB: Total wasted MB for all datasets with this HLQ.

Dataset Attributes Inefficient Blocksize Summary

The Dataset Attributes Inefficient Blocksize Summary attributes provide monitoring data for Top N datasets in terms of used space, where the BLKSIZE is less than or equal to 128 bytes (in other words, the 'inefficient' threshold that the product defines). You define the value of n in the Dataset Attributes Data Collection dialog box. See "Dataset Attributes Database" on page 204 for more information.

Address: Address of detail information in the data space of the dataset.

ALET: ALET (Address List Entry Token) of information in the data space of the dataset.

Allocated GB: Allocated GB for the dataset.

Allocated MB: Allocated MB for the dataset.

Allocated Tracks: Allocated tracks for the dataset.

Allocated Volumes: The number of volumes on which space has been allocated for the dataset.

BLKSIZE: The block size of the dataset.

Creation Date: The date that the dataset was created.

Dataset Name: The name of the dataset.

DSORG: The dataset organization.

Extents: The number of extents.

LRECL: The logical record length.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

RECFM: The record format of the dataset.

Reference Date: The date that the dataset was last referenced.

Timestamp: The time that the row was created.

Unused GB: Unused GB for the dataset.

Unused MB: Unused MB for the dataset.

Unused Tracks: Unused tracks for the dataset.

Used GB: Used GB for the dataset.

Used MB: Used MB for the dataset.

Used Tracks: Used tracks for the dataset.

Volser: The VOLSER on which the dataset is located.

Dataset Attributes Mature Dataset Summary

The Dataset Attributes Mature Dataset Summary attributes provide monitoring data for the Top N datasets in terms of allocated space, where the number of days since last reference is greater than 366, which is the threshold that the product uses to categorize a dataset as mature. You define the value of n in the Dataset Attributes Data Collection dialog box. See "Dataset Attributes Database" on page 204 for more information.

Address: Address of detail information in the data space of the dataset.

ALET: ALET (Address List Entry Token) of information in the data space of the dataset.

Allocated GB: Allocated GB for the dataset.

Allocated MB: Allocated MB for the dataset.

Allocated Tracks: Allocated tracks for the dataset.

Allocated Volumes: The number of volumes on which space has been allocated for the dataset.

Creation Date: The date the dataset was created.

Dataset Name: The name of the dataset.

Days Since Creation: The number of days since creation.

Days Since Reference: The number of days since last reference.

Extents: The number of extents.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Reference Date: The date the dataset was last referenced.

Timestamp: The time that the row was created.

Unused GB: Unused GB for the dataset.

Unused MB: Unused MB for the dataset.

Unused Tracks: Unused tracks for the dataset.

Used GB: Used GB for the dataset.

Used MB: Used MB for the dataset.

Used Tracks: Used tracks for the dataset.

Volser: The VOLSER on which the dataset is located.

Dataset Attributes Never Referenced Dataset Summary

The Dataset Attributes Never Referenced Dataset Summary attributes provide monitoring data for the Top N datasets, in terms of allocated space, where the last reference date has not been set. (In other words, the dataset has never been opened, or it has been opened using an access method that does not maintain the last reference date field.) You define the value of *n* in the Dataset Attributes Data Collection dialog box. See "Dataset Attributes Database" on page 204 for more information.

Address: Address of detail information in the data space of the dataset.

ALET: ALET (Address List Entry Token) of information in the data space of the dataset.

Allocated GB: Allocated GB for the dataset.

Allocated MB: Allocated MB for the dataset.

Allocated Tracks: Allocated tracks for the dataset.

Allocated Volumes: The number of volumes on which space has been allocated for the dataset.

Creation Date: The date the dataset was created.

Dataset Name: The name of the dataset.

Days Since Creation: The number of days since creation.

Extents: The number of extents.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Timestamp: The time that the row was created.

Unused GB: Unused GB for the dataset.

Unused MB: Unused MB for the dataset.

Unused Tracks: Unused tracks for the dataset.

Used GB: Used GB for the dataset.

Used MB: Used MB for the dataset.

Used Tracks: Used tracks for the dataset.

Volser: The VOLSER on which the dataset is located.

Dataset Attributes Newborn Dataset Summary

The Dataset Attributes Newborn Dataset Summary attributes provide monitoring data for the Top N non-temporary datasets, where the creation date is within 1 day of the collection cycle start time. You define the value of *n* in the Dataset Attributes Data Collection dialog box. See "Dataset Attributes Database" on page 204 for more information.

Address: Address of detail information in the data space of the dataset.

ALET: ALET (Address List Entry Token) of information in the data space of the dataset.

Allocated GB: Allocated GB for the dataset.

Allocated MB: Allocated MB for the dataset.

Allocated Tracks: Allocated tracks for the dataset.

Allocated Volumes: The number of volumes on which space has been allocated for the dataset.

Dataset Name: Name of the dataset.

Extents: Number of extents.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Timestamp: The time that the row was created.

Unused GB: Unused GB for the dataset.

Unused MB: Unused MB for the dataset.

Unused Tracks: Unused tracks for the dataset.

Used GB: Used GB for the dataset.

Used MB: Used MB for the dataset.

Used Tracks: Used tracks for the dataset.

Volser: VOLSER on which the dataset is located.

Dataset Attributes Reference Date Summary

The Dataset Attributes Reference Date Summary attributes provide the count of the number of datasets that exist in each product-defined unique Reference Date (in other words, number of days since last reference) range. See "Dataset Attributes Database" on page 204 for more information.

Allocated GB: Allocated GB for all datasets in the range.

Allocated MB: Allocated MB for all datasets in the range.

Allocated Tracks: Allocated tracks for all datasets in the range.

High Days Since Reference: The highest value in the reference date range.

Low Days Since Reference: The lowest value in the reference date range.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Number of Datasets: The number of datasets in the range.

Percent of Allocated Space: The percentage of total allocated space that datasets in this reference date range represent.

Percent of Total Datasets: The percentage of total processed datasets that datasets in this reference date range represent.

Timestamp: The time that the row was created.

Unused GB: Total unused GB for all datasets in the range.

Unused MB: Total unused MB for all datasets in the range.

Unused Tracks: Total unused tracks for all datasets in the range.

Used GB: Total used GB for all datasets in the range.

Used MB: Total used MB for all datasets in the range.

Used Tracks: Total used tracks for all datasets in the range.

Wasted GB: Total wasted GB for all datasets in this range.

Wasted MB: Total wasted MB for all datasets in this range.

Dataset Attributes SMS Class Detail

The Dataset Attributes SMS Class Detail attributes provide data on the Top N datasets within an SMS Data/Storage/Management Class or Storage Group in terms of allocated space. You define the value of *n* in the Dataset Attributes Data Collection dialog box. See "Dataset Attributes Database" on page 204 for more information.

Address: Address of detail information in the data space of the dataset.

ALET: ALET (Address List Entry Token) of information in the data space of the dataset.

Allocated GB: Allocated GB for the dataset.

Allocated MB: Allocated MB for the dataset.

Allocated Tracks: Allocated tracks for the dataset.

Allocated Volumes: The number of volumes on which space has been allocated for the dataset.

Dataset Name: The name of a dataset in this class.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

SMS Class: Name of the SMS class.

Timestamp: The time that the row was created.

Type: The type of construct. The values are as follows: D (Data Class); S (Storage Class); M (management Class); G (Storage Group).

Unused GB: Unused GB for the dataset.

Unused MB: Unused MB for the dataset.

Unused Tracks: Unused tracks for the dataset.

Used GB: Used GB for the dataset.

Used MB: Used MB for the dataset.

Used Tracks: Used tracks for the dataset.

Volser: The VOLSER on which this dataset is located.

Dataset Attributes SMS Class Summary

The Dataset Attributes SMS Class Summary attributes provide a summary of the number of SMS-managed datasets encountered in each of the DFSMS Data, Storage and Management classes and Storage Groups defined on the system. See "Dataset Attributes Database" on page 204 for more information.

Allocated GB: Total allocated GB represented by all of the datasets in the class.

Allocated MB: Total allocated MB represented by all of the datasets in the class.

Allocated Tracks: Total allocated tracks represented by all of the datasets in the class.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Number of Datasets: The number of datasets encountered in the indicated class.

Percent of Total Allocated: The percentage of total allocated space represented by datasets in the class.

Percent of Total Datasets: The percentage of total processed datasets represented by datasets in the class.

SMS Class: The name of the SMS class. A blank class indicates that there is no construct of this type defined for the dataset. A class of **Unknown** indicates that the class name is not defined in the active DFSMS configuration.

Timestamp: The time that the row was created.

Type: The type of construct. The values are as follows: D (Data class); S (Storage class); M (Management class); G (Storage Group)

Unused GB: Total unused GB for all datasets in this class.

Unused MB: Total unused MB for all datasets in this class.

Unused Tracks: Total unused tracks for all datasets in the class.

Used GB: Total used GB for all datasets in this class.

Used MB: Total used MB for all datasets in this class.

Used Tracks: Total used tracks for all datasets in this class.

Wasted GB: Total wasted GB for all datasets in this SMS class.

Wasted MB: Total wasted MB for all datasets in this SMS class..

Dataset Attributes Space Allocated Summary

The Dataset Attributes Space Allocated Summary attributes provide data on the Top N datasets in terms of total allocated space. You define the value of n in the Dataset Attributes Data Collection dialog box. See "Dataset Attributes Database" on page 204 for more information.

Address: Address of detail information in the data space of the dataset.

ALET: ALET (Address List Entry Token) of information in the data space of the dataset.

Allocated GB: Allocated GB for the dataset.

Allocated MB: Allocated MB for the dataset.

Allocated Tracks: Allocated tracks for the dataset.

Allocated Volumes: The number of volumes on which space has been allocated for the dataset.

Dataset Name: The name of the dataset.

DSORG: The dataset organization.

Extents: The number of extents.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Timestamp: The time that the row was created.

Unused GB: Unused GB for the dataset.

Unused MB: Unused MB for the dataset.

Unused Tracks: Unused tracks for the dataset.

Used GB: Used GB for the dataset.

Used MB: Used MB for the dataset.

Used Tracks: Used tracks for the dataset.

Volser: The VOLSER on which the dataset is located.

Dataset Attributes Space Unused Summary

The Dataset Attributes Space Unused Summary attributes provide data on the Top N datasets in terms of total unused space. You define the value of n in the Dataset Attributes Data Collection dialog box. See "Dataset Attributes Database" on page 204 for more information.

Address: Address of detail information in the data space of the dataset.

ALET: ALET (Address List Entry Token) of information in the data space of the dataset.

Allocated GB: Allocated GB for the dataset.

Allocated MB: Allocated MB for the dataset.

Allocated Tracks: Allocated tracks for the dataset.

Allocated Volumes: The number of volumes on which space has been allocated for the dataset.

Dataset Name: The name of the dataset.

DSORG: The dataset organization.

Extents: The number of extents.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Timestamp: The time that the row was created.

Unused GB: Unused GB for the dataset.

Unused MB: Unused MB for the dataset.

Unused Tracks: Unused tracks for the dataset.

Used GB: Used GB for the dataset.

Used MB: Used MB for the dataset.

Used Tracks: Used tracks for the dataset.

Volser: The VOLSER on which the dataset is located.

Dataset Attributes System Summary

The Dataset Attributes System Summary attributes provide summary data representing all volumes and datasets encountered in a collection cycle. One row exists in this table. The data in the row is refreshed at each collection cycle. See "Dataset Attributes Database" on page 204 for more information.

Allocated GB: The total GB allocated to all datasets on all volumes.

Allocated MB: The total MB allocated to all datasets on all volumes.

Allocated Tracks: The total tracks allocated to all datasets on all volumes.

Available GB: The total GB available on all volumes encountered.

Available MB: The total MB available on all volumes encountered.

Available Tracks: The total tracks available (in other words, device capacity) on all volumes encountered.

Collection End Time: The date and time of day when the last collection cycle ended.

Collection Start Time: The date and time of day when the last collection cycle began.

Datasets Processed: The total number of datasets on all volumes that were processed during the collection cycle.

Largest Dataset Allocated GB: The size of the dataset with the largest amount of allocated space in GB.

Largest Dataset Allocated MB: The size of the dataset with the largest amount of allocated space in MB.

Largest Dataset Allocated Tracks: The size of the dataset with the largest amount of allocated space in Tracks.

Largest Dataset Allocated Volumes: The number of volumes on which space has been allocated for the dataset with the largest amount of allocated space.

Largest Dataset Name: The name of the dataset with the largest amount of allocated space.

Largest Dataset Volume: Volume on which the dataset with the largest amount of allocated space is located.

Largest Unused Space Dataset Allocated Volumes: The number of volumes on which space has been allocated for the dataset with the largest amount of allocated and unused space.

Largest Unused Space Dataset GB: The size of unused space in GB within the dataset with the largest amount of allocated and unused space.

Largest Unused Space Dataset MB: The size of unused space in MB within the dataset with the largest amount of allocated and unused space.

Largest Unused Space Dataset Name: The name of the dataset with the largest amount of allocated and unused space.

Largest Unused Space Dataset Tracks: The size of unused space in Tracks within the dataset with the largest amount of allocated and unused space.

Largest Unused Space Dataset Volume: Volume on which the dataset with the largest amount of allocated and unused space is located.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Oldest Reference Dataset Age: The number of days that oldest dataset has gone unreferenced. The reference date must be non-zero. In other words, datasets that have never been referenced are not counted.

Oldest Reference Dataset Allocated Volumes: The number of volumes on which space has been allocated for the oldest dataset (where reference date is non-zero, that is, not counting never referenced datasets).

Oldest Reference Dataset Name: The name of the dataset with the largest number of days since last reference. The reference date must be non-zero. In other words, datasets that have never been referenced are not counted.

Oldest Reference Dataset Volser: The volume on which the dataset with the largest number of days since last reference is located.

Percent Allocated: The percentage of all available space that is allocated.

Percent Free: The percentage of all allocated space that is allocated but unused.

Percent Used: The percentage of all allocated space that is in use.

Status Message: Status of the Dataset Attributes Database data collection process.

Timestamp: The time that the table row was created.

Uncataloged Datasets: The total number of uncataloged datasets encountered on all volumes. This attribute is valid only if 'Collect Catalog Data' option is on in collection control table.

> **Note:** The SYS1.VVDS and SYS1.VTOCIX datasets are known to always be uncataloged. They are not included in the Uncataloged Dataset total.

Unused GB: The total GB allocated but unused by all datasets on all volumes.

Unused MB: The total MB allocated but unused by all datasets on all volumes.

Unused Tracks: The total tracks allocated but unused by all datasets on all volumes.

Used GB: The total GB used by all datasets on all volumes.

Used MB: The total MB used by all datasets on all volumes.

Used Tracks: The total tracks allocated and used by all datasets on all volumes.

Volumes Online: Total number of DASD volumes that were online at the time that collection ran on the system.

Volumes Processed: The total number of volumes processed.

Dataset Metrics

This Web-based user interface lets you examine large volumes of data about datasets.

The **Run Query** button queries the default server for dataset information, filtered according to the **Filter By** criteria specified in the filtering section, and sorted by the current sorting column in the specified sort order.

The **Log Off** button lets you exit Dataset Metrics.

View By

You can explore dataset information through different views. The user can select which view to display by selecting one of the "VIEW BY" options.

Column Groups

The Dataset view has a large number of columns. To give users manageable access to more columns, the Dataset view provides a column group selector.

Filter By

These settings let you specify search criteria for the datasets you want to analyze.

Column Groups

The Dataset view has a large number of columns. To give users manageable access to more columns, the Dataset view provides a column group selector.

Dataset

The Dataset view shows the information for individual datasets. The Dataset view has eight column groups.

HLQ

This view shows data summarized by High Level Qualifier, or HLQ, which is the dataset name up to and not including the first period (.), collected by Dataset Metrics data collection. The data in this view should correspond with the data in a new TEP workspace called Dataset HLQ Summary.

The other views (Management Class, Data Class, Storage Class, and Storage Group) show information about the corresponding SMF classes.

Results Table

When you click a column header, the table is sorted in ascending order by that column. Clicking a sorted column reverses the sort order.

Dataset Attributes Uncataloged Detail

The Dataset Attributes Uncataloged Detail attributes provide data on the Top N uncataloged datasets on a volume in terms of allocated space.

> **Note:** The SYS1.VVDS and SYS1.VTOCIX datasets are known to always be uncataloged. They are not included in the Uncataloged Detail for a volume.

When the Collect Catalog Data option is enabled, the catalog data is retrieved for each dataset encountered in the VTOC of the volume. See “Dataset Attributes Database” on page 204 for more information.

Address: Address of detail information in the data space of the dataset.

ALET: ALET (Address List Entry Token) of information in the data space of the dataset.

Allocated GB: Allocated GB for the dataset.

Allocated MB: Allocated MB for the dataset.

Allocated Tracks: Allocated tracks for the dataset.

Dataset Name: Name of the uncataloged dataset.

Extents: Number of extents.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Timestamp: The time that the row was created.

Unused GB: Unused GB for the dataset.

Unused MB: Unused MB for the dataset.

Unused Tracks: Unused tracks for the dataset.

Used GB: Used GB for the dataset.

Used MB: Used MB for the dataset.

Used Tracks: Used tracks for the dataset.

Volser: The VOLSER on which this dataset is located.

Wasted GB: Wasted GB for the dataset.

Wasted MB: Wasted MB for the dataset.

Dataset Attributes Uncataloged Summary

The Dataset Attributes Uncataloged Summary attributes provide the number of uncataloged datasets by volume. If no catalog entry is found for a dataset, the count of uncataloged datasets on that volume is incremented. The space totals are also incremented.

> **Note:** The SYS1.VVDS and SYS1.VTOCIX datasets are known to always be uncataloged. They are not included in the Uncataloged Detail for a volume.

When the Collect Catalog Data option is enabled, the catalog data is retrieved for each dataset encountered in the VTOC of the volume. See "Dataset Attributes Database" on page 204 for more information.

Allocated GB: Total allocated GB for all uncataloged datasets on the volume.

Allocated MB: Total allocated MB for all uncataloged datasets on the volume.

Allocated Tracks: Total allocated tracks for all uncataloged datasets on the volume.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Timestamp: The time that the row was created.

Total Datasets: Total number of uncataloged datasets on this volume.

Unused GB: Total unused GB for all uncataloged datasets on the volume.

Unused MB: Total unused MB for all uncataloged datasets on the volume.

Unused Tracks: Total unused tracks for all uncataloged datasets on the volume.

Used GB: Total used GB for all uncataloged datasets on the volume.

Used MB: Total used MB for all uncataloged datasets on the volume.

Used Tracks: Total used tracks for all uncataloged datasets on the volume.

Volser: The name of the volume on which at least one uncataloged dataset was found.

Wasted GB: Total wasted GB for uncataloged datasets on the volume.

Wasted MB: Total wasted MB for uncataloged datasets on the volume.

DSN Attr Group Detail

The DSN Attr Group Detail attributes provide detailed information for datasets in a specified group.

> **Note:** Although a Dataset Attributes Group can summarize more than 5000 datasets, a query against the DSN Attr Group Detail table for a specified group returns a maximum of 5000 datasets. Additionally, be aware that in a large enterprise the Dataset Attributes Database can be tracking millions of datasets.

When you create a group, define the properties and attributes such that they limit the number of datasets that make up the group in order to prevent it from requiring excessive processing, network, and space resources.

Allocated GB: The allocated GB for the dataset.

Allocated MB: The allocated MB for the dataset.

Allocated Tracks: The allocated tracks for the dataset.

Allocated Volumes: The number of volumes on which space is allocated for the dataset.

Associated Entry Name: The entry name associated with the dataset (for example, cluster name for a data component).

Average LRECL: The VSAM dataset average logical record length.

Bytes Saved By CA-Reclaim The number of bytes saved by using the CA-Reclaim feature.

CA Reclaim: CA Reclaim.

CA Reclaimed: Number of Reclaimed CAs.

CA Reused: Number of Reused CAs.

CA Splits: The number of CA splits.

Catalog: The dataset is a catalog.

Cataloged: The dataset is cataloged.

Catalog Entry Type: The catalog entry type.

CIs per CA: The number of CIs per CA.

CI Size: The control interval size.

CI Splits: The number of CI splits.

CMS: The dataset is using Cylinder Managed Space on an Extended Address Volume.

Collection Timestamp: The date and time of data collection for this row (info-only column).

Compressed Size: The VSAM dataset compressed user data size in bytes.

Compressed Size MB: The VSAM dataset compressed user data size in MB.

Compressed Size GB: The VSAM dataset compressed user data size in GB.

Component Free Space: The bytes of free space in a VSAM dataset component.

DAEAddress: The address of detail information in the dataset dataspace.

DAEALET: ALET of detail information in the dataset dataspace.

Data Component: The dataset is a VSAM data component.

Dataset Name: The name of the dataset.

Creation Date: The date when the dataset was created.

Days Since Creation: The days since the dataset was created.

Days Since Until Expiration: The days since (positive number) or until (negative number) the specified expiration date.

Days Since Last Backup: The number of days since the last backup.

Days Since Reference: The days since the last reference of the dataset.

Encrypted: The dataset is encrypted.

Encryption Algorithm: The dataset encryption algorithm.

Encryption Key Length: The dataset encryption key length.

Encryption Mode: The dataset encryption mode.

Erase: The VSAM attribute is erase.

EXCPS: The number of VSAM EXCPS.

Extents: The number of extents allocated by the dataset.

Extended Format: The extended format dataset.

Extended Address Space Eligible: The dataset is eligible to use Extended Address Space of an Extended Address Volume.

Free Bytes per CI: The free bytes per CI.

Free CIs per CA: The free CIs per CA.

GDG Attributes: The GDG attributes.

GDG Level: The GDG level.

Generation Dataset: The dataset is a generation dataset.

Group Name: The name of the user-defined dataset group.

High Allocated RBA: The high allocated RBA for the dataset on this volume.

High Used RBA: The high used RBA for the dataset on this volume.

Index Component: The dataset is a VSAM index component.

Key Label: The Key Label for encrypted dataset.

Key Length: For VSAM datasets, this is the KSDS key length; and for non-VSAM datasets, this is the DASD key length.

Managed System: The internal node name used to identify the system from which this row of data was obtained.

Maximum Buffer Size: The VSAM dataset maximum buffer size.

Maximum Datasets in GDG: The maximum number of generation datasets that can be associated with the generation data group.

Maximum Record Size: The VSAM dataset maximum record size.

Owner: The VSAM dataset owner.

Page Space: The VSAM attribute is page space.

Percent Free CIs per CA: The percentage of CIs free in each CA.

Percent Free Bytes per CI: The percentage of bytes free in each CI.

Percent Free: The percentage of the dataset's allocated space free.

Percent Used: The percentage of the dataset's allocated space in use.

Primary GB: The GB in the first extent.

Primary MB: The MB in the first extent.

Primary Tracks: The number of tracks in the first extent.

RACF: The RACF indicator is set for the dataset in the VTOC or, for VSAM, the catalog.

Records Deleted: The number of deleted records in a VSAM dataset.

Records Inserted: The number of inserted records in a VSAM dataset.

Records Retrieved: The number of retrieved records in a VSAM dataset.

Records Total: The number of logical records in a VSAM dataset.

Records Updated: The number of updated records in a VSAM dataset.

Recovery: The VSAM attribute is recovery.

Reference Date: The date when the dataset was last referenced.

Reuse: The VSAM attribute is reuse.

Secondary Allocation Record Units: The secondary space record allocation units.

Secondary Allocation Record Value: The secondary space allocation value when specified in records.

Secondary Allocation Units: The secondary space record allocation units.

Secondary Allocation Quantity: The secondary space record allocation quantity in specified record allocation units.

Share Option: The VSAM dataset share options (cross-region, cross-system).

SMS Data Class: The DFSMS data class.

SMS Storage Class: The DFSMS storage class.

SMS Storage Group: The DFSMS storage group.

SMS Management Class: The DFSMS management class.

Spanned: The VSAM attribute is spanned.

Speed: The VSAM attribute is speed.

Stripe Count: The stripe count.

Temporary: The temporary dataset.

Unique: The VSAM attribute is unique.

Unmovable: The dataset is unmovable.

Unused Tracks: The unused tracks for the dataset.

Unused MB: The unused MB for the dataset.

Unused GB: The unused GB for the dataset.

Used GB: The used GB for the dataset.

Used MB: The used MB for the dataset.

Used Tracks: The used tracks for the dataset.

Volser: The volser on which the dataset is located.

Volume Count: The number of volumes on which the dataset resides.

Volume Sequence: The sequence number of this volume relative to all volumes on which the dataset is located.

VSAM Dataset Organization: The VSAM dataset organization.

Wasted GB: Wasted GB for the dataset.

Wasted MB: Wasted MB for the dataset.

ZFS: The dataset is ZFS.

DSN Attr Group Summary

The DSN Attr Group Summary attributes provide summary data representing the datasets encountered in a collection cycle that make up each group. The data in each row is refreshed at each collection cycle. See "Dataset Attributes Database" on page 204 for more information.

> **Note:** Be aware that in a large enterprise the Dataset Attributes Database can track millions of datasets. When you create a group, define the properties and attributes such that they limit the number of datasets that make up the group to prevent the group from requiring excessive processing, network, and space resources.

Collection Timestamp: The time of data collection.

Dataset Count: The number of datasets in the group.

Description: The description of the user defined dataset group.

Disable Collection: Indicates if the collection is disabled for the user defined dataset group.

Group Name: The name of the group.

Managed System: The internal node name used to identify the system from which this row of data was obtained.

Maximum Allocated GB: The largest amount of allocated GB of the datasets in the group.

Maximum Allocated MB: The largest amount of allocated MB of the datasets in the group.

Maximum Allocated Tracks: The largest amount of allocated tracks of the datasets in the group.

Maximum Days Since Referenced: The largest number of days since the last referenced of the datasets in the group.

Maximum Unused GB: The largest amount of unused GB of the datasets in the group.

Maximum Unused MB: The largest amount of unused MB of the datasets in the group.

Maximum Unused Tracks: The largest amount of unused tracks of the datasets in the group.

Maximum Used GB: The largest amount of used GB of the datasets in the group.

Maximum Used MB: The largest amount of used MB of the datasets in the group.

Maximum Used Tracks: The largest amount of used tracks of the datasets in the group.

Maximum Wasted GB: Largest amount of wasted GB of all datasets in the group.

Maximum Wasted MB: Largest amount of wasted MB of all datasets in the group.

Minimum Allocated GB: The smallest amount of allocated GB of the datasets in the group.

Minimum Allocated MB: The smallest amount of allocated MB of the datasets in the group.

Minimum Allocated Tracks: The smallest amount of tracks of the datasets in the group.

Minimum Unused GB: The smallest amount of unused GB of the datasets in the group.

Minimum Unused MB: The smallest amount of unused MB of the datasets in the group.

Minimum Unused Tracks: The smallest number of unused tracks of the datasets in the group.

Minimum Used GB: The smallest amount of used GB of the datasets in the group.

Minimum Used MB: The smallest amount of used MB of the datasets in the group.

Minimum Used Tracks: The smallest amount of used tracks of the datasets in the group.

Minimum Wasted GB: The smallest amount of wasted GB of all datasets in the group.

Minimum Wasted MB: The smallest amount of wasted MB of all datasets in the group.

Oldest Creation Date: This is the create date of the oldest dataset of the datasets in this group.

Persist: Indicates if the group definition will be saved across monitoring server restarts.

Total Allocated GB: The sum of the allocated space in GB of the datasets in the group.

Total Allocated MB: The sum of the allocated space in MB of the datasets in the group.

Total Allocated Tracks: The sum of the allocated space in tracks of the datasets in the group.

Total Unused GB: The total amount of unused GB of the datasets in the group.

Total Unused MB: The total amount of unused MB of the datasets in the group.

Total Unused Tracks: The total amount of unused tracks of the datasets in the group.

Total Used GB: The total amount of used GB of the datasets in the group.

Total Used MB: The total amount of used MB of the datasets in the group.

Total Used Tracks: The total amount of used tracks of the datasets in the group.

Total Wasted GB: Total Wasted GB for all datasets in the group.

Total Wasted MB: Total Wasted MB for all datasets in the group.

Unused GB Percent: The unused GB as a percentage of total allocated GB of the datasets in the group.

Unused MB Percent: The unused MB as a percentage of total allocated MB of the datasets in the group.

Unused Tracks Percent: The unused tracks as a percentage of total allocated tracks of the datasets in the group.

Used GB Percent: The used GB as a percentage of total allocated GB of the datasets in the group.

Used MB Percent: The used MB as a percentage of total allocated MB of the datasets in the group.

Used Tracks Percent: The used tracks as a percentage of total allocated tracks of the datasets in the group.

Dataset Group Defs

The Dataset Group Defs attributes provide information to other attribute groups. Users do not use these attributes. The attributes are documented for the purpose of background reference only.

Dataset Mask Data: Group mask bulk data. This contains a mask name status and space collection flag. These fields are repeated for each Group Mask being maintained. Mask name: max 44 characters; Status : 0=inactive, 1=active, 2=stopped

Dataset Group Name: This is a user defined name used to identify a group of dataset name masks that comprise the DSN group.

Dataset Mask Flag: Dataset mask flag.

Dataset Mask Sequence: Dataset mask sequence field.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Message: Dataset mask message.

Dataset Group Details

The Dataset Group Details attributes provide detailed information on your dataset. The attributes in this group include details regarding control intervals (CIs) and control areas (CAs).

% Free Bytes In CI: The percentage of bytes free across all CIs. (VSAM only.) You can enable aggregation for this attribute in Tivoli® Data Warehouse.

% Free CIS In CA: The percentage of free CIs across all CAs in the dataset. (VSAM only.) You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Alternate Index Attributes: The alternate index attributes. Values are UPGRADE, and NO UPGRADE.

Associated Entry Name: The associated entry names.

Average LRECL: The VSAM dataset average logical record size.

Average MSR: The average MSR (millisecond response time) for this dataset.

Bytes Free In Compression: The number of bytes of free space in compression. (VSAM only.)

CA Splits: The number of CA splits. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

A CA split is the movement of half of the records in an existing CA to a new CA because a record add or update cannot be accommodated in the existing CA. This condition results in two approximately half-full CAs instead of one full and one empty CA. A CA split is inefficient because it requires a large number of I/O operations. (VSAM only.)

Catalog Name: The name of the catalog that the dataset belongs.

CI Size: The control interval (CI) size of the dataset. A CI is a unit of data that is transferred between auxiliary storage and virtual storage when an I/O request is made. It contains records, free space, and control information.

CI Splits: The number of control interval (CI) splits. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

A CI split is the movement of some records from an existing CI to another free CI in the same control area because a record add or update cannot be accommodated in the existing one. This condition results in two half-empty CIs instead of one full and one empty CI. (VSAM only.)

Compressed User Data Size: The size of compressed user data.

Compression Indicator: The compression of the dataset. Values are EXTENDED FORMAT, COMPRESSIBLE, or EXTENDED FORMAT.

Creation Date: The date when the dataset was created.

Dataset % Free: The percentage of space that is free in this dataset. This value is calculated using: `(Dataset High Allocated RBA - Dataset High Used RBA) \ Dataset High Allocated RBA`. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Dataset % Used: The percentage of space that is used in this dataset. This is calculated using: `Dataset High Used RBA \ Dataset High Allocated RBA` . You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Dataset Flags: The dataset entry type. Values are PRIMARY and NOT PRIMARY.

Dataset Group Name: A user-defined name that identifies a group of dataset name masks that make up the Dataset Group.

Dataset High Allocated RBA: The dataset high-allocated RBA.

Dataset High Used RBA: The dataset high-used RBA.

Dataset Mask: The user-defined dataset mask that matches the dataset.

Dataset Name: The name of the dataset for which the information was collected.

Dataset Type: The dataset entry type. Values are: NON-VSAM DATASET, GENERATION DATA GROUP, CLUSTER, DATA COMPONENT, ALTERNATE INDEX, GENERATION DATASET, INDEX COMPONENT, PATH, ALIAS, USER CAT CONNECTOR ENTRY.

Days Since Creation: The number of days since the dataset was created.

Days Unreferenced: The number of days since the dataset was last referenced.

EXCPS: The number of EXCPS for the dataset. (VSAM only.) You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Expiration Date: The date when the dataset expires.

Extended Attributes: The extended attributes.

Extended Format Attributes (HEX): The extended format attributes expressed as hexadecimal values.

Extended Format Attributes: The extended format attributes.

First Extent Tracks Allocated: The primary space allocation.

Free CIS Per CA: The number of free control intervals (CIs) per control area (CA). (VSAM only.) You can enable aggregation for this attribute in Tivoli® Data Warehouse.

GDG Dataset Attributes: The Dataset Generation Data Group Attributes. Values are as follows:

Delete oldest GDS GATLIMIT exceeded.
Delete all GDSs when GATLIMIT exceeded.
Do not scratch dataset when rolled off.
Scratch dataset when rolled off if volume mounted.

GDG Last Altered Date: The date the dataset Generation Data Group (GDG) was last changed.

GDG Maximum Datasets: The maximum datasets in the Generation Data Group (GDG).

KSDS Key Length: The length of the KSDS key.

KSDS Key Position: The relative position of key within the dataset.

Last Backup Date: The date of the most recent dataset backup.

Log Parameter Set: The value of the log parameter set. Values are NEVER SET, LOG NONE, LOG UNDO, LOG ALL.

Managed System: The internal name used to identify the system from which this row of data was obtained.

Maximum Buffer Size: The maximum buffer size.

Maximum I/Os: The maximum I/Os for a volume for this dataset.

Maximum MSR: The maximum MSR (millisecond response time) for this dataset.

Maximum Record Size: The maximum record size for this dataset.

NON-VSAM Attributes: The non-VSAM attributes information for this dataset. Values are Active GDS, Deferred GDS, Rolled off GDS, PDSE, POSIX DS, NON VSAM.

Number Bytes per CI: The number of free bytes across all control intervals (CIs). (VSAM only.)

Number CIS per CA: Contains the number of control intervals (CIs) that are contained in a control area (CA) for this dataset. A control area is a group of control intervals. (VSAM only.)

Owner: Identifies the dataset owner.

Primary Volume: The volume serial number for primary space allocation.

Records Deleted: The number of deleted records for the dataset. (VSAM only.) You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Records Inserted: The number of inserted records for the dataset. (VSAM only.) You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Records Retrieved: The number of retrieved records for the dataset. (VSAM only.) You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Records Updated: The number of updated records for the dataset. (VSAM only.) You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Recovery Time: The recovery time, TOD value, local.

Reference Date: The date of the last reference for the dataset.

Request Type: Request type.

Secondary Space Allocation: The secondary space allocation.

SMFID: The SMFID of the system where the data was collected.

SMS Data Class: The SMS data class for the dataset.

SMS Management Class: The SMS management class for the dataset.

SMS Storage Class: The SMS storage class for the dataset.

Space Allocation Units: The space units used for allocation. Values are AVG BLOCK, AVG BLOCK ROUNDED, CYLINDER, TRACK, and n/a.

Striping Counts: The striping count for a striped dataset.

Timestamp: The timestamp of when the data was collected. Displayed in UTC (Coordinated Universal Time).

Total I/Os: The total sum of I/Os across all volumes for this dataset.

Total Records: The number of logical records for the dataset. (VSAM only.) You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Total Track Allocation: The total sum of tracks allocated across all volumes for this dataset.

Total Track Free: The total sum of tracks free across all volumes for this dataset.

Total Track Used: The total sum of tracks used across all volumes for this dataset.

Volume Count: The number of volumes for this dataset.

VSAM Dataset Attributes (HEX): The VSAM dataset attributes. The attributes are Speed, Unique, Reusable, Erase, Inhibit update, Temporary export, and Track overflow expressed as hexadecimal values.

VSAM Dataset Attributes: The VSAM dataset attributes. The attributes are Speed, Unique, Reusable, Erase, Inhibit update, Temporary export, and Track overflow.

VSAM Dataspace Attributes (HEX): Defines the attributes of the VSAM data space for the dataset, in hexadecimal. The values are Swap, No Swap, and Page Space (indicates the dataset is a pagespace).

VSAM Dataspace Attributes: Defines the attributes of the VSAM data space for the dataset. The values are Swap, No Swap, and Page Space (indicates the dataset is a pagespace).

VSAM Reuse Attributes (HEX): Defines the VSAM reuse attributes for the dataset in hexadecimal. Values (in hexadecimal) can be: Dataset has RACF® discrete profile, Index component dataset, Reusable dataset, Erase specified (cluster only), Not defined, swap space (cluster only), page space (cluster only).

VSAM Reuse Attributes: Defines the VSAM reuse attributes for the dataset. Values can be: Dataset has RACF® discrete profile, Index component dataset, Reusable dataset, Erase specified (cluster only), Not defined, swap space (cluster only), page space (cluster only).

VSAM Share Attributes: The share attributes of a VSAM dataset. The share attribute is used when a data or index component of a cluster, alternate index, or the data component of a catalog can be shared among users. Cross Region specifies the amount of sharing allowed among regions within the same system or within multiple

systems using global resource serialization (GRS). Cross System specifies the amount of sharing allowed among systems.

VSAM Type Attributes (HEX): The type attributes of a VSAM dataset in hexadecimal. Values (HEX) can be: KSDS or non-KSDS, Write check, Imbed, Replicate, Key-range dataset, RRDS (Relative Record Dataset), Spanned records allowed, Non-unique or unique keys allowed, LDS, or VRRDS. A LDS is a VSAM dataset that contains data but no control information. A VRRDS is a Relative Record Dataset and is similar to a non-VSAM relative file.

VSAM Type Attributes: The type attributes of a VSAM dataset. Values can be: KSDS or non-KSDS, Write check, Imbed, Replicate, Key-range dataset, RRDS (Relative Record Dataset), Spanned records allowed, Non-unique or unique keys allowed, LDS, or VRRDS. An LDS is a VSAM dataset that contains data but no control information. A VRRDS is a Relative Record Dataset and is similar to a non-VSAM relative file.

Dataset Group Detail Attrs

The Dataset Group Details attributes provide a variety of detailed information on your dataset.

Allocation values: The values of the dataset allocations.

Allocations: The dataset allocations. Values are Owner, Volume, Space Type, Primary Space, Secondary Space, High Allocated RBA, High Used RBA, Percent Free, and Percent Used.

Attribute values: The values of the dataset attributes.

Attributes: The dataset attributes.

The dataset attributes values are listed in three columns for readability:		
Column 1	**Column 2**	**Column 3**
Key Length Relative Key Position Average LRECL Maximum LRECL Buffer Space EXCP Exit Share Options Association UPGRADE NOUPGRADE SPEED RECOVERY UNIQUE NON-UNIQUE ERASE NOERASE INDEXED NON-INDEXED IMBED	NOIMBED REUSE NOREUSE SPANNED NONSPANNED WRITECHK NOWRITECHK REPLICATE NOREPLICATE NUMBERED LINEAR NONUNIQKEY RACF® discrete profile Index Component Dataset Reusable Erase Specified SWAP Space Page Space	Invalid COMUDSIZ and UDATASIZ Block level compression NON-VSAM Dataset Generation data Group Cluster Data Component Alternate Index Generation Dataset Index Component Path Alias User Cat Connector Entry NON VSAM Active GDS PDSE Rolled off GDS Deferred GDS POSIX DS

Dataset Group Name: A user-defined name that identifies a group of dataset name masks that make up the Dataset Group.

Dataset Name: A name of the dataset.

Managed System: The internal name used to identify the system from which this row of data was obtained.

Request Type: Request type.

Statistic Values: The values of the dataset statistics.

Statistics: The dataset statistics. Values are Total Records, Deleted Records, Inserted Records, Updated Records, Retrieved Records, CI Size, CI Splits, CA Splits, Number CIs per CA, Free CIs per CA, Free Space per CI, Percent Free CIs, Percent Free Bytes, EXCPS, Extents, and System Timestamp.

Dataset Group Details Volume

The Dataset Group Details Volume attributes provide detailed information on the dataset of a volume.

Allocation Flag: The dataset volume flags. The following are valid values:

CANDIDATE VOLUME
CANDIDATE VOLUME\CATALOG ENTRY ONLY
CANDIDATE VOLUME\OVERFLOW VOLUME
CANDIDATE VOLUME\OVERFLOW VOLUME\CATALOG ENTRY ONLY
CATALOG ENTRY ONLY
OVERFLOW VOLUME
OVERFLOW VOLUME\CATALOG ENTRY ONLY
PRIMARY VOLUME
PRIMARY VOLUME\CANDIDATE VOLUME
PRIMARY VOLUME\CANDIDATE VOLUME\CATALOG ENTRY ONLY
PRIMARY VOLUME\CATALOG ENTRY ONLY
PRIMARY VOLUME\INDIRECT
PRIMARY VOLUME\OVERFLOW VOLUME
PRIMARY VOLUME\OVERFLOW VOLUME\CATALOG ENTRY ONLY
PRIMARY VOLUME\ SYMBOLIC
n/a

Associated Entry Name: The associated entry name. (Example: VSAM Cluster component.)

Average MSR: The average MSR (millisecond response time) for the dataset on this volume. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Dataset Group Name: A user-defined name that identifies a group of dataset name masks that comprise the Dataset Group.

Dataset Name: The name of the dataset for which the information was collected.

Extended Address Space Eligible: Dataset is eligible to use Extended Address Space of an Extended Address Volume.

Extents: The number of extents for this dataset.

GDG Generation Level: The Dataset Generation Group Level.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Percent Free: The percentage of tracks free for this dataset on this volume. This is calculated using this formula: `(Tracks Allocated - Tracks Free ) / Tracks Free`

Percent Used: The percentage of tracks used for this dataset on this volume. This is calculated using this formula: `(Tracks Allocated - Tracks used ) / Tracks Used`

Physical Block Size: The number of bytes per track.

Request Type: Request type.

Sequence Number: Identifies the file sequence number.

Sequence Order: The sequence order of the volume.

SMFID: The SMFID of the system where the data was collected.

Timestamp: The time stamp of when the data was collected.

Total I/Os: The total I/Os for this dataset on this volume. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Tracks Allocated: The number of tracks allocated by this dataset on this volume. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Tracks Free: The number of tracks free for this dataset on this volume. This is calculated using: Tracks Allocated - Tracks Used.

Tracks Used: The number of tracks used by this dataset on this volume. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Using Cylinder Managed Space: Dataset is using Cylinder Managed Space on an Extended Address Volume.

Volume: The Volume Serial Number for primary or candidate space allocation.

VSAM Associated Entry Name: The name of the VSAM Dataset Associated entry. For example, VSAM Cluster component.

VSAM High Allocation RBA: The VSAM high-allocated Relative Byte Address.

VSAM High CI RBA: The VSAM Relative Byte Address data control interval with high key.

VSAM High Used RBA: The VSAM high-used Relative Byte Address.

VSAM Percent Free: The percentage of space that is free in this VSAM dataset. This is calculated using: `(High Allocated RBA - High Used RBA) \ High Allocated RBA`.

VSAM Percent Used: The percentage of space that is used in this VSAM dataset. This is calculated using: `High Used RBA \ High Allocated RBA`.

Dataset Group Summary

The Dataset Group Summary attributes provide summary information on your dataset.

Dataset Group Name: A user-defined name that identifies a group of dataset name masks that make up the dataset group.

Dataset with Highest MSR: The dataset with the slowest MSR (millisecond response time) within the group.

Dataset with Largest Allocation: The dataset with largest space allocation within the group.

Dataset with Largest Free Space: The dataset with largest amount of unused space within the group.

Dataset with Largest Usage: The name of the dataset with largest allocation, maximum highest space allocation, and oldest referenced date within the group, in tracks.

Dataset with Oldest Referenced: The dataset with oldest referenced date (read, write, or opened) within the group.

Enable Detail History Collection: The detail history collection status for the dataset group.

Group Total Average MSR: The total average MSR (millisecond response time) for all the active datasets within the group, in milliseconds. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Group Total Free Space: The total free space for all the datasets within the group, in tracks. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Group Total I/Os: The total I/Os for all the active datasets within the group. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Group Total Space Allocation: The total allocated space for all the datasets within the group, in tracks. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Group Total Used Space: The total used space for all the datasets within the group, in tracks. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Highest MSR: The highest MSR (millisecond response time) for the identified dataset within the group, in milliseconds. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Largest Space Allocation: The space allocation for the identified dataset, in tracks. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Largest Space Free: The free space for the identified dataset, in tracks. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Largest Space Usage: The largest allocation for the identified dataset within the group.

Largest Usage Oldest Reference Date: The reference date for the identified dataset within the group.

Managed System: An internal name used to identify the system from which this row of data was obtained.

Oldest Date Referenced: The last referenced date for the identified dataset.

Percent Free: The percentage of the group total space allocation not in use. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Percent Used: The percentage of the group total space allocation in use. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

RLS Dataset: Used to enable navigation to the Dataset Group RLS Summary workspace.

RLS Group: Used to build a dataset group that contains strictly RLS dataset masks. All masks in the group are subject to RLS masking rules.

Total Datasets: The total number of datasets in this dataset group. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Dataset Performance

The Dataset Performance attributes provide performance information for a dataset by application or summarized for all applications accessing the dataset. They provide metrics at the I/O component level.

ASID: The address space ID of the job using this dataset.

Cache Hit Percent: The sum of cache read hits plus cache write hits divided by the sum of cache-eligible reads plus cache-eligible writes.

Cache Read Hit Percent: The number of cache read hits divided by cache-eligible reads.

Cache Write Hit Percent: The number of cache write hits divided by the number of cache-eligible writes.

Connect Time Percent: The ratio of connect time divided by MSR (millisecond response time), expressed as a percentage.

Connect Time: Represents time in milliseconds spent locating data on the track and transferring between the DASD device and central storage.

Dataset Name: The MVS™ dataset name.

Device Active Only Percent: The average device active only time per I/O in milliseconds.

Device Active Only Time: That portion in milliseconds of total MSR (millisecond response time) that occurs after 'CHANNEL-END' but before 'DEVICE-END'.

Device Address: The MVS™ device number, assigned by HCD, associated with the device on which the dataset is located.

Disconnect Time Percent: The ratio of disconnect time divided by MSR (millisecond response time), expressed as a percentage.

Eligible Cache Percent: The percentage of total I/Os that were cache eligible.

Total I/O Count: The total number of I/Os (read and write) processed for the dataset.

Interrupt Delay Time: The average interrupt delay time for an I/O. This is the elapsed time from the completion of an I/O until z/OS issues test subchannel to retrieve the results.

Interrupt Delay Time Percent: The percent of response time due to interrupt delay time. This percentage is based on the average volume interrupt delay time on which the dataset resides.

I/O Per Second: The average number of I/O operations per second for this dataset during this interval.

IOSQ Time Percent: The ratio of IOSQ time divided by MSR (millisecond response time), expressed as a percentage.

IOSQ Time: The average time in milliseconds that an I/O waits because the device is already busy (signified by the UCBBUSY bit for a device being on).

JOB_START: Job Start value.

Jobname: The MVS™ JOBNAME that is using this dataset.

Low SC Objective: The lower value of direct or sequential storage class MSR (millisecond response time) objective.

Managed System: The SMS name of this system.

NON_C_READ: Non C Read value.

NON_C_WRIT: Non C Writ value.

Pend Time Percent: The ratio of pending time divided by MSR (millisecond response time), expressed as a percentage.

Pend Time: The time in milliseconds that I/O is delayed in the path to the device. Pending time can be attributable to the channel or control unit path being busy.

Read Disconnect Time: The average time in milliseconds that read I/O waits because the device is not connected to the channel.

Read I/O Count: Total number of read I/Os processed for the dataset.

Read Write Percent: The sum of cache-eligible reads plus non-cache reads divided by the total I/O.

Response Time Index: The total response time in milliseconds for this dataset during this interval.

Response Time: The total millisecond response time for this dataset during this interval.

SC Direct MSR Objective: The target direct response time specified for this storage class. This value is used during allocation to select a volume with acceptable performance characteristics. This value also influences caching.

SC Sequential MSR Objective: A construct in the SMS storage class that is used during allocation and indicates the desired response time for sequential accesses to this dataset.

Storage Class Name: The name of the SMS storage class to which this dataset belongs.

Time Stamp: The date and time used to control the scope of the query.

Total Disconnect Time: The average time in milliseconds that all I/O (read and write) waits because the device is not connected to the channel.

Volume: The six-character name of the DASD volume where this dataset is located.

Write I/O Count: Total number of write I/Os processed for the dataset.

Write Disconnect Time: The average time in milliseconds that write I/O waits because the device is not connected to the channel.

Dataset SMS

The Dataset SMS attributes provide information about the SMS constructs associated with a dataset. Additionally, the number of stripes, if applicable, is provided.

Data Class: The SMS data class. This value describes attributes and space requirements for SMS managed datasets. Data Classes contain attributes that govern dataset allocation. These attributes are set by the storage administrator and users cannot change them.

Dataset Name: The MVS™ dataset name.

Managed System: The SMS name of this system.

Management Class: The SMS management class. Management class is a collection of management attributes, defined by the storage administrator, used to control the release of allocated but unused space; to control the retention, migration and backup of datasets; to control the retention and backup of aggregate groups; to control the retention, backup and class transition of objects.

Storage Class: The SMS storage class.

A collection of storage attributes that identify performance goals and availability requirements, defined by the storage administrator, used to select a device that can meet those goals and requirements.

Stripe Count: The number of stripes which comprise this dataset.

A stripe is the portion of a striped dataset that is located on one volume. The records in that portion are not always logically consecutive. The system distributes records among the stripes such that the volumes can be read from or written to simultaneously.

Dataset Space

The Dataset Space attributes provide information about the space allocation and usage for a dataset.

Block Size: The physical record size for storage of data in a dataset and transfer of that data during an I/O operation. For a DASD dataset, block size is maintained in the dataset control block (DSCB).

Because more data is being transferred, large block sizes result in higher connect time. Larger blocks with fewer buffers might be appropriate for sequential datasets. Consider using DASD cache control units for heavily used, read-only sequential datasets.

Random access files need smaller block sizes with a larger number of index buffers (when indexed). For PDS libraries, block sizes should usually reflect average number size.

Dataset Name: The 44-character MVS™ name of this dataset.

Dataset Type: The dataset organization type for this dataset. Values are as follows:

Direct_Access
Direct_Access_Unmovable
Extended_Seq_Unmovable

Extended_Sequential
Hierarchical_File
ISAM
ISAM_Unmovable
Partitioned
Partitioned_Unmovable
PDS_Extended
Physical_Seq_Unmovable
Physical_Sequential
VSAM
VSAM_Unmovable
Unknown

Date Created: The date that the current DASD space for this dataset was first allocated.

Date Referenced: The most recent date when this dataset was opened.

Expiration Date: A date stamp kept on the DASD device where this dataset is allocated that indicates when this dataset is no longer required. This date stamp is written when the dataset is first allocated. Attempting to delete the dataset before this date results in a message to the MVS™ system operator.

Extended Address Space Eligible: Dataset is eligible to use Extended Address Space on an Extended Address Volume.

Logical Record Length: The DCB parameter that specifies the length, in bytes, of each record in the dataset (fixed format), or the maximum record length (variable or undefined format).

Managed System: The SMS name of this system. An internal node name used to identify the system from which this row of data was obtained.

Number of Extents: The number of separate areas of DASD space allocated to this dataset.

Record Format: The DCB parameter that specifies the type and blocking characteristics of records in the dataset. These values are:

F Fixed Length Records
FB Fixed Length Blocked Records
V Variable Length Records
VB Variable Length Blocked Records
U Undefined Length Records
D ISCII/ASCII Variable Length Records

See the IBM® publication "MVS™ DFP 3.3 Using Datasets" for more information on this and other DCB parameters.

Record Secondary Allocation Unit: The allocation unit for secondary space requests.

Record Secondary Allocation: The amount of space allocated to this dataset in secondary allocations.

Secondary Allocation Quantity: The number of space units which are allocated for each dataset extent beyond the primary allocation.

Secondary Allocation Unit: The unit in which space is allocated for each dataset extent beyond the primary allocation.

Tracks Allocated: The number of tracks reserved for the use of this dataset.

Tracks Used Percent: The ratio of tracks used divided by the number of tracks allocated. The value is obtained by dividing the value of the **Tracks Used** attribute by the value of the **Tracks Allocated** attribute.

Tracks Used: The number of tracks that actually contain data for this dataset.

Unused Tracks: The number of tracks allocated to the dataset that are not being used. For VSAM datasets, this is the number of tracks that are not being used at the current high RBA.

Using Cylinder Managed Space: Dataset is using Cylinder Managed Space on an Extended Address Volume.

Volume Count: The number of DASD volumes used to contain this dataset.

Volume Extents: The count of extents for this dataset on this volume.

Volume Sequence: The sequence number of the volume in a multi-volume dataset.

Volume: The six-character name assigned to this DASD volume associated with this row of data.

VSAM Key Displacement: The distance in bytes of the key from the beginning of the data record. See the IBM® publication "MVS™ DFP 3.3 Using Datasets" for more information on this parameter.

VSAM Key Length: The DCB parameter that specifies the length in bytes (0 - 255) of a record key. See the IBM® publication "MVS™ DFP 3.3 Using Datasets" for more information on this and other DCB parameters.

Dataset Users

The Dataset Users attributes provide the names of applications using or waiting for access to a dataset. Also provided are the name of the application, the system ID, and the type of control (shared or exclusive) that the application requires.

Application: The application name associated with this row.

Application Type: Indicates the type of application.

Application Waiting: Indicates whether the application is waiting or not.

- No = The application is not waiting because it has either exclusive or shared control of the dataset.
- Yes = The application is waiting because control of the dataset is owned by another user.

Control: Indicates whether this application shares this dataset with other applications (Shared) or not (Exclusive).

Dataset Name Index: The name of the dataset that is enqueued.

Managed System: The SMS name of this system.

System: The SMFID (Systems Management Facility ID) of the system on which this application is running.

FICON Director Control Units

The FICON Director Control Units attributes provide information on the port.

Device Number: Switch device number.

Port Address: The address of this port.

Port Number: The number of this port.

Control Unit ID: Control Unit ID as 4 hex digits.

FICON Director Ports

The FICON Director Ports attributes provide information on port activity.

Device Number: Switch device number.

Port Address: The address of this port.

Port Number: The number of this port.

Connection Unit: The type of unit connected to the port.

Connection ID: The connection ID of the channel or first control unit.

Connection Serial Number: The serial number of the unit connected on this port.

Control Units Connected: The number of control units connected to the port.

Average Frame Pacing: The average frame pacing time for the port on the switch in units of microseconds.

Error Count: The count of errors for the port.

Read Bandwidth: The rate in .01 MB/Second that data was read for the port.

Write Bandwidth: The rate in .01 MB/Second that data was written for the port.

Read Frame Size: The average frame size for received data frames in bytes.

Write Frame Size: The average frame size for transmitted data frames in bytes.

FICON Director Summary

The FICON Director Summary attributes provide summary information on ports and port activity.

Device Number: Switch device number.

Logical Switch ID: The logical ID of the switch.

Type: The type of switch.

Model: The model of the switch.

Manufacturer: Manufacturer of the switch.

Plant: Manufacturing plant.

Serial Number: Switch serial number.

Collection Time: Time when data was collected.

Maximum Ports: Maximum number of ports that can be installed.

Installed Ports: The number of ports that are installed.

Max Average Frame Pacing: The maximum average frame pacing time for any port on the switch in units of microseconds.

Maximum Error Count: The maximum count of errors for any port on the switch.

Max Read Bandwidth: The maximum rate in .01 MB/second that data was read for any port on the switch.

Max Write Bandwidth: The maximum rate in .01 MB/second that data was written for any port on the switch.

Min Read Frame Size: The minimum average read frame size for any port on the switch in bytes.

Min Write Frame Size: The minimum average write frame size for any port on the switch in bytes.

Read Bandwidth: The average received .01 MB/second of data for the switch.

Write Bandwidth: The average transmitted .01 MB/second of data for the switch.

S3 App Defs

Internal use only.

Description: This is a user defined description used to identify a group.

Group Definition Data: Group definition bulk data.Format is: ATTRIBUTENAME,OPERATION,VALUE. These fields will be repeated for each Group Definition, maximum of 100 constraints.

Attribute

maximum of 10 characters

Operation

=, >, >=, <>, <, <=

Value

value being compared

Group Name: This is a user defined name used to identify a group.

Group Type: Internal flag specifying the group type.

0

User DASD group

1

Volume Group

Persistance Flag: This is the user defined Persistance Flag.

0

No persist after TEMS restart

1

Persist after TEMS restart

S3 Dag Defs

Internal use only.

Description: This is a user defined description used to identify a group.

Group Definition Data: Group definition bulk data.Format is: ATTRIBUTENAME,OPERATION,VALUE. These fields will be repeated for each Group Definition, maximum of 100 constraints.

Attribute

maximum of 10 characters

Operation

=, >, >=, <>, <, <=

Value

value being compared

Group Name: This is a user defined name used to identify a group.

Group Type: Internal flag specifying the group type.

0

User DASD group

1

Volume Group

Persistance Flag: This is the user defined Persistance Flag.

0

No persist after TEMS restart

1

Persist after TEMS restart

S3 DASD Defs

Internal use only.

Description: This is a user defined description used to identify a group.

Group Definition Data: Group definition bulk data.Format is: ATTRIBUTENAME,OPERATION,VALUE. These fields will be repeated for each Group Definition, maximum of 100 constraints.

Attribute

maximum of 10 characters

Operation

=, >, >=, <>, <, <=

Value

value being compared

Group Name: This is a user defined name used to identify a group.

Group Type: Internal flag specifying the group type.

0

User DASD group

1

Volume Group

Persistance Flag: This is the user defined Persistance Flag.

0

No persist after TEMS restart

1

Persist after TEMS restart

S3 Group Defs

The S3 Group Defs attributes provide definition information about the group.

Description: This is a user defined description used to identify a group.

Disable Collection Flag: This is the user defined Disable Collection Flag.

0

Enable Collection for current group

1

Disable Collection for current group

Group Definition Data: Group definition bulk data.Format is: ATTRIBUTENAME,OPERATION,VALUE. These fields will be repeated for each Group Definition, maximum of 100 constraints.

Attribute

maximum of 10 characters

Operation

=, >, >=, <>, <, <=

Value

value being compared

Group Name: This is a user defined name used to identify a group.

Group Type: Internal flag specifying the group type.

0

User DASD group

1

Volume Group

Persistence Flag: This is the user defined Persistence Flag.

0

No persist after TEMS restart

1

Persist after TEMS restart

S3 Device Class Details

The Device Class Details attributes provide detailed information on devices for specific Device Class.

Device Class: Device class.

Device Number: Device address in hexadecimal.

Device Ready: Device ready.

Device Status: Device status.

Device Type: Device type.

Managed System: Internal node name used to identify the system from which this row of data was obtained.

Volser: The volume serial number of a disk device.

S3 Device Summary

The Device Summary attributes provide summarized information on number of devices for different Device Class.

Device Class: Device class.

Managed System: Internal node name used to identify the system from which this row of data was obtained.

Not Ready Device Number: Number of not ready devices for specific class.

Offline Device Number: Number of offline devices for specific class.

Online Device Number: Number of online devices for specific class.

Ready Device Number: Number of ready devices for specific class.

Total Device Number: Total number of devices for specific class.

S3 Storage Group Configuration

These attributes provide information on how and which volumes to set and add for specific Storage Group configurations.

Managed System: An internal node name used to identify the system.

Configuration Name: A user defined name used for a storage group configuration definition.

Description: A user defined description used for a storage group configuration definition.

Model Name: A user model name used for a storage group configuration definition.

Storage Groups With Issues: A number of storage groups which have issues, such as no free slots.

Storage Groups With Warnings: A number of storage groups which have warnings, such as a few free slots.

Storage Groups: A storage group used for the current configuration.

Volume Pools: A volume pools used for the current configuration.

S3 User Storage Groups

The User Storage Groups attributes provide summarized information on custom groups for SMS Storage groups.

Free Space Cylinders: The total amount of free space accumulated in all specified storage groups in cylinders. This value does not include the space on volumes in an SMS storage group that have not been converted.

Free Space GB: The total amount of free space accumulated in all specified storage groups in GB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Free Space MB: The total amount of free space accumulated in all specified storage groups in MB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Free Space Percent: The average free space across volumes in all specified groups expressed as a percentage. This value does not include volumes which are unavailable for new dataset allocation.

Free Space Tracks: The total amount of free space accumulated in all specified storage groups in tracks. This value does not include the space on volumes in an SMS storage group that have not been converted.

Inactive Freespace Cylinders: The total amount of inactive space accumulated in all specified storage groups in cylinders. This value does not include the space on volumes in an SMS storage group that have not been converted.

Inactive Freespace Tracks: The total amount of inactive space accumulated in all specified storage groups in tracks. This value does not include the space on volumes in an SMS storage group that have not been converted.

Inactive Freespace GB: The total amount of inactive space accumulated in all specified storage groups in GB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Inactive Freespace MB: The total amount of inactive space accumulated in all specified storage groups in MB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Largest Free Extent MB: The size of the largest free extent within all specified storage groups specified in the group, in MB. This value does not include volumes which are unavailable for new dataset allocation.

Low Volume Free Space %: The lowest free space value for a volume within all specified storage groups in the group, expressed as a percentage. This value does not include volumes which are unavailable for new dataset allocation.

Managed System: Internal node name used to identify the system from which this row of data was obtained.

Storage Group Count: Storage group count within the group.

Total Space Cylinders: The total amount of space accumulated in all specified storage groups in cylinders. This value does not include the space on volumes in an SMS storage group that have not been converted.

Total Space GB: The total amount of space accumulated in all specified storage groups in GB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Total Space MB: The total amount of space accumulated in all specified storage groups in MB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Total Space Tracks: The total amount of space accumulated in all specified storage groups in tracks. This value does not include the space on volumes in an SMS storage group that have not been converted.

Total Volumes: Volume count accumulated in all specified groups.

Track Managed Free Space Cylinders: The total amount of track managed free space in all specified storage groups in cylinders. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Free Space GB: The total amount of track managed free space in all specified storage groups in GB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Free Space MB: The total amount of track managed free space in all specified storage groups in MB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Free Space Percent: The lowest percent of track managed area free space in all specified storage groups. This value does not include volumes which are unavailable for new dataset allocation.

Track Managed Free Space Tracks: The total amount of track managed free space in all specified storage groups in tracks. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Inact Freespace Cylinders: The total amount of track managed inactive space in all specified storage groups in cylinders. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Inact Freespace GB: The total amount of track managed inactive space in all specified storage groups in GB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Inact Freespace MB: The total amount of track managed inactive space in all specified storage groups in MB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Inact Freespace Tracks: The total amount of track managed inactive space in all specified storage groups in tracks. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Largest Free Extent MB: Size of the largest free track managed area extent within all specified storage groups in MB. This value does not include volumes which are unavailable for new dataset allocation.

Track Managed Total Space Cylinders: The total amount of track managed space in all specified storage groups in cylinders. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Total Space GB: The total amount of track managed space in all specified storage groups in GB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Total Space MB: The total amount of track managed space in all specified storage groups in MB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Total Space Tracks: The total amount of track managed space in all specified storage groups in tracks. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Used Space MB: The total amount of track managed used space in all specified storage groups in MB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Used Space Cylinders: The total amount of track managed used space in all specified storage groups in cylinders. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Used Space Tracks: The total amount of track managed used space in all specified storage groups in tracks. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Used Space GB: The total amount of track managed used space in all specified storage groups in GB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Used Space Cylinders: The total amount of used space accumulated in all specified storage groups in cylinders. This value does not include the space on volumes in an SMS storage group that have not been converted.

Used Space GB: The total amount of used space accumulated in all specified storage groups in GB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Used Space MB: The total amount of used space accumulated in all specified storage groups in MB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Used Space Tracks: The total amount of used space accumulated in all specified storage groups in tracks. This value does not include the space on volumes in an SMS storage group that have not been converted.

S3 Volume Event

The Volume Events attributes provide information on when Volume Online and Volume Offline events happened.

Event: An event for the volume.

Is in History: A flag which indicates if the record was added to history.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

SMFID: SMF ID of the system on which this event happened.

Timestamp: Date and time of the event.

Volume: The volume serial number.

S3 Volume Model

These attributes provide definition for models that are used during volume initialization.

Managed System: An internal node name used to identify the system.

Model Name: A user defined name used for a model definition.

Description: A user defined description used for a model definition.

VTOC Size: A VTOC size in tracks used for a model definition.

Device Model: A device model used for a model definition.

S3 Volume Pool

These attributes provide definition for volumes included or excluded from the pools.

Managed System: An internal node name to identify the system.

Pool Name: A user defined name to identify a pool.

Description: A user defined description to identify a pool.

Applicable Volumes: A number of volumes applicable in defined pool.

Total Volumes: A total number of volumes.

Volumes: A volume specified for the pool.

Exclude Volumes: A volume excluded for the pool.

Hitachi Data Systems 2107 Array Configuration

The Hitachi Data Systems 2107 Array Configuration attributes provide configuration information about the arrays in a Hitachi Data Systems 2107 storage facility.

Array ID: The array Identifier.

Array Type: The RAID type of the array.

Array Width: The width of the array (number of Disk Drive Modules in the array).

DDM Capacity: The capacity in GB of each Disk Drive Module in the array.

DDM Speed: The speed, in 1000 RPM, of the Disk Drive Modules in the array.

DDM Throttling: Indicates if a near-line DDM in the array is throttling performance due to temperature or workload.

Managed System: The internal node name used to identify the system from which this row of data was obtained.

Parity Group: The name of the parity group. This is for HDS storage facilities only.

RAID Degraded: Indicates if one or more DDMs in the array need rebuilding.

Rank ID: The rank identifier.

RPM Exception: Indicates if a DDM with a slower RPM than the normal array DDMs is a member of the array as a result of a sparing action.

SSD Array: Indicates if the array is a solid state device.

Storage Facility ID: The storage facility identifier of the device (Type-Model-Serial Number).

Time Stamp: The date and time associated with this row of data.

Hitachi Data Systems 2107 Configuration

The Hitachi Data Systems 2107 Configuration attributes provide configuration information about a Hitachi Data Systems 2107 storage facility.

Available Cache: The available cache in MB.

Configured NVS: The configured Non-Volatile Storage in MB.

Logical Subsystems: The total number of logical subsystems defined.

Logical Volumes: The total number of logical volumes defined.

Managed System: The internal node name used to identify the system from which this row of data was obtained.

Number of arrays with problems: The number of arrays with error conditions of Raid Degraded (one or more array members need rebuilding), DDM Throttling (a near-line DDM in the array is throttling performance due to temperature or workload), or RPM Exception (a DDM with an slower RPM than the normal array DDMs is a member of the array as a result of a sparing action).

Pinned NVS: The pinned Non-Volatile Storage in MB.

Storage Facility ID: The storage facility identifier of the device (Type-Model-Serial Number).

Storage Facility Serial Number: The storage facility serial number of the device.

Time Stamp: The date and time associated with this row of data.

Total Arrays: The total number of arrays defined.

Total Cache: The total cache in MB.

Total Parity Groups: The number of parity groups configured on an HDS 2107 storage facility.

Total Ranks: The total number of ranks defined.

Total SSD Arrays: The total number of solid state device arrays.

Hitachi Data Systems 2107 Parity Group

The Hitachi Data Systems 2107 Parity Group attributes provide performance information about the parity groups in an HDS storage facility.

Average Read Resp: The average read response time to the parity group over the interval, in milliseconds.

Average Write Resp: The average write response time to the parity group over the interval, in milliseconds.

Avg Bytes/ Read Op: The average number of bytes per read operation.

Avg Bytes/ Write Op: The average number of bytes per write operation.

Bytes Read Per Sec: The rate per second at which bytes were read from the parity group over the interval.

Bytes Written Per Sec: The rate per second at which bytes were written to the parity group over the interval.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Number of Arrays: The total number of arrays that make up the parity group.

Number of arrays with problems: The number of arrays with error conditions of Raid Degraded (one or more array members need rebuilding), DDM Throttling (a near-line DDM in the array is throttling performance due to temperature or workload), or RPM Exception (a DDM with an slower RPM than the normal array DDMs is a member of the array as a result of a sparing action).

Parity Group: The name of the parity group for HDS storage facilities only.

Rank ID: The identification number of the rank.

Read Ops Per Sec: The rate per second that read operations were issued to the rank over the interval.

SSD Rank: Indicates that the parity group is backed by SSD arrays.

Storage Facility ID: The storage facility identifier for the device (Type-Model-Serial Number).

Time Stamp: The date and time associated with this row of data.

Total Bytes Read: The total number of bytes read from the parity group over the interval.

Total Bytes Written: The total number of bytes written to the parity group over the interval.

Total Read Ops: The total number of read operations issued to the parity group over the interval.

Total Read Resp: The total read response time to the parity group over the interval, in milliseconds.

Total Write Ops: The total number of write operations issued to the parity group over the interval.

Total Write Resp: The total write response time to the parity group over the interval, in milliseconds.

Write Ops Per Sec: The rate per second that write operations were issued to the parity group over the interval.

HSM CDS

The HSM CDS attributes provide space allocation and utilization information for each of the HSM control datasets.

DDNAME: The DDNAME of this dataset or cluster. Specific DDNAMEs are associated with specific HSM functions:

DDNAMEs MIGCAT, MIGCAT2-4, BAKCAT, and BAKCAT2-4 are associated with migration.
DDNAME OFFCAT is associated with the OCDS, which contains information about tape backup volumes and tape migration level 2 volumes.
DDNAME Journal is associated with the HSM Journal.
The journal contains a duplicate of each record in the control datasets that has changed since the BCDS, and OCDS were last backed up.

Display Order: The order in which the datasets are displayed.

HSMplex Name: The name of the HSMplex to which the HSM hosts on this system belong.

Managed System: The SMS name of this system.

Number of Extents Data Component: The number of extents in the data component of this cluster. A high number of extents can indicate a need for re-allocation or reorganization of this component. Use IDCAMS LISTCAT to review splits.

Number of Extents Index Component: The number of extents in the index component of this cluster. A high number of extents can indicate a need for re-allocation or reorganization of this component. Use IDCAMS LISTCAT to review splits.

Percent Available Space Data Component: The amount of available space in an HSM control dataset, expressed as a percentage of the total available space. This value is based on the high allocated RBA and the available space in the component. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Percent Available Space Index Component: The amount of available space in an HSM control dataset, expressed as a percentage of the total available space. This value is based on the high allocated RBA and the available space in the component. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Percent Free Space Data Component: The amount of unused space in an HSM control dataset, expressed as a percentage of the total available space. This is based on the high allocated RBA and high used RBA. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Percent Free Space Index Component: The amount of unused space in an HSM control dataset, expressed as a percentage of the total available space. This is based on the high allocated RBA and high used RBA. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Size (KB) Data Component: The total amount of space, in units of 1024 bytes, allocated to this component. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Size (KB) Index Component: The total amount of space, in units of 1024 bytes, allocated to this component. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

HSM Common Storage

The HSM Common Storage attributes provide details about common service area storage allocation, including information regarding migration work elements (MWE).

Active Limit Percent: DFSMShsm™ active limit for the common service area storage allocation.This limit is a percent of the specified maximum limit. The default is 90% of the specified maximum limit. After the specified percentage of common service area has been allocated and DFSMShsm™ is active, only batch WAIT MWEs are added to the common service area queue.

HSMplex Name: The name of the HSMplex to which the HSM hosts on this system belong.

Inactive Limit Percent: DFSMShsm™ inactive limit for the common service area storage allocation. This limit is a percent of the specified maximum limit. After the percent of the specified maximum limit has been allocated and DFSMShsm™ is inactive, no MWEs are added to the common service area queue. The default is 30% usage of the maximum limit common service area.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Maximum Common Storage: Maximum amount of common storage area (in bytes) to be allocated to all MWEs. After the limit has been reached, no MWEs are added to the common service area queue. The default is 102400 (100k) bytes.

Maximum MWE Address Space: Maximum number of MWEs to allow per address space. If any NOWAIT MWEs are added to the common service area queue for this address space after the maximum number has been reached, they are flagged to indicate the common service area storage allocated to the MWE is to be freed when the ARCCTL has copied the MWE into DFSMShsm™ address space. The default is 4 MWEs.

Used Common Storage: The number of bytes of the common storage area that are currently being used.

HSM Cross System CRQ Hosts

The HSM Cross System CRQ Hosts attributes provide detailed data on the HSM Host CRQ.

CRQplex Base Name: The base name of the CRQplex to which the HSM host belongs (or n/a, if not in a CRQplex).

HSM Host ASID: The ASID of the HSM host address space.

HSM Host CRQ Held: Indicates whether CRQ functions are held for this host, including n/a unless the state is connected.

HSM Host CRQ Lost Connection: Indicates whether the host lost its connection to the Coupling Facility.

HSM Host CRQ Recall Held: Indicates that recall from the CRQ is held for this host, including n/a unless the state is connected. (I = Internal; E= External; N= No; B = Both; and U = n/a.)

HSM Host CRQ Recall Place Held: Indicates that recall place from the CRQ is held for this host, including n/a unless the state is connected. (I = Internal; E= External; N= No; B = Both; and U = n/a.)

HSM Host CRQ Recall Select Held: Indicates that recall select from the CRQ is held for this host, including n/a unless the state is connected. (I = Internal; E= External; N= No; B = Both; and U = n/a.)

HSM Host CRQ State: The state of the host in relation to the CRQ. (0=Unconnected; 1=Connecting; 2=Connected; 3=Disconnecting; 4=Retry; 5=Failed; and 100=n/a)

HSM Host CRQ Structure Init: Indicates whether the CRQ structure has been initialized.

HSM Host CRQ Task Abended: Indicates that a common queue task has abended.

HSM Host CRQ Tasks Identified: Indicates whether the host CRQ tasks have been identified.

HSM Host CRQ Tasks Started: Indicates whether the host CRQ tasks have been started.

HSM Host CRQ User Connect: Indicates user-specified common queue connect.

HSM Host CRQ User Disconnect: Indicates user-specified common queue disconnect.

HSM Host ID: The ID of the HSM host.

HSM Host Name: The name of the HSM host started task.

HSMplex Name: The name of the HSMplex to which the HSM hosts on this z/OS image belong.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Structure Use Terminating: Indicates whether the usage of the structure is terminating.

System: The SMF ID of the system on which the HSM host is active.

Threshold Exceeded: Indicates whether the high threshold for the structure was exceeded.

HSM Cross System CRQplex

The HSM Cross System CRQplex attributes provide detailed data on the CRQplex.

Active Requests: The number of active requests on the CRQ.

CRQplex Base Name: The base name of the CRQplex.

Current Elements: The number of elements currently in use.

Current Entries: The number of entries currently in use.

Element Percent Full: The percent of elements in the CRQ currently being used, displayed in tenths of a percent.

Entry Percent Full: The percent of entries in the CRQ currently being used, displayed in tenths of a percent.

HSM Hosts Connected: The number of HSM hosts connected to this CRQ.

HSM Hosts Not Connected: The number of HSM hosts in this CRQplex that are not in a connected state.

HSMplex Name: The name of the HSMplex to which all HSM hosts in this CRQplex belong.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Maximum Elements: The maximum number of elements available.

Maximum Entries: The maximum number of entries available.

Oldest Request Age: The number of minutes the oldest request has waited to be completed. A large number of queued requests or a large completion time can be caused by an insufficient number of tasks started of this function type, a high rate of requests, or this function type being held.

Queued Requests: The number of requests on the CRQ waiting for execution.

Total Requests: The number of requests on the CRQ.

HSM CRQ Hosts

The HSM CRQ Hosts attributes provide detailed data on HSM CRQ hosts.

CRQplex Base Name: The base name of the CRQplex to which the HSM host belongs (or n/a, if not in a CRQplex).

HSM Host ASID: The ASID of the HSM host.

HSM Host CRQ Held: Indicates whether the function is held for this host, including n/a unless the state is connected.

HSM Host CRQ Lost Connection: Indicates whether the host lost its connection to the CF.

HSM Host CRQ Recall Held: Indicates that recall from the CRQ is held for this host, including n/a unless the state is connected. (I = Internal; E= External; N= No; B = Both; and U = n/a.)

HSM Host CRQ Recall Place Held: Indicates that recall place from the CRQ is held for this host, including n/a unless the state is connected. (I = Internal; E= External; N= No; B = Both; and U = n/a.)

HSM Host CRQ Recall Select Held: Indicates that recall select from the CRQ is held for this host, including n/a unless the state is connected. (I = Internal; E= External; N= No; B = Both; and U = n/a.)

HSM Host CRQ State: The state of the host in relation to the CRQ. (0=Unconnected; 1=Connecting; 2=Connected; 3=Disconnecting; 4=Retry; 5=Failed; and 100=n/a)

HSM Host CRQ Structure Init: Indicates whether the CRQ structure has been initialized.

HSM Host CRQ Task Abended: Indicates that a common queue task has abended.

HSM Host CRQ Tasks Identified: Indicates whether the host CRQ tasks have been identified.

HSM Host CRQ Tasks Started: Indicates whether the host CRQ tasks have been started.

HSM Host CRQ User Connect: Indicates user-specified common queue connect.

HSM Host CRQ User Disconnect: Indicates user-specified common queue disconnect.

HSM Host ID: The ID of the HSM host.

HSM Host Name: The name of the started task HSM host.

HSMplex Name: The name of the HSMplex to which the HSM hosts on this z/OS image belong.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Structure Use Terminating: Indicates whether the usage of the structure is terminating.

System: The SMF ID of the system on which the HSM host is active.

Threshold Exceeded: Indicates whether the high threshold for the structure was exceeded.

HSM CRQ Requests

The HSM CRQ Requests attributes provide detailed data on CRQ requests, including status.

Aggregate Group Name: The name of the HSM aggregate group to which this dataset belongs.

CRQplex Base Name: The base name of the CRQplex.

Dataset Name: The name of the dataset being recalled.

HSMplex Name: The name of the HSMplex to which the HSM hosts in this CRQplex belong.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Originating TEMS Node: The node of the request's originating TEMS. Used for routing toolkit requests.

Processor HSM Host ID: The ID of the HSM host that is processing the request (blank for requests that are not active).

Processor HSM Host Name: The name of the HSM host started task that is processing the request (blank for requests that are not active).

Processor Origin Node: The Origin node associated with the HSM host processing the request. This is used to direct toolkit actions to the correct node for execution.

Processor System: The Systems Management Facility identifier of the system on which the HSM host processing the request is executing.

Processor TEMS Node: The name of the TEMS node on the z/OS image on which the request is processing if active (otherwise, blanks). Used for routing toolkit requests.

Queue Time: The time when DFHSM received this request.

Request Age: The number of minutes this request has waited to be serviced. A large number of queued requests or a large wait time can be caused by an insufficient number of tasks started of this function type, a high rate of requests, or this function type being held.

Request Number: The number assigned to this request.

Request Status: The status of the request, including Queued, Active, Abended, Complete and Cancelled. A queued request has been received and is waiting for higher priority requests to complete. An active request is currently being serviced. An abended request has completed unsuccessfully. A complete request has completed successfully. A cancelled request was cancelled before it completed

Request Type: The type of request, either dataset or volume.

Request Volume: The specific DASD volume related to this request.

Requester HSM Host ID: The ID of the HSM host that originated the request.

Requester HSM Host Name: The name of the HSM host started task that originated the request.

Requester Origin Node: The Origin node associated with the request originator. This is used to direct toolkit actions to the correct node for execution.

Requester System: The Systems Management Facility identifier of the system on which the requesting HSM host is executing.

Requester TEMS Node: The TEMS node assocated with the request originator, used to direct toolkit actions to the correct TEMS for execution.

Requesting Jobname: The name of job/started task/TSO userid which submitted this request.

Requesting User ID: The user ID associated with this request.

User Waiting: Indicates whether this is a synchronous or asynchronous request

HSM CRQplex

The HSM CRQplex attributes provide detailed status of requests.

Active Requests: The number of active requests on the CRQ.

CRQplex Base Name: The base name of the CRQplex to which the HSM host belongs (or n/a, if not in a CRQplex).

Current Elements: The number of elements currently in use.

Current Entries: The number of entries currently in use.

Element Percent Full: The percent of elements in the CRQ currently being used, displayed in tenths of a percent.

Entry Percent Full: The percent of entries in the CRQ currently being used, displayed in tenths of a percent.

HSM Hosts Connected: The number of HSM hosts connected to this CRQ.

HSM Hosts Not Connected: The number of HSM hosts in this CRQplex that are not in a connected state.

HSMplex Name: The name of the HSMplex to which the HSM hosts on this z/OS image belong.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Maximum Elements: The maximum number of elements available.

Maximum Entries: The maximum number of entries available.

Oldest Queue Time: The date and time when the oldest request was submitted.

Oldest Request Age: The number of minutes the oldest request has waited to be completed. A large number of queued requests or a large completion time can be caused by an insufficient number of tasks started of this function type, a high rate of requests, or this function type being held.

Queued Requests: The number of requests on the CRQ waiting for execution.

Total Requests: The number of requests on the CRQ.

HSM Dataset Recover Version

The HSM Dataset Recover Version attributes provide summary information regarding the backup versions of HSM datasets.

Backup Date Time: The date and time of the backup.

Catalog Status: The catalog status of the dataset.

Generation: Generation number of the backup.

Version: Version number of the backup.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

HSM Function Details

The HSM Function Detail attributes provide a variety of details associated with the HSM function.

Automatic Requests: The number of DFHSM requests of this function type issued automatically. This metric is extracted from the DFHSM daily statistics record.

Average Allocate Time: The average allocation time (in seconds) for a request to be satisfied on this HSM host.

Average Completion Time: The average completion time (in seconds) for a request to be satisfied on this HSM host.

Average Process Time: The average processing time (in seconds) for a request to be satisfied on this HSM host.

Average Queue Time: The average queue time (in seconds) for a request to be satisfied on this HSM host.

Failed Requests: Count of failed requests of this function type extracted from the DFHSM daily statistics record.

Function Index: Function type, as extracted from the DFHSM daily statistics record, to which this row of data applies.

Function Type: The function that was executed. Values are as follows:

- Primary_to_Level1
- Level1_to_Level2
- Primary_to_Level2
- Level1_to_Primary
- Level2_to_Primary
- Delete_Migrated_Dataset
- Daily_Backup
- Spill_Backup
- Recovery
- Recycle_Backup_Volume
- Delete_Dataset
- Recycle_L2_Tape_Volumes
- Delete_Backup_Versions

GB Read: GB of data read for this function type, extracted from DFHSM's daily statistics record.

GB Written: GB of data written for this function type, extracted from DFHSM's daily statistics record.

HSM Host Name: The name of the HSM host. This is the started task name.

HSM_Host_ASID: The ASID of the HSM host.

HSM_Host_ID: The ID of the HSM host.

HSMplex Name: The name of the HSMplex to which the HSM hosts on this system belong.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

MB Read: Total MB read for this function type, extracted from the DFHSM daily statistics record. (_=-2; n/a=-1)

MB Written: Total MB written for this function type, extracted from the DFHSM daily statistics record. (_=-2; n/a=-1)

Tracks Read: Total DASD tracks read for this function type, extracted from the DFHSM daily statistics record.

Tracks Written: Total DASD tracks written for this function type, extracted from the DFHSM daily statistics record.

User Requests: Number of requests initiated by user command for this function type as extracted from the DFHSM daily statistics record on this HSM host.

HSM Function Statistics

The HSM Function Statistics attributes provide summary statistics for the HSM activities, such as movement of datasets from level 1 to level 2 archives.

Automatic Requests: The number of DFHSM requests issued automatically by this function type, as extracted from the DFSMShsm daily statistics record. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Average Allocate Time: The average allocation time (in seconds) for a request to be satisfied across all HSM hosts.

Average Completion Time: The average completion time (in seconds) for a request to be satisfied across all HSM hosts.

Average Process Time: The average processing time (in seconds) for a request to be satisfied across all HSM hosts.

Average Queue Time: The average queue time (in seconds) for a request to be satisfied across all HSM hosts.

Failed Requests: Represents the count of failed requests of this function type extracted from the DFSMShsm daily statistics record. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Function Index: The function type, as extracted from the DFSMShsm daily statistics record, to which this row of data applies.

Function Type: The type of function used. Valid types are as follows:

- Daily_Backup
- Delete_Backup_Versions
- Delete_Dataset
- Delete_Migrated_Dataset
- Level1_to_Level2
- Level1_to_Primary
- Level2_to_Primary
- Primary_to_Level1
- Primary_to_Level2
- Recovery
- Recycle_Backup_Volume
- Recycle_L2_Tape_Volumes
- Spill_Backup

Function Type String: The function type displayed as a string.

GB Read: GB of data read for this function type, extracted from DFHSM's daily statistics record.

GB Written: GB of data written for this function type, extracted from DFHSM's daily statistics record.

HSMplex Name: The name of the HSMplex to which the HSM hosts on this system belong.

Managed System: The SMS name of this system.

MB Read: The total number of MB read for this function type, extracted from the DFSMShsm daily statistics record. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

MB Written: The total number of MB written for this function type, extracted from the DFSMShsm daily statistics record. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Tracks Read: The total number of DASD tracks read for this function type, extracted from the DFSMShsm daily statistics record. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Tracks Written: The total number of DASD tracks written for this function type, extracted from the DFSMShsm daily statistics record. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

User Requests: The number of requests initiated by user command for this function type as extracted from the DFSMShsm daily statistics record. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

HSM Function Summary

The HSM Function Summary attributes provide status and queue information about each of the HSM functions.

Active Requests: The count of requests currently active for this function. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Dataset Requests: The count of requests either queued or active for this function for datasets. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Function: The status of this HSM function (Held or Not Held).

Function Str: The name of the function for this data row. Valid names are Migration, Recall, Backup, Recovery or Dump.

Function Status: The status of this HSM function. Values are Held, Not Held, or Held EOV.

HSMplex Name: The name of the HSMplex to which the HSM hosts on this system belong.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Volume Requests: The count of requests either queued or active for this function for volumes. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Waiting Requests: The count of requests currently waiting for this function. Requests are queued either when the function is held, or maxtasks for this function is reached. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

HSM Host Function Status

The HSM Host Function Status attributes provide detailed status of the HSM host, related to the functions it performs.

Aggregate Backup Status: The status of the HSM host to do aggregate backups. (Not_Held=0; Held=1)

Aggregate Recovery Status: The status of the HSM host to do aggregate recoveries. (Not_Held=0; Held=1)

Audit Activity: The HSM host is currently running audit (Active) or not running audit. (Inactive=0; Active=1)

Audit Status: The status of the audit function on this HSM host. (Not_Held=0; Held=1)

Automatic Backup Status: The status of the HSM host to do automatic backups. (Not_Held=0; Held=1; Held_EOV=2)

Automatic Dump Status: The status of the HSM host to do automatic dumps. (Not_Held=0; Held=1; Held_EOV=2)

Automatic Migration Status: The status of the HSM host to do automatic migrations. (Not_Held=0; Held=1; Held_EOV=2)

Backup Status: The status of the backup function on this HSM host. (Not_Held=0; Held=1; Held_EOV=2)

BCDS Backup Status: Indicates if the backup control dataset is currently being backed up on this HSM host. (Inactive=0; Active=1)

CRQplex Base Name: The base name of the CRQplex to which the HSM host belongs (or n/a, if not in a CRQplex).

Dataset Backup Status: The status of the HSM host to run dataset backup functions. (Not_Held=0; Held=1; Held_DASD=2; Held_Tape=3)

Dump Status: The status of the dump function on this HSM host. (Not_Held=0; Held=1; Held_EOV=2)

Expire Backup Version Status: The status of the host to expire backup versions. (Not_Held=0; Held=1)

HSM Host ASID: The ASID of the HSM host.

HSM Host ID: The ID of the HSM host.

HSM Host Name: The name of the HSM host. This is the started task name.

HSM Host Type: The HSM host type. (Unknown=0; Main=1; Main(Primary)=2; Aux=5; and Aux(Primary)=6)

HSMplex Name: The name of the HSMplex to which the HSM hosts on this z/OS image belong.

HSM Plex Type: The HSM Plex type (Full, Classic or File).

Journal CDS Backup Status: Indicates if the journal control dataset is currently being backed up on this HSM host. (Inactive=0; Active=1)

List Activity: Indicates if the list function is active on this HSM host. (Inactive=0; Active=1)

List Status: The status of the HSM host to run list functions. (Not_Held=0; Held=1)

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Migration CDS Backup Status: Indicates if the migration control dataset is currently being backed up on this HSM host. (Inactive=0; Active=1)

Migration Status: The status of the HSM host to run migration functions. (Not_Held=0; Held=1; Held_EOV=2)

Offline CDS Backup Status: Indicates if the offline control dataset is currently being backed up on this HSM host. (Inactive=0; Active=1)

Recall Status: The status of the HSM host to run recall functions. (Not_Held=0; Held=1; Held_TSO=2)

Recover Status: The status of the HSM host to run recover functions. (Not_Held=0; Held=1; Held_EOV=2)

Recycle Activity: Indicates if the recycle function is active on this HSM host. (Inactive=0; Active=1)

Recycle Status: The status of the HSM host to run recycle functions. (Not_Held=0; Held=1)

Report Activity: Indicates if the report function is active on this HSM host. (Inactive=0; Active=1)

Report Status: The status of the HSM host to run report functions. (Not_Held=0; Held=1)

Tape Copy Status: The status of the HSM host to run tape copy functions. (Not_Held=0; Held=1)

Tape Recall Status: The status of the HSM host to run tape recall functions. (Not_Held=0; Held=1; Held_TSO=2)

Tape Recovery Status: The status of the HSM host to run tape recovery functions. (Not_Held=0; Held=1)

Tape Replace Status: The status of the HSM host to run tape replace functions. (Not_Held=0; Held=1)

HSM Host Request Summary

The HSM Host Request Summary attributes provide a summary of requests queued to an HSM host.

Backup Requests Queued: The number of backup requests queued to this HSM host.

Delete Requests Queued: The number of delete requests queued to this HSM host.

Dump Requests Queued: The number of dump requests queued to this HSM host.

HSM Host ASID: The ASID of the HSM host.

HSM Host ID: The ID of the HSM host.

HSM Host Name: The name of the HSM host. This is the started task name.

HSM Host Type: The HSM host type. Unknown=0; Main=1; Main(Primary)=2; Aux=5; and Aux(Primary)=6

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Migrate Requests Queued: The number of migrate requests queued to this HSM host.

Recall Requests Queued: The number of recall requests queued to this HSM host.

Recovery Requests Queued: The number of recovery requests queued to this HSM host.

HSM Host Status

The HSM Host Status attributes provide a variety of metrics associated with the status of the HSM host.

Avg Dump Alloc Time: The average time (in seconds) spent waiting for allocation for dump requests on this host.

Avg Dump Elapsed Time : The average time (in seconds) it took to process dump requests on this host.

Avg Dump Process Time: The average time (in seconds) spent processing dump requests on this host.

Avg Dump Queue Time: The average time (in seconds) spent on the queue for dump requests on this host.

CRQplex Base Name: The base name of the CRQplex to which the HSM host belongs (or n/a, if not in a CRQplex).

Failed Aggregate Backups: The number of failed aggregate backups on this host.

Failed Dataset Restores: The number of failed dataset restores on this host.

Failed Fast Replic Recovers: The number of failed fast replicate recovers on this host.

Failed Fast Replicat Backups: The number of failed fast replicate backups on this host.

Failed Volume Dump Copies: The number of failed volume dump copies on this host.

Failed Volume Dumps: The number of failed volume dumps.

Failed Volume Restores: The number of volume restores that failed on this host.

HSM Host ASID: The ASID of the host DFSMShsm address space.

HSM Host ID: The ID of the HSM host.

HSM Host Name: The name of the HSM host. This is the started task name.

HSM Host Start Time: Start time and date for the HSM address space.

HSM Host Status: The status of the HSM host. This value is available only if the host has been previously started and another host has not been started in its place. (Active=0; Inactive=1)

HSM Host Type: The HSM host type. (Unknown=0; Main=1; Main(Primary)=2; Aux=5; and Aux(Primary)=6)

HSMplex Name: The name of the HSMplex to which the HSM hosts on this z/OS image belong.

Interval Abends: Count of HSM requests that abended during this interval.

Interval CPU Usage: CPU time consumed by DFHSM to service requests during this interval.

Interval Requests: Total number of HSM requests during this interval.

Interval Start Time: Start time for the current HSM interval.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Recycled Back Up Volumes: Total number of BACKUP volumes recycled during this interval.

Recycled Migration Volumes: Total number of MIGRATION volumes recycled during this interval.

System: The SMF ID of the system on which the HSM host is active.

System Name: The system name of the system.

System Requested Dump Copies: The number of volume dump copies on this host requested by the system.

System Requested Dumps: The number of volume dumps on this host requested by the system.

Total Aggregate Backups: The number of aggregate backups done on this host.

Total Dataset Restores: The number of dataset restores done on this host.

Total Fast Replic Recovers: The number of fast replicate recovers done on this host.

Total Fast Replicat Backups: The number of fast replication backups on this host.

Total Volume Dump Copies: The number of volume dump copies on this host.

Total Volume Dumps: The number of complete volume dumps attempted on this host.

User Requested Dump Copies: The number of volume dump copies on this host requested by a user.

User Requested Dumps: The number of volume dumps on this host requested by a user.

Volume Backups: Total number of requests for volume BACKUP during this interval.

Volume Migrations: Total number of requests for volume MIGRATIONs during this interval.

Volume Restores: Total number of requests for volume restores during this interval.

HSM Private Storage

The HSM Private Storage attributes provide metrics associated with the private storage of the HSM host.

Central Storage Fixed Frames: The number of central storage fixed frames in use by the HSM host. (n/a=-1)

Central Storage Real Frames: The number of frames of central storage in use by the HSM host. (n/a=-1)

ESTOR Frames: The number of expanded storage frames in use by the HSM host. (n/a=-1)

HSM Host ASID: The ASID of the HSM host.

HSM Host ID: The ID of the HSM host.

HSM Host Name: The name of the HSM host. This is the started task name.

HSM Host Type: The HSM host type. (Unknown=0; Main=1; Main(Primary)=2; Aux=5; and Aux(Primary)=6)

HSMplex Name: The name of the HSMplex to which the HSM hosts on this system belong.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

HSM Requests

The HSM Requests attributes provide detailed information about the current and queued requests for HSM functions, such as datasets queued for recall.

Aggregate Group Name: The name of an HSM aggregate group to which this dataset belongs.

An aggregate group is an SMS grouping of datasets where all datasets in the group have identical backup or recovery attributes (such as medium, retention period and destination).

You can define all datasets related to an application in an aggregate group so that they all can be treated alike and at the same time by SMS.

Dataset Name: The name of the dataset for which a request has been made.

Function: The type of request. Each request type has a different queue. The types are as follows: Migration, Recall, Backup, Recovery, Delete, Vol_Recover, and Dump.

Function Index: The HSM function type to which this data row pertains.

HSM Host ID: The ID associated with this HSM host.

HSM Host Name: The name of HSM host associated with the request.

HSMplex Name: The name of the HSMplex to which the HSM hosts on this system belong.

Managed System: The SMS name of this system.

Queue Time: The time that the request was queued.

Request Age: The number of minutes this request has waited to be serviced. A large number of queued requests or a large wait time can be caused by an insufficient number of tasks started of this function type, a high rate of requests or this function type being held.

Requesting Jobname: The name of the job/started task/TSOUSERID that submitted this request.

Request Number: The number assigned to this request. Requests are numbered in ascending sequence by DFHSM as they are received.

Request Status: The status of request. The request status defines the state of the request. The following states are valid:

Queued = Request has been received and is waiting for higher priority requests to complete.
Active = Request is currently being serviced.
Abended = The request has completed unsuccessfully.
Complete = The request has completed successfully.
Cancelled = The request was cancelled before it completed.

Request Time: The time that DFHSM received this request.

Request Type: Denotes whether this request is for a specific volume or specific dataset.

Request Volume: The specific DASD volume related to this request.

Requesting User ID: The user ID associated with this request.

User Waiting: Specifies whether this is a synchronous (waiting) or asynchronous (not waiting) request.

HSM Status

The HSM Status attributes provide information on the status of HSM, including whether it is active, the start time, and some summary activity information.

Active HSM Hosts: The number of active HSM hosts on this z/OS® image.

Avg Dump Alloc Time: The average time (in seconds) that a dump request spent waiting for allocation across all HSM hosts on the z/OS ® image.

Avg Dump Elapsed Time: The average time (in seconds) that a dump took to complete across all HSM hosts on the z/OS ® image.

Avg Dump Process Time: The average time (in seconds) that a dump request spent processing across all HSM hosts on the z/OS ® image.

Avg Dump Queue Time: The average time (in seconds) that a dump request spent waiting for execution across all HSM hosts on the z/OS ® image.

CRQplex Member: Indicates whether at least one HSM host on the z/OS image is in a CRQplex.

Failed Aggregate Backups: The number of failed fast aggregate backups across all HSM hosts on the z/OS ® image.

Failed Dataset Restores : The number of failed dataset restores across all HSM hosts on the z/OS ® image.

Failed Fast Replic Recovers: The number of failed fast replicate recovers across all HSM hosts on the z/OS ® image.

Failed Fast Replicat Backups: The number of failed fast replicate backups across all HSM hosts on the z/OS ® image.

Failed Volume Dump Copies: The number of failed fast volume dump copies across all HSM hosts on the z/OS ® image.

Failed Volume Dumps: The number of failed volume dumps.

Failed Volume Restores: The number of failed volume restores.

HSM ASID: The ASID of the host DFSMShsm address space.

HSM Plex Type: The HSM Plex type (Full, Classic or File).

HSM Start Time: The start time and date for the HSM address space.

HSM Status: The current status of the HSM product. Status can be one of the following:

Inactive = The product is licensed, but not currently running.
Active = The product is licensed, installed, and ready to run.
Notinstalled = The product is not licensed or not defined as part of the system.

HSMplex Name: The name of the HSMplex to which the HSM hosts on this z/OS image belong.

Inactive HSM Hosts: The number of inactive HSM hosts on this z/OS® image.

Interval Abends: The count of HSM requests that abended during this interval. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Interval CPU Usage: The CPU time consumed by HSM to service requests during this interval. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Interval Requests: The total number of HSM requests during this interval. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Interval Start Time: The start time for the current HSM interval.

Managed System: The SMS name of this system.

Recycled Back Up Volumes: The total number of BACKUP volumes recycled during this interval. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Recycled Migration Volumes: The total number of MIGRATION volumes recycled during this interval. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

System Requested Dump Copies: The number of system requested fast volume dump copies taken by all HSM hosts on this z/OS ® image.

System Requested Dumps: The number of system requested dumps across all HSM hosts on the z/OS ®.

Total Aggregate Backups: The number of aggregate backups taken by all HSM hosts on this z/OS ® image.

Total Dataset Restores: The number of dataset restores taken by all HSM hosts on this z/OS ® image.

Total Fast Replic Recovers: The number of fast replicate recovers taken by all HSM hosts on this z/OS ® image.

Total Fast Replicat Backups : The number of fast replicate backups taken by all HSM hosts on this z/OS ® image.

Total Volume Dump Copies: The number of fast volume dump copies taken by all HSM hosts on this z/OS ® image.

Total Volume Dumps: The total number of volume dumps.

User Requested Dump Copies: The number of user requested fast volume dump copies taken by all HSM hosts on this z/OS ® image.

User Requested Dumps: The number of user requested dumps across all HSM hosts on the z/OS ® image.

Version: The version and release level of DFHSM on this system.

Volume Backups: The total number of requests for volume BACKUP during this interval. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Volume Migrations: The total number of requests for volume MIGRATIONs during this interval. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Volume Restores: The number of volume restores.

Logical Control Unit

The Logical Control Unit attributes provide performance information for a logical control unit.

All Channels Busy Percent: The percentage of the time all channel paths to this LCU were busy.

Attached Channels: The total number of channel paths that have defined connections to the LCU.

Average Delay Queue: The average number of delayed requests on the control unit header (CU-HDR). This field is calculated as the accumulated Q length minus the number of enqueued requests divided by the number of enqueued requests.

Channel Path I/O Rate: The Channel Path I/O rate for I/Os to this LCU. This field is calculated as the number of I/O operations accepted on this channel path divided by (sample time minus interval start time).

Contention Rate: The rate per second at which the SAP places delayed I/O requests on the CU-HDR for this LCU. This is done when all paths to the subchannel are busy and at least one path to the control unit is busy.

Control Unit Busy Percent: The percentage of requests deferred due to control unit busy.

DCM Channels Attached: The flag indicating that there are DCM channels attached to this LCU.

DCM Channels Channel Path I/O Rate: The Channel Path I/O rate for I/Os to this LCU. This field is computed as the number of I/O operations accepted on this channel path divided by (sample time minus interval start time).

DCM Channels Control Unit Busy Percent: Shows the relationship for each channel path of the LCU, between requests deferred due to control unit busy and total successful requests serviced by that path. Each CHPID of the LCU measures the distribution of control unit contention. This number is derived from: the number of deferred I/O requests due to **control unit busy** divided by the following combined values:

(the number of deferred I/O requests due to director port busy +
the number of deferred I/O requests due to control unit busy +
the number of successful I/O requests on that path) * 100

DCM Channels Defined: The number of DCM channels defined for this LCU.

DCM Channels Director Port Busy Percent: Indicates director port contention. This value represents the number of times an I/O request was deferred because the director port was busy during the measurement interval. This value is derived from the number of deferred I/O requests due to **director port busy** divided by the following combined values:

(the number of deferred I/O requests due to director port busy +
the number of deferred I/O requests due to control unit busy +
the number of successful I/O requests on that path) * 100

DCM Channels Maximum: The maximum number of DCM channels used for this LCU.

DCM Channels Minimum: The minimum number of DCM channels used for this LCU.

Director Port Busy Percent: The percentage the director port is busy. This field indicates director port contention. This value is the number of times an I/O request was deferred because the director port was busy during the measurement interval.

High DFW Retry Percent: The highest DFW retry percentage of any volume on this LCU.

High Volume MSR: The highest MSR (millisecond response time) of any volume on this LCU.

High Volume Utilization: The highest percentage busy of any volume on this LCU.

I/O Per Second: The channel Path I/O Rate for I/Os to this.

LCU Hex Number: The character representation to the hexadecimal number of this LCU.

LCU Number: The HCD-assigned number for this logical control unit. An LCU is a logical representation of either a single control unit, with or without attached devices, or a group of control units that share devices.

Low Read Hit Percent: The lowest read hit percentage of any volume on this LCU.

Low Write Hit Percent: The lowest write hit percentage of any volume on this LCU.

Managed System: The SMS name of this system.

Maximum HyperPAV Unavailable Alias Percent: The maximum percent of the number of I/Os that could not start for a control unit because a HyperPAV alias was not available.

Maximum I/Os Queued: The maximum number of I/Os queued for a control unit.

NonSMS Volumes: The total number of nonSMS-managed volumes on this LCU.

PDS Record: A single record of trend data retrieved from persistent data store.

SMS Volumes: The total number of SMS-managed volumes on this LCU.

Time Stamp: The date and time associated with this row.

Total Volumes: The count of the total number of volumes in this LCU.

Logical Control Unit Channel Paths

The Logical Control Unit Channel Paths attributes provide performance information for the channels associated with a logical control unit.

Avg Initial Cmd Response Time: The average number of milliseconds of delay that a successfully initiated start or resume function needs until the first command is indicated as accepted by the device.

Channel Path IO Rate: The number of I/O operations per second accepted on this channel path.

CHPID: The channel path identifier. A channel path is the physical connection from the processor to I/O or peripheral devices. Some devices can operate with up to eight paths.

Complex Percent Utilized: The aggregate utilization of this CHPID by all processors in the complex.

Control Units Percent Busy: The control unit utilization is expressed as the percentage of time that all control units attached to the channel path are busy for I/O operations to this LCU.

Control Units: The ID of the physical control units that make up the logical control unit.

DCM Status: The dynamic Channel Path Management Status. Managed or unmanaged, dynamic channel path management provides the capability to dynamically assign channels to control units to respond to peaks in demand.

Description: The type of channel path identified by this CHPID.

Director Port Percent Busy: The percentage of initial selection attempts that were unsuccessful because the director port was busy.

Init Cmd Response Time Ratio: The ratio of the average initial command response time for a chpid compared to the lowest (greater than 0) average initial command response time for a chpid. A ratio greater than 6 is an indication of an out of balance condition.

LCU Number H: The system-assigned identifier of this logical control unit. An LCU is the logical representation of either a single control unit, with or without attached devices, or a group of control units that share devices.

LCU Number: The system-assigned identifier of this logical control unit. An LCU is the logical representation of either a single control unit, with or without attached devices, or a group of control units that share devices.

Managed System: The SMS name of this system.

Path ID: The Channel Path Identifier. A channel path is the physical connection from the processor to I/O or peripheral devices. Some devices can operate with up to eight paths.

Status: The status of this channel path on this LPAR.

LCU Control Unit

The LCU Control Unit attributes provides data to monitor I/O requests for a control unit that are related to the HyperPAV (Hyper parallel access volume) feature.

Control Unit: The physical control unit number.

HyperPAV Requests: The total number of HyperPAV I/O requests for the control unit.

I/Os Queued Alias Unavailable: The number of times an I/O could not start for the control unit because no HyperPAV aliases were available.

LCU Number: The HCD-assigned number for this logical control unit. An LCU is a logical representation of either a single control unit, with or without attached devices, or a group of control units that share devices.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Maximum HyperPAV Aliases Used: The maximum number of HyperPAV aliases in use concurrently.

Maximum HyperPAV Volume Aliases: The maximum number of HyperPAV aliases in use concurrently by a single logical volume.

Maximum I/Os Queued: The maximum number of I/Os concurrently queued for the control unit.

Percent I/Os Queued Alias Unavailable: The percent of I/Os that could not start for the control unit because no HyperPAV aliases were available.

Time Stamp: Date and time this data was reported. Primarily for use in historical reporting.

RLS Appl Dataset Details

The RLS Appl Dataset Details attributes provide performance statistics for each dataset allocated to an application.

Application: The application which is accessing this dataset.

Average Response Time: Average response time for the dataset in .01 milliseconds.

Avg Resp Time Direct Access: Average response time for all direct access requests in .01 milliseconds

Avg Resp Time Sequential Access: Average response time for all sequential access requests in .01 milliseconds.

BMF Dir Acc False Invalids per Minute: Number of BMF direct access false invalids per minute in hundredths.

BMF Direct Acc False Invalids: Number of BMF direct access false invalids.

BMF Direct Access Reqs: Number of BMF direct access requests.

BMF Direct Access Reqs Minute: Number of BMF direct access requests per minute in hundredths.

BMF False Invalid percent: Percent of BMF requests that were false invalids in tenths of a percent.

BMF False Invalids: Number of BMF false invalids.

BMF False Invalids per Minute: Number of BMF false invalids per minute in hundredths.

BMF Requests: Number of BMF requests.

BMF Requests per Minute: Number of BMF requests per minute in hundredths.

BMF Seq Acc False Invalid Min: Number of BMF sequential access false invalids per minute in hundredths.

BMF Seq Acc False Invalids: Number of BMF sequential access false invalids.

BMF Sequential Access Reqs: Number of BMF sequential access requests.

BMF Sequential Access Reqs Min: Number of BMF sequential access requests per minute in hundredths.

CA Splits per Minute: The number of control area splits per minute in hundredths.

CA Splits Percent: The percent of writes that resulted in control area splits.

Cache Set Name: Cache set used by this dataset.

CI Splits per Minute: The number of control interval splits per minute in hundredths.

CI Splits Percent: The percent of writes that resulted in a control interval split.

Collection Interval: Time span over which data was collected in seconds.

Collection Time: Start time when data was collected.

Control Area Splits: The number of control area splits.

Control Interval Splits: The number of control interval splits.

Dataset Name: Dataset name for which the statistics are provided.

Direct Access Reads CR: Number of direct access reads with Consistent Read.

Direct Access Reads CR Minute: Number of direct access reads with CR per minute in hundredths.

Direct Access Reads NRI: Number of direct access reads with No Read Integrity.

Direct Access Reads NRI Minute: Number of direct access reads with NRI per minute in hundredths.

Direct Access Requests Minute: Direct access requests per minute in hundredths.

Direct Access Requests Total: Total number of direct access requests.

Direct Access Writes: Number of direct access writes.

Direct Access Writes Minute: Number of direct access writes per minute in hundredths.

DIWA Lock Percent: Percent of write requests that had Data Insert Work Area (DIWA) lock requests in tenths of a percent.

DIWA Lock Reqs True Contention: Number of Data Insert Work Area (DIWA) lock requests that had true contention.

DIWA Lock Requests: Number of Data Insert Work Area (DIWA) lock requests.

DIWA Lock Requests per Minute: Number of Data Insert Work Area (DIWA) lock requests per minute in hundredths.

DIWA Lock True Contention Minute: Number of Data Insert Work Area (DIWA) lock requests that had true contention per minute in hundredths.

DIWA Lock True Contention Pct: Percent of Data Insert Work Area (DIWA) lock requests that had true contention in tenths of a percent.

ESDS Add To End Lock Requests: Number of ESDS add to end lock requests.

ESDS Add To End Pct: Percent of write requests that had ESDS add to end lock requests in tenths of a percent.

ESDS Add To End Reqs Minute: Number of ESDS add to end lock requests per minute in hundredths.

ESDS Add To End True Cont Minute: Number of esds add to end lock requests in true contention per minute in hundredths.

ESDS Add To End True Contention: Number of esds add to end lock requests in true contention.

Lock Contention per Minute: Number of lock requests in contention per minute in hundredths.

Lock Contention Percent: Percent of lock requests in contention in .1 percent.

Lock False Contention Pct: Percent of lock requests in false contention in .1 percent.

Lock Req False Content Min: Number of lock requests in false contention per minute in hundredths.

Lock Req True Content Min: Number of lock requests in true contention per minute in hundredths.

Lock Requests: Number of lock requests.

Lock Requests False Contention: Number of lock requests in false contention.

Lock Requests per Minute: Number of lock requests per minute in hundredths.

Lock Requests True Contention: Number of lock requests in true contention.

Lock Structure: Lock structure name.

Lock True Contention Pct: Percent of lock requests in true contention in .1 percent.

ORIGINNODE: An internal node name used to identify the system.

Seq Access Requests Minute: Sequential access requests per minute in hundredths.

Seq Access Requests Total: Total number of sequential access requests.

Sequential Access Reads CR: Number of sequential access reads with CR Integrity.

Sequential Access Reads CR Min: Number of sequential access reads with CR per minute in hundredths.

Sequential Access Reads NRI: Number of sequential access reads with No Read Integrity.

Sequential Access Reads NRI Min: Number of sequential access reads with NRI per minute in hundredths.

Sequential Access Writes: Number of sequential access writes.

Sequential Access Writes Minute: Number of sequential access writes per minute in hundredths.

SMF-ID: SMF id if this is a system level record. Blank if this is for the sysplex.

Statistics Level: Indicator if this record is for the entire sysplex or for an individual system.

Storage Class: Storage Class Name

System Sysplex Name: The system or sysplex name from which the data is derived. If a sysplex name, the data is derived from all systems in the sysplex.

Total Reads: Number of reads.

Total Reads per Minute: Number of reads per minute in hundredths.

Total Requests: Total number of requests.

Total Requests per Minute: Number of requests per minute in hundredths.

Total Writes: Number of write requests.

Total Writes per Minute: Number of write requests per minute in hundredths.

Upgr Lock Reqs True Cont Min: Number of upgrade lock requests in true contention per minute in hundredths.

Upgr Lock Reqs True Contention: Number of upgrade lock requests in true contention.

Upgrade Lock Reqs per Minute: Number of upgrade lock requests per minute in hundredths.

Upgrade Lock Requests: Number of upgrade lock requests.

VSAM Cluster: The name of the cluster for this dataset.

VSAM Component: The component type of the dataset.

RLS Buffer LSU Summary

The RLS Buffer LSU Summary attributes provide the status of the activity for all buffers.

Average CF Cache Read Hit Pct: The average CF Cache read hit percent over the collection interval.

Average DASD Hit Pct: The average DASD hit percent over the collection interval.

Average Local CI Refresh Pct: The average percentage of RE-DO operations over the collection interval.

Avg Recursive Local CI Ref Pct: The average percentage of recursive RE-DO operations across the sysplex.

BMF Accelerated Mode: Total number of buffer manager LRU intervals processed, where BMF was over the goal, accelerated the aging, but did not go into panic mode. This is the LRU intervals in accelerated mode.

BMF Accelerated Mode Pct: Percentage of average intervals where BMF was over the goal, accelerated the aging but did not go into panic mode in tenths of a percent. This is the percent of intervals in accelerated mode.

BMF Panic Mode: Total number of Buffer Manager LRU intervals where BMF was over the goal, and normal algorithms were bypassed to reclaim buffers. This is panic mode.

BMF Panic Mode Pct: The percent of LRU intervals during which BMF was in panic mode in tenths of a percent.

BMF Write Requests: Total write requests through BMF.

Collection Interval: Time span over which data was collected in seconds.

Collection Time: Time when data was collected.

Curr Recursive Local CI Ref Pct: The percentage of recursive RE-DO operations over the current LRU interval.

Current BMF Read Hit Pct: The BMF read hit percent for the current LRU interval.

Current CF Cache Read Hit Pct: The percentage of CF cache read hits over the current LRU interval.

Current DASD Hit Pct: The percentage of DASD hits over the current LRU interval.

Current Local CI Refresh Pct: The percentage of RE-DO operations over the current LRU interval.

Location: Indicates if the statistics are for buffers above or below the bar.

Max BMF Read Hit Pct: The maximum BMF read hit percent for any LRU interval over the collection interval.

Max Recursive Local CI Ref Pct: The maximum percentage of recursive RE-DO operations for any LRU interval over the collection interval.

Maximum CF Cache Read Hit Pct: The maximum CF Cache read hit percent for any LRU interval over the collection interval.

Maximum DASD Hit Pct: The maximum DASD hit percent for any LRU interval over the collection interval.

Maximum Local CI Refresh Pct: The maximum percentage of RE-DO operations for any LRU interval over the collection interval.

Min BMF Read Hit Pct: The minimum BMF read hit percent for any LRU interval over the collection interval.

Min Recursive Local CI Ref Pct: The minimum percentage of recursive RE-DO operations for any LRU interval over the collection interval.

Minimum CF Cache Read Hit Pct: The minimum CF Cache read hit percent for any LRU interval over the collection interval.

Minimum DASD Hit Pct: The minimum DASD hit percent for any LRU interval over the collection interval.

Minimum Local CI Refresh Pct: The minimum percentage of RE-DO operations for any LRU interval over the collection interval.

ORIGINNODE: An internal node name used to identify the system.

Recursive Local CI Refresh: Total number of recursive RE-DO operations.

Recursive Local CI Refresh Pct: The percentage of RE-DO operations that went recursive in tenths of a percent.

SMF-ID: SMF ID.

Statistics Level: Statistics level, indicating if the data is for the sysplex or an individual system.

System Sysplex Name: Either the sysplex or system name.

Total BMF LRU Intervals: Total number of Buffer Manager LRU intervals processed.

Total BMF Read Hits: Number of BMF read hits.

Total BMF Requests: Total number of times that BMF was called.

Total Buffer Size High Goal: Total maximum buffer size goal.

Total Buffer Size High Value: Total high buffer size over the LRU interval.

Total CF Cache Read Hits: Total number of CF cache read hits.

Total DASD Hits: The number of times requests were satisfied from DASD.

Total Refresh Local CI Copy: Total number of redos.RE-DO is an operation to refresh your down level local copy of the CI for write integrity when there are concurrent write requests from multiple threads.

XCF Castout Lock Cont Retry Pct: Percentage of castout contentions retried in tenths of a percent.

XCF Castout Lock Contention: Total number of CF read requests that encountered XCF cast-out lock contention. The cast-out lock is a lock used with the data item (CI) to serialize the update to DASD.

XCF Castout Lock Retries: Total number of retries for castout-lock. Total number of CF read requests that had to be retried more than once because they encountered contention on the cast out lock.

RLS Buffer Pool Detail

The RLS Buffer Pool Detail attributes provide detailed information about buffer manager pool statistics.

Avg Current Buffers: Average current value of the number of BMF buffers for this pool during this interval. This value will be n/a for system level metrics.

Avg Current Extents: Average current value of the number of extents for this pool during this interval. This value will be n/a for system level metrics.

Avg Maximum Buffers: Average high value of the number of BMF buffers for this pool during this interval. This value will be n/a for system level metrics.

Avg Maximum Extents: Average high value of the number of extents for this pool during this interval. This value will be n/a for system level metrics.

Avg Minimum Buffers: Average low value of the number of BMF buffers for this pool during this interval. This value will be n/a for system level metrics.

Avg Minimum Extents: Average low value of the number of extents for this pool during this interval. This value will be n/a for system level metrics.

Buffer Pool Size KB: The size of each buffer in the pool.

Collection Interval: Time span over which data was collected in seconds.

Current Buffers: Current value of the number of BMF buffers for this pool during this interval.

Current Extents: Current value of the number of extents for this pool during this interval.

Location: Indicates whether the statistics are for buffers that are on or off the bar.

Maximum Buffers: Maximum value of the number of BMF buffers for this pool during this interval.

Maximum Extents: Maximum value of the number of extents for this pool during this interval.

Minimum Buffers: Minimum value of the number of BMF buffers for this pool during this interval.

Minimum Extents: Minimum value of the number of extents for this pool during this interval.

ORIGINNODE: An internal node name used to identify the system.

SMF-ID: SMF ID.

Statistics Level: Statistics level, indicating if the data is for the sysplex or an individual system.

System Sysplex Name: Either the sysplex or system name.

RLS Collection Interval

The RLS Collection Interval attributes provide collection control settings for the collection on RLS data for subtypes 15 and 17-19 data.

Managed System: An internal node name used to identify the system.

RLS Collection Interval: The interval. How often the data should be collected in the background.

RLS Collection Type: The type of collection being done.

RLS Dataset Details

The RLS Dataset Details attributes provide performance statistics for ad-hoc RLS dataset requests that are not included in the RLS Application Dataset Performance attribute group.

Average Response Time: Average response time for the dataset in .01 milliseconds.

Avg Response Time Direct Access: Average response time for all direct access requests in .01 milliseconds

Avg Response Time Sequential Access: Average response time for all sequential access requests in .01 milliseconds.

BMF Dir Acc False Invalids: Number of BMF direct access false invalids.

BMF Dir Acc False Invalids per Minute: Number of BMF direct access false invalids per minute in hundredths.

BMF DIR Access Reqs per Minute: Number of BMF direct access requests per minute in hundredths.

BMF Direct Access Reqs: Number of BMF direct access requests.

BMF False Invalid percent: Percent of BMF requests that were false invalids in tenths of a percent.

BMF False Invalids: Number of BMF false invalids.

BMF False Invalids per Minute: Number of BMF false invalids per minute in hundredths.

BMF Requests: Number of BMF requests.

BMF Requests per Minute: Number of BMF requests per minute in hundredths.

BMF Seq Acc False Invalids: Number of BMF sequential access false invalids.

BMF Seq Acc False Invalids per Minute: Number of BMF sequential access false invalids per minute in hundredths.

BMF Seq Access Reqs per Minute: Number of BMF sequential access requests per minute in hundredths.

BMF Sequential Access Reqs: Number of BMF sequential access requests.

CA Splits per Minute: The number of control area splits per minute in hundredths.

CA Splits Percent: The percent of writes that resulted in control area splits.

Cache Set Name: Cache set used by this dataset.

Cache Structure Name: DFSMS Cache Structure name.

CI Splits per Minute: The number of control interval splits per minute in hundredths.

CI Splits Percent: The percent of writes that resulted in a control interval split.

Collection Interval: Time span over which data was collected in seconds.

Collection Time: Time when data was collected.

Control Area Splits: The number of control area splits.

Control Interval Splits: The number of control interval splits.

Dataset Name: Dataset name for which the statistics are provided.

Dir Acc Reads CR per Minute: Number of direct access reads with CR per minute in hundredths.

Dir Acc Reads NRI per Minute: Number of direct access reads with NRI per minute in hundredthss.

Direct Access Reads CR: Number of direct access reads with Consistent Read.

Direct Access Reads NRI: Number of direct access reads with No Read Integrity.

Direct Access Requests Minute: Direct access requests per minute in hundredths.

Direct Access Requests Total: Total number of direct access requests.

Direct Access Writes: Number of direct access writes.

Direct Access Writes Minute: Number of direct access writes per minute in hundredths.

DIWA Lock Percent: Percent of write requests that had Data Insert Work Area (DIWA) lock requests in tenths of a percent.

DIWA Lock Reqs True Contention: Number of Data Insert Work Area (DIWA) lock requests that had true contention.

DIWA Lock Requests: Number of Data Insert Work Area (DIWA) lock requests.

DIWA Lock Requests per Minute: Number of Data Insert Work Area (DIWA) lock requests per minute in hundredths.

DIWA Lock True Contention Minute: Number of Data Insert Work Area (DIWA) lock requests that had true contention per minute in hundredths.

DIWA Lock True Contention Pct: Percent of Data Insert Work Area (DIWA) lock requests that had true contention in tenths of a percent.

ESDS Add To End Lock Requests: Number of ESDS add to end lock requests.

ESDS Add To End Pct: Percent of write requests that had ESDS add to end lock requests in tenths of a percent.

ESDS Add To End Reqs Minute: Number of ESDS add to end lock requests per minute in hundredths.

ESDS Add To End True Cont Minute: Number of ESDS add to end lock requests in true contention per minute in hundredths.

ESDS Add To End True Contention: Number of ESDS add to end lock requests in true contention.

Lock Contention per Minute: Number of lock requests in contention per minute in hundredths.

Lock Contention Percent: Percent of lock requests in contention in .1 percent.

Lock False Contention Pct: Percent of lock requests in false contention in .1 percent.

Lock Req False Content per Min: Number of lock requests in false contention per minute in hundredths.

Lock Req True Content per Min: Number of lock requests in true contention per minute in hundredths.

Lock Requests: Number of lock requests.

Lock Requests False Contention: Number of lock requests in false contention.

Lock Requests per Minute: Number of lock requests per minute in hundredths.

Lock Requests True Contention: Number of lock requests in true contention.

Lock Structure: Lock structure name.

Lock True Contention Pct: Percent of lock requests in true contention in .1 percent.

ORIGINNODE: An internal node name used to identify the system.

Seq Acc Reads CR per Minute: Number of sequential access reads with CR per minute in hundredths.

Seq Acc Reads NRI per Minute: Number of sequential access reads with NRI per minute in hundredths.

Seq Access Requests Minute: Sequential access requests per minute in hundredths.

Seq Access Requests Total: Total number of sequential access requests.

Sequential Access Reads CR: Number of sequential access reads with CR Integrity.

Sequential Access Reads NRI: Number of sequential access reads with No Read Integrity.

Sequential Access Writes: Number of sequential access writes.

Sequential Access Writes Minute: Number of sequential access writes per minute in hundredths.

SMF-ID: SMF id if this is a system level record. Blank if this is for the sysplex.

Statistics Level: Indicator if this record is for the entire sysplex or for an individual system.

Storage Class: Storage Class Name

System Sysplex Name: The system or sysplex name from which the data is derived. If a sysplex name, the data is derived from all systems in the sysplex.

Total Reads: Number of reads.

Total Reads per Minute: Number of reads per minute in hundredths.

Total Requests: Total number of requests.

Total Requests per Minute: Number of requests per minute in hundredths.

Total Writes: Number of write requests.

Total Writes per Minute: Number of write requests per minute in hundredths.

Upgrade Lock Reqs per Minute: Number of upgrade lock requests per minute in hundredths.

Upgrade Lock Reqs True Cont Minute: Number of upgrade lock requests in true contention per minute in hundredths.

Upgrade Lock Reqs True Contention: Number of upgrade lock requests in true contention.

Upgrade Lock Requests: Number of upgrade lock requests.

VSAM Cluster: The name of the cluster for this dataset.

VSAM Component: The component type of the dataset.

Dataset Collection Interval

The Dataset Collection Interval attributes provide information about the dataset group.

Dataset Group Collection Interval: This is the Collection Interval for DSN Group processing.

Dataset Group Collection Type: This is the Collection type for DSN Group processing.The type can be:

MIN

minutes

RMF

RMF interval

ODN

on-demand

Managed System: An internal node name used to identify the system from which this row of data was obtained.

RLS Dataset Group Details

The RLS Dataset Group Details attributes provide performance statistics for the datasets in a Dataset Group.

Average Response Time: Average response time for the dataset in .01 milliseconds.

Avg Response Time Direct Access: Average response time for all direct access requests in .01 milliseconds

Avg Response Time Sequential Access: Average response time for all sequential access requests in .01 milliseconds.

BMF Dir Acc False Invalids: Number of BMF direct access false invalids.

BMF Dir Acc False Invalids per Minute: Number of BMF direct access false invalids per minute in hundredths.

BMF DIR Access Reqs per Minute: Number of BMF direct access requests per minute in hundredths.

BMF Direct Access Reqs: Number of BMF direct access requests.

BMF False Invalid percent: Percent of BMF requests that were false invalids in tenths of a percent.

BMF False Invalids: Number of BMF false invalids.

BMF False Invalids per Minute: Number of BMF false invalids per minute in hundredths.

BMF Requests: Number of BMF requests.

BMF Requests per Minute: Number of BMF requests per minute in hundredths.

BMF Seq Acc False Invalids: Number of BMF sequential access false invalids.

BMF Seq Acc False Invalids per Minute: Number of BMF sequential access false invalids per minute in hundredths.

BMF Seq Access Reqs per Minute: Number of BMF sequential access requests per minute in hundredths.

BMF Sequential Access Reqs: Number of BMF sequential access requests.

CA Splits per Minute: The number of control area splits per minute in hundredths.

CA Splits Percent: The percent of writes that resulted in control area splits.

Cache Set Name: Cache set used by this dataset.

CI Splits per Minute: The number of control interval splits per minute in hundredths.

CI Splits Percent: The percent of writes that resulted in a control interval split.

Collection Interval: Time span over which data was collected in seconds.

Collection Time: Time when data was collected.

Control Area Splits: The number of control area splits.

Control Interval Splits: The number of control area splits.

Dataset Group Name: Name of the dataset group.

Dataset Name: Dataset name for which the statistics are provided.

Dir Acc Reads CR per Minute: Number of direct access reads with CR per minute in hundredths.

Dir Acc Reads NRI per Minute: Number of direct access reads with NRI per minute in hundredths.

Direct Access Reads CR: Number of direct access reads with Consistent Read.

Direct Access Reads NRI: Number of direct access reads with No Read Integrity.

Direct Access Requests Minute: Direct access requests per minute in hundredths.

Direct Access Requests Total: Total number of direct access requests.

Direct Access Writes: Number of direct access writes.

Direct Access Writes Minute: Number of direct access writes per minute in hundredths.

DIWA Lock Percent: Percent of write requests that had Data Insert Work Area (DIWA) lock requests in tenths of a percent.

DIWA Lock Reqs True Contention: Number of Data Insert Work Area (DIWA) lock requests that had true contention.

DIWA Lock Requests: Number of Data Insert Work Area (DIWA) lock requests.

DIWA Lock Requests per Minute: Number of Data Insert Work Area (DIWA) lock requests per minute in hundredths.

DIWA Lock True Contention Minute: Number of Data Insert Work Area (DIWA) lock requests that had true contention per minute in hundredths.

DIWA Lock True Contention Pct: Percent of Data Insert Work Area (DIWA) lock requests that had true contention in tenths of a percent.

ESDS Add To End Lock Requests: Number of ESDS add to end lock requests.

ESDS Add To End Pct: Percent of write requests that had ESDS add to end lock requests in tenths of a percent.

ESDS Add To End Reqs Minute: Number of ESDS add to end lock requests per minute in hundredths.

ESDS Add To End True Cont Minute: Number of esds add to end lock requests in true contention per minute in hundredths.

ESDS Add To End True Contention: Number of esds add to end lock requests in true contention.

Lock Contention per Minute: Number of lock requests in contention per minute in hundredths.

Lock Contention Percent: Percent of lock requests in contention in .1 percent.

Lock False Contention Pct: Percent of lock requests in false contention in .1 percent.

Lock Req False Content per Min: Number of lock requests in false contention per minute in hundredths.

Lock Req True Content per Min: Number of lock requests in true contention per minute in hundredths.

Lock Requests: Number of lock requests.

Lock Requests False Contention: Number of lock requests in false contention.

Lock Requests per Minute: Number of lock requests per minute in hundredths.

Lock Requests True Contention: Number of lock requests in true contention.

Lock Structure: Lock structure name.

Lock True Contention Pct: Percent of lock requests in true contention in .1 percent.

ORIGINNODE: An internal node name used to identify the system.

Seq Acc Reads CR per Minute: Number of sequential access reads with CR per minute in hundredths.

Seq Acc Reads NRI per Minute: Number of sequential access reads with NRI per minute in hundredths.

Seq Access Requests Minute: Sequential access requests per minute in hundredths.

Seq Access Requests Total: Total number of sequential access requests.

Sequential Access Reads CR: Number of sequential access reads with CR Integrity.

Sequential Access Reads NRI: Number of sequential access reads with No Read Integrity.

Sequential Access Writes: Number of sequential access writes.

Sequential Access Writes Minute: Number of sequential access writes per minute in hundredths.

SMF-ID: SMF id if this is a system level record. Blank if this is for the sysplex.

Statistics Level: Indicator if this record is for the entire sysplex or for an individual system.

Storage Class: Storage Class Name

System Sysplex Name: The system or sysplex name from which the data is derived. If a sysplex name, the data is derived from all systems in the sysplex.

Total Reads: Number of reads.

Total Reads per Minute: Number of reads per minute in hundredths.

Total Requests: Total number of requests.

Total Requests per Minute: Number of requests per minute in hundredths.

Total Writes: Number of write requests.

Total Writes per Minute: Number of write requests per minute in hundredths.

Upgrade Lock Reqs per Minute: Number of upgrade lock requests per minute in hundredths.

Upgrade Lock Reqs True Cont Minute: Number of upgrade lock requests in true contention per minute in hundredths.

Upgrade Lock Reqs True Contention: Number of upgrade lock requests in true contention.

Upgrade Lock Requests: Number of upgrade lock requests.

VSAM Cluster: The name of the cluster for this dataset.

VSAM Component: The component type of the dataset.

RLS Dataset Group Summary

The RLS Dataset Group Summary attributes provide a roll-up of the RLS statistics for the datasets in the dataset group.

ORIGINNODE: Originnode.

Dataset Group Name: Name of the dataset group.

Max Average Response Time: The maximum response time for any dataset in the group in .01 milliseconds.

Max Avg Resp Time Direct Access: The maximum average response time for all direct access requests in .01 milliseconds for any dataset in the dataset group across the sysplex.

Max Avg Response Time Seq Access: The maximum average response time for all sequential access requests in .01 milliseconds for any dataset in the dataset group across the sysplex.

Max BMF False Invalid percent: The maximum percent of BMF false invalid requests for any dataset in the dataset group across the sysplex in tenths of a percent.

Max BMF False Invalids: The maximum number of BMF false invalid requests for any dataset in the dataset group across the sysplex.

Max BMF False Invalids per Min: The maximum number of BMF false invalid requests per minute in hundredths for any dataset in the dataset group across the sysplex.

Max BMF Requests: The maximum number of BMF requests for any dataset in the dataset group across the sysplex.

Max BMF Requests per Minute: The maximum number of BMF requests per minute in hundredths for any dataset in the dataset group across the sysplex.

Max CA Splits: The maximum number of control area splits for any dataset in the group across the sysplex.

Max CA Splits pct: The maximum percent of write requests that caused a control area split for any dataset in the dataset group across the sysplex in tenths of a percent.

Max CA Splits per Minute: The maximum number of control area splits per minute in hundredths for any dataset in the dataset group across the sysplex.

Max CI Splits: The maximum number of control interval splits for any dataset in the group across the sysplex.

Max CI Splits pct: The maximum percent of write requests that caused a control interval split for any dataset in the dataset group across the sysplex in tenths of a percent.

Max CI Splits per Minute: The maximum number of control interval splits per minute in hundredths for any dataset in the dataset group across the sysplex.

Max DIWA Lock Reqs per Minute: The maximum number of Data Insert Work Area (DIWA) lock requests per minute in hundredths for any dataset in the dataset group across the sysplex.

Max DIWA Lock Request Pct: The maximum percent of write requests that had Data Insert Work Area (DIWA) lock requests for any dataset in the dataset group across the sysplex in tenths of a percent.

Max DIWA Lock Requests: The maximum number of Data Insert Work Area (DIWA) lock requests for any dataset in the dataset group across the sysplex.

Max ESDS Add To End Pct: The maximum percent of write requests that had ESDS add to end lock requests for any dataset in the dataset group across the sysplex in tenths of a percent.

Max Lock Contention per Minute: The maximum number of lock requests in contention per minute in hundredths for any dataset in the dataset group across the sysplex.

Max Lock Contention Percent: The maximum percent of lock requests in contention for any dataset in the dataset group across the sysplex in tenths of a percent.

Max Total Reads: The maximum number of read requests for any dataset in the dataset group across the sysplex.

Max Total Reads per Minute: The maximum number of read requests per minute in hundredths for any dataset in the dataset group across the sysplex.

Max Total Reqs per Minute: The maximum number of requests per minute in hundredths for any dataset in the dataset group across the sysplex.

Max Total Requests: The maximum total requests for any dataset in the dataset group across the sysplex.

Max Total Writes: The maximum number of write requests for any dataset in the dataset group across the sysplex.

Max Total Writes per Minute: The maximum number of write requests per minute in hundredths for any dataset in the dataset group across the sysplex.

Sysplex Name: The name of the sysplex from which the data was derived.

RLS Lock Structure

The RLS Lock Structure attributes provide a roll-up of the activity for all lock structures.

Add To End Lock Reqs False Cont: Total number of ESDS Add to End lock requests in false contention.

Add To End Lock Reqs Per Min: Number of ESDS Add to End lock requests per minute in hundredths.

Add To End Lock Reqs True Cont: Total number of ESDS Add to End lock requests in true contention.

ATE Lock False Cont Pct: Percent of ESDS Add to End lock requests in false contention in tenths of a percent.

ATE Lock False Cont Per Min: Number of ESDS Add to End lock requests in false contention per minute in hundredths.

ATE Lock True Cont Pct: Percent of ESDS Add to End lock requests in true contention in tenths of a percent.

ATE Lock True Cont Per Min: Number of ESDS Add to End lock requests in true contention per minute in hundredths.

Average Response Time: Average response time for the dataset in .01 milliseconds.

Collection Interval: Time span over which data was collected in seconds.

Collection Time: Time when data was collected.

DIWA Lock False Cont Minute: Number of Data Insert Work Area (DIWA) lock requests that had false contention per minute in hundredths.

DIWA Lock False Contention: Number of Data Insert Work Area (DIWA) lock requests that had false contention.

DIWA Lock False Contention Pct: Percent of Data Insert Work Area (DIWA) lock requests that had false contention in tenths of a percent.

DIWA Lock Reqs True Contention: Number of Data Insert Work Area (DIWA) lock requests that had true contention.

DIWA Lock Requests: Number of Data Insert Work Area (DIWA) lock requests.

DIWA Lock Requests per Minute: Number of Data Insert Work Area (DIWA) lock requests per minute in hundredths.

DIWA Lock True Contention Minute: Number of Data Insert Work Area (DIWA) lock requests that had true contention per minute in hundredths.

DIWA Lock True Contention Pct: Percent of Data Insert Work Area (DIWA) lock requests that had true contention in tenths of a percent.

ESDS Add To End Lock Requests: Total number of component 2 (ESDS Add to End) lock requests.

Exception Path Lock Req Minute: VSAM RLS exception path lock requests per minute in hundredths.

Exception Path Lock Requests: Total number of lock requests. VSAM RLS exception path.

Excp Path False Cont Lock Req: Total number of false contention VSAM RLS exception path lock requests.

Excp Path False Cont Lock Req Per Minute: Number of false contention VSAM RLS exception path lock requests per minute in hundredths.

Excp Path False Contention Pct: Percent of exception path lock requests that were in false contention in tenths of a percent.

Excp Path True Cont Lock Req: Total number of true contention VSAM RLS exception path lock requests.

Excp Path True Cont Lock Req Per Minute: Number of true contention VSAM RLS exception path lock requests per minute in hundredths.

Excp Path True Contention Pct: Percent of exception path lock requests that were in true contention in tenths of a percent.

Lock Table Name: DFSMS lock table name.

Main Path False Cont Lock Req: Total number of false contention VSAM RLS mainline lock requests.

Main Path False Cont Lock Req Per Minute: Number of false contention VSAM RLS mainline lock requests per minute in hundredths.

Main Path False Contention Pct: Percent of mainline lock requests that were in false contention in tenths of a percent.

Main Path True Cont Lock Req: Total number of true contention VSAM RLS mainline lock requests.

Main Path True Cont Lock Req Per Minute: Number of true contention VSAM RLS mainline lock requests per minute in hundredths.

Main Path True Contention Pct: Percent of mainline lock requests that were in true contention in tenths of a percent.

Mainline Path Lock Req Minute: VSAM RLS mainline path lock requests per minute in hundredths.

Mainline Path Lock Requests: Total number of lock requests. VSAM RLS mainline path.

ORIGINNODE: An internal node name used to identify the system.

Rec Lock False Contention Pct: Percent of record lock requests that were in false contention in tenths of a percent.

Rec Lock Req False Con Minute: Number of record lock requests in false contention per minute in hundredths.

Rec Lock Req False Contention: Number of record lock requests (obtain, alter, promote) in false contention.

Rec Lock Req True Con Minute: Number of record lock requests in true contention per minute in hundredths.

Rec Lock Req True Contention: Number of record lock requests (obtain, alter, promote) in true contention.

Rec Lock True Contention Pct: Percent of record lock requests that were in true contention in tenths of a percent.

Record Lock Reqs Per Minute: Number of record lock requests per minute in hundredths.

Record Lock Requests: Number of record lock requests (obtain, alter, promote)

SMF-ID: SMF ID.

Statistics Level: Statistics level, indicating if the data is for the sysplex or an individual system.

System Sysplex Name: Either the sysplex or system name.

Upg Locks False Cont Per Min: Number of upgrade locks in false contention per minute in hundredths.

Upg Locks True Cont Per Min: Number of upgrade locks in true contention per minute in hundredths.

Upgrade Lock Requests: Number of upgrade lock requests (component 1 class 2).

Upgrade Locks False Cont Pct: Percentage of upgrade lock requests that had false contention in tenths of a percent.

Upgrade Locks False Contention: Number of upgrade locks in false contention.

Upgrade Locks Requests Per Minute: Number of upgrade lock requests per minute in hundredths.

Upgrade Locks True Cont Pct: Percentage of upgrade lock requests that had true contention in tenths of a percent.

Upgrade Locks True Contention: Number of upgrade locks in true contention.

RLS Lock Structure CF Detail

The RLS Lock Structure CF Detail attributes return information about the lock structure.

Duplex Time: Duplex elapsed time in seconds.

Entries Used Pct: Percent of lock structure entries in use in tenths of a percent.

Lock Mode: Lock structure mode. Valid values are duplex, duplexing, and simplex.

Lock Name: The name of the lock structure.

Lock Status: Lock status. Valid values are Not-Inited, Enabling, Enable, Qui-ing, Qui-ed.

ORIGINNODE: An internal node name used to identify the system.

Primary Version: Primary lock structure version. Display as hex character string.

Secondary Version: Secondary lock structure version. Only valid when mode is duplex or duplexing. Display as hex character string.

Structure Size: Size of the lock structure in kb.

Sysplex Name: The sysplex name from which the data is derived.

Total Entries: Total number of entries in the lock structure.

Used Entries: Number of entries in use in the lock structure.

RLS Performance Overview

The RLS Performance Overview attributes provide performance data for each interval for each system in the sysplex with activity and one row of data is given for each interval across the sysplex.

Avg Queue Length|: Average number of requests waiting for locks in thousandths.

Contention Rate: Percentage of lock requests that were globally managed in thousandths of a percent.

False Cont Rate: Percentage of lock requests that were falsely globally managed in thousandths of a percent.

Interval: Time interval for which data was captured. Either 1 minute, 1 hour, 8 hours, or 1 day.

Lock Name: The name of the lock structure.

Lock Rate: Number of lock requests per second in thousandths

ORIGINNODE: An internal node name used to identify the system.

SMF-ID: SMF ID.

Statistics Level: Statistics level, indicating if the data is for the sysplex or an individual system.

System Sysplex Name: Either the sysplex or system name.

Systems Reporting: The number of systems included in the sysplex statistics. This will be 1 for a single system report.

RLS Storage Class

The RLS storage class attributes provide performance statistics for each Storage Class with RLS activity in the sysplex.

Average Response Time: Average response time for the Storage class in .01 milliseconds.

Avg Response Time Direct Access: Average response time for all direct access requests in .01 milliseconds

Avg Response Time Sequential Access: Average response time for all sequential access requests in .01 milliseconds.

BMF Dir Acc False Invalids: Number of BMF direct access false invalids.

BMF Dir Acc False Invalids per Minute: Number of BMF direct access false invalids per minute in hundredths.

BMF DIR Access Reqs per Minute: Number of BMF direct access requests per minute in hundredths.

BMF Direct Access Reqs: Number of BMF direct access requests.

BMF False Invalid percent: Percent of BMF requests that were false invalids in tenths of a percent.

BMF False Invalids: Number of BMF false invalids.

BMF False Invalids per Minute: Number of BMF false invalids per minute in hundredths.

BMF Requests: Number of BMF requests.

BMF Requests per Minute: Number of BMF requests per minute in hundredths.

BMF Seq Acc False Invalids: Number of BMF sequential access false invalids.

BMF Seq Acc False Invalids per Minute: Number of BMF sequential access false invalids per minute in hundredths.

BMF Seq Access Reqs per Minute: Number of BMF sequential access requests per minute in hundredths.

BMF Sequential Access Reqs: Number of BMF sequential access requests.

CA Splits per Minute: The number of control area splits per minute in hundredths for all datasets in the storage class.

CA Splits Percent: The percent of writes that resulted in control area splits for all datasets in the storage class.

Cache Set Name: Cache set used by this dataset.

CI Splits per Minute: The number of control interval splits per minute in hundredths for all datasets in the storage class.

CI Splits Percent: The percent of writes that resulted in a control interval split for all datasets in the storage class.

Collection Interval: Time span over which data was collected in seconds.

Collection Time: Time when data was collected.

Control Area Splits: The number of control area splits for all datasets in the storage class.

Control Interval Splits: The number of control area splits for all datasets in the storage class.

Dir Acc Reads CR per Minute: Number of direct access reads with CR per minute in hundredths.

Dir Acc Reads NRI per Minute: Number of direct access reads with NRI per minute in hundredths.

Direct Access Reads CR: Number of direct access reads with Consistent Read.

Direct Access Reads NRI: Number of direct access reads with No Read Integrity.

Direct Access Requests Minute: Direct access requests per minute in hundredths.

Direct Access Requests Total: Total number of direct access requests.

Direct Access Writes: Number of direct access writes.

Direct Access Writes Minute: Number of direct access writes per minute in hundredths.

DIWA Lock Percent: Percent of write requests that had Data Insert Work Area (DIWA) lock requests in tenths of a percent.

DIWA Lock Reqs True Contention: Number of Data Insert Work Area (DIWA) lock requests that had true contention.

DIWA Lock Requests: Number of Data Insert Work Area (DIWA) lock requests.

DIWA Lock Requests per Minute: Number of Data Insert Work Area (DIWA) lock requests per minute in hundredths.

DIWA Lock True Contention Minute: Number of Data Insert Work Area (DIWA) lock requests that had true contention per minute in hundredths.

DIWA Lock True Contention Pct: Percent of Data Insert Work Area (DIWA) lock requests that had true contention in tenths of a percent.

ESDS Add To End Lock Requests: Number of ESDS add to end lock requests.

ESDS Add To End Pct: Percent of write requests that had ESDS add to end lock requests in tenths of a percent.

ESDS Add To End Reqs Minute: Number of ESDS add to end lock requests per minute in hundredths.

ESDS Add To End True Cont Minute: Number of esds add to end lock requests in true contention per minute in hundredths.

ESDS Add To End True Contention: Number of esds add to end lock requests in true contention.

Lock Contention per Minute: Number of lock requests in contention per minute in hundredths.

Lock Contention Percent: Percent of lock requests in contention in .1 percent.

Lock False Contention Pct: Percent of lock requests in false contention in .1 percent.

Lock Req False Content per Min: Number of lock requests in false contention per minute in hundredths.

Lock Req True Content per Min: Number of lock requests in true contention per minute in hundredths.

Lock Requests: Number of lock requests.

Lock Requests False Contention: Number of lock requests in false contention.

Lock Requests per Minute: Number of lock requests per minute in hundredths.

Lock Requests True Contention: Number of lock requests in true contention.

Lock Structure: Lock structure name.

Lock True Contention Pct: Percent of lock requests in true contention in .1 percent.

ORIGINNODE: An internal node name used to identify the system.

Seq Acc Reads CR per Minute: Number of sequential access reads with CR per minute in hundredths.

Seq Acc Reads NRI per Minute: Number of sequential access reads with NRI per minute in hundredths.

Seq Access Requests Minute: Sequential access requests per minute in hundredths.

Seq Access Requests Total: Total number of sequential access requests.

Sequential Access Reads CR: Number of sequential access reads with CR Integrity.

Sequential Access Reads NRI: Number of sequential access reads with No Read Integrity.

Sequential Access Writes: Number of sequential access writes.

Sequential Access Writes Minute: Number of sequential access writes per minute in hundredths.

SMF-ID: SMF id if this is a system level record. Blank if this is for the sysplex.

Statistics Level: Indicator if this record is for the entire sysplex or for an individual system.

Storage Class: Storage Class Name

System Sysplex Name: The system or sysplex name from which the data is derived. If a sysplex name, the data is derived from all systems in the sysplex.

Total Reads: Number of reads.

Total Reads per Minute: Number of reads per minute in hundredths.

Total Requests: Total number of requests.

Total Requests per Minute: Number of requests per minute in hundredths.

Total Writes: Number of write requests.

Total Writes per Minute: Number of write requests per minute in hundredths.

Upgrade Lock Reqs per Minute: Number of upgrade lock requests per minute in hundredths.

Upgrade Lock Reqs True Cont Minute: Number of upgrade lock requests in true contention per minute in hundredths.

Upgrade Lock Reqs True Contention: Number of upgrade lock requests in true contention.

Upgrade Lock Requests: Number of upgrade lock requests.

RMM Configuration

The RMM Configuration attributes provide RMM configuration state and option data.

Accounting Source: Accounting is from the STEP or the JOB. (Job=0; Step=1)

Backup In Progress: Indicates whether backup is in progress. (No=0; Yes=1)

Backup Procedure: Specifies the name of the backup procedure.

BLP: Indicates whether RMM mode is active for Bypass Label Processing (BLP) mount requests. (RMM=0; NoRMM=1)

Cat Sync In Progress: Indicates whether catalog synchronization is in progress. (No=0; Yes=1)

Catalog Retention Period: Indicates the number of hours assigned as the catalog retention period.

Catalog SysID: Indicates whether the catalog SysID is set. (Set=0; NotSet=1)

CDS ID: Identifies the ID of the control dataset.

Common Time: Indicates whether the Common Time feature is enabled. (Disabled=0; Enabled=1)

Date Format: Specifies the date format. (American=1; European=2; ISO=3; Julian=4)

Days Since Last Low on Scratch Processing: Specifies the days since the last low on the processing of scratch space.

Days Since Last Expiration Processing: Specifies the days since last expiration processing event.

Days Since Last Storage Location Proc: Specifies the days since last storage location processing event.

Days Since Last VRS Update: Specifies the days since the last VRS update.

Days Since CDS Extract: Specifies the days since a CDS extract.

Days Since Last Cat Sync: Specifies the days since last catalog synchronization.

Days Since Last Update: Specifies the days since the last update.

Default Retention Period: Specifies the number of days in the default retention period.

Distant Bins Free: Specifies the number of distant bins that are free.

Distant Bins Total: Specifies the total number of distant bins.

DSN Cmd Auth: Indicates whether DSN command authorization is enabled. (No=0; Yes=1)

EDGUX100 Status: Indicates the enablement status of the EDGUX100 installation exit. (None=0; Enabled=1; Disabled=2)

EDGUX200 Status: Indicates the enablement status of the EDGUX200 installation exit. (None=0; Enabled=1; Disabled=2)

Expiration In Progress: Indicates whether expiration is in progress. (No=0; Yes=1)

Extended Bin: Indicates the enablement status of extended bin support. (Disabled=0; Enabled=1)

Extract In Progress: Indicates whether an extract operation is in progress. (No=0; Yes=1)

IPL Date Check: Indicates the enablement status of the IPL date check. (No=0; Yes=1)

Journal Full Value: Specifies the value that defines the journal full condition.

Journal Status: Indicates the enablement status of the journal. (Disabled=0; Enabled=1; Locked=2)

Last Catalog Sync: Date and time of the last catalog synchronization.

Last CDS Extract: Specifies the last extraction performed on the CDS.

Last Expiration Processing: Days since last VRS update.

Last Low on Scratch Processing: Last low on the processing of scratch space.

Last Storage Location Processing: Last storage location processing.

Last Update Date: Specifies the date of the last update.

Last VRS Update: Specifies the last update to the vital record specification.

Library Racks Free: Specifies the number of free Library racks.

Library Racks Total: Specifies the total number of library racks.

Local Bins Free: Specifies the number of free local bins.

Local Bins Total: Specifies the total number of local bins.

Local Tasks: Specifies the number of local tasks.

Location In Progress: Indicates whether the Location In Progress option is enabled. (No=0; Yes=1)

Managed System: Specifies the name of the managed system.

Master Overwrite: Specifies the Master Overwrite setting that is in effect. (Add=1; Last=2; Match=3; User=4)

Max Hold: Specifies the Max Hold value that is in effect.

Maximum Retention Period: Maximum retention period in days. (Unlimited=-1)

Media Name: Specifies the name assigned to the media.

Move By: Move VRS volumes individual or by multi-volume set. (Volume=0; Set=1)

Operating Mode: Specifies the operating mode. (Manual=1; Record=2; Warning=3; Protect=4)

Owner Cmd Auth: Indicates whether owner command authorization is enabled. (No=0; Yes=1)

Parmlib Suffix: Specifies the suffix of the parameter library.

Pre ACS: Specifies whether Pre-ACS is enabled. (No=0; Yes=1)

RACF®: Specifies the RACF® option that is in use. (None=1; Predefined=2; Automatic=3; Cleanup=4)

Remote Bins Free: Specifies the number of free remote bins.

Remote Bins Total: Specifies the total number of remote bins.

Restore In Progress: Indicates whether a restore is in progress. (No=0; Yes=1)

Retain By: Specifies the retention policy for volumes. (Volume=0; Set=1)

Reuse Bin: Specifies the type of reuse allowed for a bin. (ConfirmMove=0; StartMove=1)

RMM System ID: Specifies the RMM system ID.

Scratch Procedure: Specifies the name of the scratch procedure.

SMF Audits Rec Num: Specifies the SMF audits rec num.

SMF Security Rec Num: Specifies the SMF security rec num.

SMSACS: Indicates whether the SMSACS option is enabled. (No=0; Yes=1)

SMSTape Command: Indicates whether the Command subparameter is enabled for the SMSTape operand. (No=0; Yes=1)

SMSTape Exits: Indicates whether the Exits subparameter is enabled for the SMSTape operand. (No=0; Yes=1)

SMSTape Purge: Indicates whether the Purge parameter is enabled for the SMSTape operand. (No=0; Yes=1; Asis=2)

SMSTape Scratch: Indicates whether the Scratch subparameter is enabled for the SMSTape operand. (No=0; Yes=1)

Stacked Volume Enable Status: Indicates the enablement status for stacked volumes. (None=0; Enabled=1; Disabled=2; Mixed=3)

Subsystem ID : Specifies the RMM subsystem ID.

Subsystem Status: Indicates the status of the RMM subsystem. (Inactive=0; Active=1)

Subsystem Type: Identifies the subsystem type. (Standard=0; Client=1; Server=2)

Tapevol Exit Purge: Specifies the Tapevol Exit Purge setting. (Release=0; Expire=1; None=2)

Timestamp: Provides a timestamp.

UnCatalog: Indicates whether the UnCatalog specifier is enabled. (No=0; Yes=1; Scratch=2)

User Notification: Indicates whether user notification is enabled. (No=0; Yes=1)

Verify In Progress: Indicates that verification is in progress. (No=0; Yes=1)

VRS Change: Indicates the response to changes to the vital record specification. (Info=0; Verify=1)

VRS Name Precedence: Indicates name precedence for the vital record specification. (Job=1; Dataset=2)

VRS In Progress: Indicates whether vital record specification is in progress. (No=0; Yes=1)

VRSEL: Indicates the type of vital record processing. (Old=0; New=1)

VRSMIN Count: Specifies the minimum number of vital record specifications.

VRSMIN Action: Specifies the action that occurs when the minimum number of vital record specifications is reached. (Fail=0; Info=1; Warn=2)

RMM Control

The RMM Control attributes provide basic RMM Control data.

Message: Information regarding the data that was obtained.

Controls: DFSMSrmm controls.

Timestamp: Time table row was created in CT timestamp format.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

RMM Control Dataset

The RMM Control Dataset attributes provide RMM Control Dataset and Journal attributes and statistics.

Actual Percent Used : The percent of the dataset that is used with reference to tracks. For control datasets, this percentage value is calculated by comparing the values of High Allocated and High Used RBA. For Journals, this percentage value is calculated by comparing Allocated Tracks and Last Track Used.

CA Splits: Number of component CA splits.

Catalog Entry Type: Catalog Entry Type for the dataset. (Cluster=C Data_Component=D Index_Component=I NonVSAM=A)

CI Splits: Number of component CI Splits.

Creation Date: Dataset Creation Date.

Dataset Type: Type of CDS Dataset. (Master_CDS=0; Satellite_CDS=1; Journal=9)

Dataset Name: Name of the dataset.

Days Since Creation: Number of days since the creation of the dataset.

Days Since Last Backup: Number of days since the last dataset backup.

Extents: Number of extents allocated to the dataset.

GB Allocated: Number of GB allocated.

GB Free: Number of GB free.

GB Used: Number of GB used.

Last Backup: Date of the last backup as reported by RMM (for the CDS) or as recorded in the catalog entry (for the Journal).

Managed System: Managed System.

MB Allocated: Number of MB allocated.

MB Free: Number of MB free.

MB Used: Number of MB used.

RMM Percent Used: Dataset percent used as reported by RMM.

Timestamp: Timestamp.

Tracks Allocated: Number of Tracks allocated to the dataset.

Tracks Free: Number of tracks free.

Tracks Used: Number of tracks used by the dataset.

Volser: Volume Serial on which dataset is located.

RMM Dataset Details

The RMM Dataset Details attributes provide data for a single RMM dataset.

Abend While Open: Dataset was open at time of a system or application abend. (No=0; Yes=1)

Block Count: Number of blocks used by the dataset on this volume.

Blocksize: The Maximum Dataset blocksize. Blocksize*Number of Blocks = Approximate dataset size.

Cataloged: Dataset is cataloged. (Unknown=0; No=1 ;Yes=2)

Creating DDname: DDname that created the dataset.

Creating Device: Name of device on which the dataset was created.

Creating Job: Name of job that created the dataset.

Creating Program: Name of program that created the dataset.

Creating Step: Name of step that created the dataset.

Creation Date: Date on which the dataset was created.

Data Class: DFSMS Data Class for the dataset.

Dataset Name: Name of the dataset.

Dataset Sequence: Relative position of the dataset within a multiple dataset group.

Date Last Read: Date when the dataset was last accessed for Input.

Date Last Written: Date when the dataset was last accessed for Output.

Days for Retention: Number of days until or since (negative number) the dataset is no longer retained by the current VRS. Nonexpiring=-2147483648; WHILECATLG=-286331154

Days Since Creation: Number of days since the dataset was created.

Days Since Last Read: Number of days since the dataset was last accessed for input.

Days Since Last Write: Number of days since the dataset was last accessed for output.

Days Since Original Expiration: Days Since/Until the Original Expiration date. Nonexpiring=-2147483648

Days Until Expiration: Number of days until or since (negative number) the dataset expires. Nonexpiring=-2147483648

Expiration Date: Date on which the dataset expires.

Last DDname: DDname that last accessed the dataset.

Last Device: Number of the device that last accessed the dataset.

Last Job: Name of the Job that last accessed the dataset.

Last Program: Name of the Program that last accessed the dataset.

Last Step: Name of the Step that last accessed the dataset.

LRECL: Logical Record Length (LRECL).

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Management Class: DFSMS Management Class for the dataset.

Original Expiration Date: The Expiration Date contained in the label when the dataset was created.

Owner: ID of the owner of the volume where the dataset is located.

Percent of Volume: How much of the physical tape media is occupied by the dataset. **Note:** The accuracy of this number is based on information reported to the host at the end of a tape volume file. The number can only be provided if the hardware supports the provision of positioning information, such as with IBM® 3590 Magstar® devices.

Physical Sequence: Position of the dataset on the volume relative to other datasets.

Primary VRS Date: Primary VRS Date.

Primary VRS Jobname: Job name mask of the primary VRS for this dataset.

Primary VRS Name: Dataset name mask of the primary VRS for this dataset.

Primary VRS Subchain Name: Name of subchain VRS.

Primary VRS Type: Type of the primary VRS selected for the dataset.

- Undefined=0 Dataset=1 SMSMC=2 VRSMV=3 DSNMV=4 DSNMC=5

Record Format: Record Format (RECFM).

Retention Date: Date after which the dataset is no longer retained by the current VRS.

Secondary VRS Date: Date on which the secondary VRS became active for the dataset.

Secondary VRS Jobname: Job name mask of the secondary VRS for this dataset.

Secondary VRS Mask: Mask that matches the VRS management value of SMS management class.

Secondary VRS Subchain Name: Name of the subchain VRS.

Security Class: Security Class name.

Storage Class: DFSMS Storage Class for the dataset.

Storage Group: DFSMS Storage Group for the dataset.

System ID: SMF ID of the system on which the dataset was created.

Timestamp: Time table row was created in CT® timestamp format.

Total Blocks: Number of blocks used by the dataset on all volumes. This value includes all volumes up to and including the current volume.

Volser: Serial number of the volume on which the dataset is located.

VRS Retained: Dataset is retained by a Vital Record Specification. (No=0; Yes=1)

RMM Location Details

The RMM Location Details attributes provide data for a single RMM location.

Location: Location name.

Location Def Exists: Indication if the Location is defined with a LOCDEF command in the RMM PARMLIB. (No=0; Yes=1)

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Management Type: The shelf management technique in use for the storage location. (Unknown=0; Bin=1; NoBins=2)

Media: Media name for the location.

Priority: Priority of this location relative to other locations.

Timestamp: Time table row was created in CT® timestamp format.

Type: Type of Location. (Shelf=0; Auto=1; Manual=2; Store=3)

RMM Owner Details

The RMM Owner Details attributes provide data for a single RMM owner.

Address1: Address Line 1.

Address2: Address Line 2.

Address3: Address Line 3.

Department: Department.

Email: Owner's email.

External Phone: External Phone number.

First Name: Owner First Name.

Internal Phone: Internal Phone number.

Last Name: Owner Last Name.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Name: Owner Name.

Number of Volumes: Total number of volumes owned.

Timestamp: Time table row was created in CT® timestamp format.

User node: User node.

Userid: Owner's Userid.

RMM Pending Actions

The RMM Pending Actions attributes provide data for a single RMM action.

Action: RMM Action Name. Values are as follows: (SCRATCH REPPLACE INIT ERASE RETURN NOTIFY)

Current Location: For Moves, the Current Location of the volumes.

Destination Location: For Moves, the Destination Location for the volumes.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Move Type: Return To Scratch move. (NoRTS=0; RTS=1)

Number of Volumes: Number of Volumes in specified action state.

Status: Status of the action. (Unknown=5; Complete=6; Confirmed=7; Pending=8)

Timestamp: Time table row was created in CT® timestamp format.

RMM System Summary

The RMM System Summary attributes provide summary data for RMM at various levels.

% Volumes R/W Previous Day: Percent of total volumes read or written on the previous day.

Average Datasets Per Volume: Average number of datasets per volume.

Average Volume Use Count: Average volume use count for all volumes associated with this summarized object.

Average Compression Ratio: The average compression ratio for all volumes associated with this summarized object.

Average compression ratio measures the effectiveness of data compression and is expressed as the percentage of the original data written by the application that was actually recorded on tape. This metric is dependent on the capacity and percent used values being reported back to RMM by the hardware, which might not be the case for older tape devices.

The compression ratio is based on the average % used for all volumes, the total capacity of the volumes, and the total usage of all the volumes and is calculated as follows:

- Total usage is the `(dataset blocksize)*(blockcount)`, which represents the amount of data written by the application for all datasets on the volume.
- Percent used (% used) is based on data actually written to the tape (assuming that the data is reported back by the hardware).
- Average capacity used is provided by this formula:

  ```
  (avg % used for all vols)*total capacity
  ```

- Average compression ratio is provided by this formula:

  ```
  100-(avg capacity used/total usage)
  ```

Average Dataset Size: Average dataset size, expressed in MB to 1 decimal place on all volumes associated with this summarized object.

Average Volume Utilization: Average volume utilization. Volume utilization (in other words, percent full) is based on where the last file on the tape ended. Only physical (not stacked or logical) volumes that are in Master or User status are included in the utilization calculation and distributions. Note that not all devices report this information back to RMM, so the value might be 0 for volumes that do in fact contain data.

Collection Status: Status of RMM Data collection process.

- Collection_completed=RMM_COLLECTION_COMPLETE
- Collection_is_waiting=RMM_COLLECTION_WAITING
- Collection_not_started=RMM_COLLECTION_INACTIVE
- Collection_parameter_errors=RMM_COLLECTION_PARAM_ERROR
- Collection_in_progress=RMM_COLLECTION_IN_PROGRESS
- Collection_halted_by_user=RMM_COLLECTION_HALTED
- Collection_halted_due_to_deadline=RMM_COLLECTION_DEADLINE
- Collection_terminated=RMM_COLLECTION_TERMINATED
- Collection_failed=RMM_COLLECTION_FAILED
- RMM_is_not_active=RMM_INACTIVE
- RMM_is_not_installed=RMM_NOT_INSTALLED

Data Created Previous Day: Data created on all volumes associated with this summarized object on the previous Julian day, expressed as MB to a single decimal place.

Datasets Created Previous Day: Total number of datasets created on the previous Julian day on all volumes associated with this summarized object.

Datasets R/W Previous Day: The number of datasets last read or written on the previous Julian day on all volumes associated with this summarized object (at the time that the last summary collection cycle ended).

Entry Volumes: Number of volumes in Entry status.

Erase Pending Volumes: Number of volumes in Erase Pending status.

Init Pending Volumes: Number of volumes in Init Pending status.

Init Volumes: Number of volumes in Init status.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Master Volumes: Number of volumes in Master status.

Name: Name of the object being summarized.

New Scratch Volumes: Number of volumes returned to scratch status during or since the last run of RMM expiration processing, as reported by the LISTCONTROL command.

Notify Pending Volumes: Number of volumes in Notify Pending status.

Permanent Read Errors Volumes: Total number of volumes associated with this summarized object with permanent read errors.

Permanent Write Errors Volumes: Total number of volumes associated with this summarized object with permanent write errors.

Permanent Read Errors: Total number of permanent read errors for all volumes associated with this summarized object.

Permanent Write Errors: Total number of permanent write errors for all volumes associated with this summarized object.

Physical Volumes: Number of physical (in other words, not LOGICAL or STACKED) volumes.

Replace Pending Volumes: Number of volumes in Replace Pending status.

Return Pending Volumes: Number of volumes in Return Pending status.

Scratch Pending Volumes: Number of volumes in Scratch Pending status.

Scratch Volumes: Number of volumes in Scratch status.

Summary Collection End Time: Time at which the last summary collection cycle ended.

Summary Collection Start Time: Time at which the last summary collection cycle started.

Temporary Read Errors Volumes: Total number of volumes associated with this summarized object with temporary read errors.

Temporary Write Errors Volumes: Total number of volumes associated with this summarized object with temporary write errors.

Temporary Read Errors : Total number of temporary read errors for all volumes associated with this summarized object.

Temporary Write Errors: Total number of Temporary write errors for all volumes associated with this summarized object.

Timestamp: Time table row was created in CT® timestamp format.

Total Error Volumes: Total number of volumes associated with this summarized object with permanent or temporary read or write errors.

Total Errors: Total number of temporary or permanent read or write errors for all volumes associated with this summarized object.

Total Permanent Errors Volumes: Total number of volumes associated with this summarized object with permanent read or write errors.

Total Temporary Errors Volumes: Total number of volumes associated with this summarized object with temporary read or write errors.

Total Usage: Total Usage in GB. Usage is the amount of data written to the tape by the application, before any compression, expressed as GB.

Total Capacity: Total Capacity in GB of all volumes associated with this summarized object. Note that some tape devices do not report capacity info to RMM and this field might therefore be zero (0).

Total Datasets: Total number of datasets on all volumes associated with this summarized object.

Total Permanent Errors: Total number of permanent read or write errors for all volumes associated with this summarized object.

Total Temporary Errors: Total number of temporary read or write errors for all volumes associated with this summarized object.

Total Volumes: Total number of volumes in RMM for this object.

Total Write Mounts: Total number of Write Mounts for all volumes associated with this summarized object.

Type: Type of the object being summarized.

User Volumes: Number of volumes in User status.

Volume 0% Utilized: Number of Volumes 0% Utilized.

Volume 1-10% Utilized: Volumes 1-10% Utilized.

Volume 11-20% Utilized: Volumes 11-20% Utilized.

Volume 21-30% Utilized: Volumes 21-30% Utilized.

Volume 31-40% Utilized: Volumes 31-40% Utilized.

Volume 41-50% Utilized: Volumes 41-50% Utilized.

Volume 51-60% Utilized: Volumes 51-60% Utilized.

Volume 61-70% Utilized: Volumes 61-70% Utilized.

Volume 71-80% Utilized: Volumes 71-80% Utilized.

Volume 81-90% Utilized: Volumes 81-90% Utilized.

Volume 91-100% Utilized: Volumes 91-100% Utilized.

Volumes R/W Previous Day: The number of volumes associated with the summarized object that were last read or written on the previous Julian day (at the time that the last summary collection cycle ended).

RMM Volume Details

The RMM Volume Details attributes provide data for a single RMM volume.

Account Number: User accounting information.

Assigned Date: For a MASTER or USER volume the Assigned date field specifies the date the volume was assigned to an owner. That is the date when the volume was first added or when a SCRATCH volume became a MASTER volume as a result of a nonspecific mount request.

Availability: The Availability status of the volume. Valid values are as follows:

- Pending Release - The volume is not available.
- On Loan - The volume is not available because it is on loan, possibly at a different site.
- Open - The volume has been opened for a write operation and has not yet been closed.
- Vital Record - The volume is a vital record.
- Unknown - The availability status of the volume is not known.

Bin: The shelf location where the volume is currently stored.

Bin Media: The media type assigned to the shelf location where the volume is currently stored.

Container: Container where the volume is currently stored.

Creating Job Name: Name of the job that created the volume.

Creating System ID: ID of the system where the volume was defined to RMM.

Creation Date : Creation date of the RMM object.

Current Location: Location where the volume is currently stored.

Current Location Type: The type of location where the volume is currently stored.

- Shelf=0
- Store_Builtin=1
- Manual=2
- Auto=3

- Store_Bins=4
- Store_NoBins=5

Dataset Name: Name of the dataset on the volume.

Dataset Record: RMM is recording details for all datasets on the volume (YES) or just the first dataset (NO). (No=0; Yes=1)

Dataset Count: Number of datasets on the volume, or sequence number for a specific dataset.

Date Last Read: Date the volume was last read from.

Date Last Written: Date the volume was last written to.

Days for Retention: Number of Days Until or since (negative number) the Retention date.

- Nonexpiring=-2147483648
- WHILECATLG=-286331154

Days Since Assigned: Number of Days Since Assigned date.

Days Since Creation: Number of days since the volume was created.

Days Since Last Read: Number of days since the volume was last read from.

Days Since Last Written: Number of days since the volume was last written to.

Days Since Movement Tracking: Number of days since the movement tracking date.

Days Until Expiration: Number of Days Until or since (negative number) the expiration date. NonExpiring=-2147483648

Days Until Original Expiration Date: Number of Days Until or since (negative number) Original Expiration date. Nonexpiring=-2147483648

density: The recording density of the volume.

- Undefined=0
- 1600=1
- 6250=2
- 3480=3
- Compact=4

Description: Descriptive text about the volume.

Destination Type: The type of location where the moving volume will be stored.

- Shelf=0
- Store_builtin=1
- Manual=2
- Auto=3
- Store_Bins=4
- Store_NoBins=5

Destination Bin Media: The media type assigned to the shelf location where the moving volume will be stored.

Destination Bin Number: The shelf location where the moving volume will be stored.

Destination Location: The name of the target location where the moving volume will be stored.

Erase on Release: Erase and initialize the volume before performing the release action. (No=0; Yes=1)

Erase Pending: Volume will be erased upon release. (No=0; Yes=1)

Expiration Date: The date on which the volume will expire and its status set to Release Pending.

Expiration Date Ignore: The expiration date of the volume is set to the current date when removed from VRS control. (No=0; Yes=1)

Home Location: The library location for the volume.

Home Location Type: The type of library location for the volume.

- Shelf=0
- Store_builtin=1
- Manual=2
- Auto=3
- Store_Bins=4
- Store_NoBins=5

In Transit: Volume is in transit to its destination location. (No=0; Yes=1)

Init on Release: Initialize the volume before performing the release action. (No=0; Yes=1)

Init Pending: Volume will be initialized upon release. (No=0; Yes=1)

Label Type: The label type of the volume. Values are as follows:

- AL - ANSI Label
- AUL - ASCII User Label
- BLP - Bypass Label Processing
- NL - No Label
- SL - Standard Label
- SUL - Standard User Label

Label Version: Version for the current label type.

- NoVersion=0
- Label_Version_1=1
- Label_Version_3=3
- Label_Version_4=4

Last Drive: Device address where the volume was last read or written.

Last Userid: Last Userid. The Last Changed by field specifies the ID of the user who last changed the details for this volume. If the volume was not last changed by a user, this field can contain one of these values:

- *OAM - volume last updated by OAM.
- *HKP - volume last updated by inventory management.
- *OCE - volume last updated by automatic volume recording when read or written to

Loan Location: Location of the volume if it is on loan.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Media Attributes: Media Attributes of the volume. Values are as follows:

- NONE - The volume has no special attributes.
- RDCOMPAT - This indicates that the volume was created using one recording format but might be mounted on a drive that supports that format for input processing, but another for output processing.

Media Compaction: A Compaction technique was used to record data to the volume. (Unknown=0; No=1; Yes=2)

Media Name: The Name of the volume's media.

Media Recording: The recording format of the volume.

- Non-Cartridge=0
- 18Trk=1
- 36Trk=2

- 128Trk=3
- 256Trk=4
- 384Trk=5
- EFMT1=6
- EFMT2=7

Media Type: The physical media type of the volume. Values are as follows:

- CST - The volume is a 3480 or 3490 cartridge.
- ECCST - Enhanced Capacity 3480 or 3490 cartridge.
- HPCT - High Performance Cartridge Tape.
- EHPCT - Extended High Performance Cartridge Tape
- ETC - IBM® TotalStorage® Enterprise Tape Cartridge.
- EWTC - IBM® TotalStorage® WORM Enterprise Tape Cartridge.
- EETC - IBM® TotalStorage® Enterprise Economy Tape Cartridge.
- EEWTC - IBM® TotalStorage® Enterprise Economy WORM Tape Cartridge.
- EXTC - IBM® TotalStorage® Enterprise Extended Tape Cartridge.
- EXWTC - IBM® TotalStorage® Enterprise Extended WORM Tape Cartridge.

Move Mode: How the volume moves between locations. Auto=0; Manual=1

Movement Tracking Date: The date when the volume started to move between two locations or when the confirmation that the last movement had occurred was performed.

MVS™ Use: Volume can be used with the MVS™ operating system. (No=0; Yes=1)

Next Volume: Next Volume Serial number for multi-volume datasets.

Notify on Release: Notify the volume owner that the volume has expired. (No=0; Yes=1)

Notify Pending: Notify the owner of the volume that the volume will be released. (No=0; Yes=1)

OCE: Volume information was recorded at the last Open, Close, or EOV event. (No=0; Yes=1)

Old Bin Media: The media type assigned to the shelf location where the volume was previously stored.

Old Bin Number: The shelf location where the volume was previously stored.

Old Location: The name of the location where the volume was previously stored.

Old Location Type: The type of location where the volume was previously stored.

- Shelf=0
- Store_builtin=1
- Manual=2
- Auto=3
- Store_Bins=4
- Store_NoBins=5

Original Expiration Date: Expiration Date in the volume label, which might have been changed in RMM using the CHANGEVOLUME command.

Owner: ID of the owner.

Owner Access: The level of access the owner has to the volume.

- Read=0
- Update=1
- Alter=2

Permanent Read Errors: Total number of permanent read errors on the volume.

Permanent Write Errors: Total number of permanent write errors on the volume.

Previous Volume: Previous Volume Serial number for multi-volume datasets.

Product Number: Number of the Product located on the volume.

Rack Number: Rack/shelf location number where the volume is located.

Release Actions: Action to take when RMM releases the volume.

- Unknown=0
- Scratch=8
- Replace=7
- Return=4

Replace Pending: Volume is to be replaced (for example, due to excessive I/O errors). (No=0; Yes=1)

Required Location: The intended target location for the volume.

Required Location Type: The intended target location type for the volume.

- Shelf=0
- Store_builtin=1
- Manual=2
- Auto=3
- Store_Bins=4
- Store_NoBins=5

Retain By Set: The volume is being retained because it is part of a multi-volume set. (No=0; Yes=1)

Retention Date: Highest date for all datasets on the volume after which the volume is no longer retained by the current VRS.

Return Pending: Volume will be returned to its owner upon release. (No=0; Yes=1)

Scratch Immediate: Return volume to scratch in a single inventory management run. (No=0; Yes=1)

Scratch Pending: Volume will return to scratch status after all other release actions are completed. (No=0; Yes=1)

Security Class: The Security Class of the volume.

Stacked Volume Count: Number of volumes contained in a stacked volume.

Storage Group: DFSMS Storage Group for the volume.

Temporary Read Errors: Total number of temporary read errors on the volume.

Temporary Write Errors: Total number of temporary write errors on the volume.

Timestamp: Time table row was created in CT® timestamp format.

Use Count: Number of times the volume has been opened for either read or write.

Vendor: Manufacturer or supplier of the media.

VM Use: Volume can be used with the VM operating system. (No=0; Yes=1)

VOL1 Label: Serial number of the volume as recorded in the standard internal volume label.

Volume Capacity: The physical capacity of the tape volume in MB as reported by the hardware.

> **Note:** The value is displayed for those devices which provide capacity information, such as the IBM® 3590 Magstar® and later devices.

Volume Serial: Volume Serial number.

Volume Status: The Status of the volume. Values are as follows:

- MASTER - The volume contains valid user data and cannot be overwritten.
- USER - The volume is assigned to a user and is in use.
- SCRATCH - The volume is free and available for use.
- INIT - The volume is a scratch volume that is waiting to be initialized. After the volume is initialized, the status becomes SCRATCH.
- ENTRY - The volume is a scratch volume that is moving into an automatic library; once entered, the status becomes SCRATCH or INIT.

Volume Type: The type of the volume as defined to RMM.

- Physical=0
- Logical=1
- Stacked=2

Volume Usage: How much of the volume has been taken up by the datasets that are located on it in KB. This value is calculated as the sum of the usage of all the datasets on the volume. The usage for a dataset is the product of the block count and the block size. Volume Usage is an indication of how much data has been written by an application prior to compression or compaction.

Volume Write Mount Count: Number of times the volume has been mounted for output and written to.

Volume Percent Full: How much of the volume has been taken up by the data written to it. The value is calculated using the reported position on the volume where the end of the last file was written.

> **Note:** This value is only available for those devices which provide positioning information, such as the IBM® 3590 Magstar® and later devices. When data is written to a volume it is most probably compressed so that the volume usage, capacity and percent full values are only an indication of the remaining, available capacity.

Volume Sequence: Sequence number of the volume in a multi-volume set.

Worldwide ID: ID value set in the volume by the manufacturer and recorded by DFSMSrmm™ when the volume is first used. DFSMSrmm™ ensures that the correct volume is mounted by checking that the WWID still matches the recorded value.

RMM VRS Details

The RMM VRS Details attributes provide details for a Vital Record Specification (VRS) for RMM (removable media manager).

Count: Specifies how many days or how many cycles of a dataset should be retained in the location specified in the Location field.

Delay Days: Specifies the number of days that a dataset is retained in its current location before it is sent to a storage location.

Delete Date: Date when the VRS is deleted.

Description: Description string.

Expiration Date Ignore: Whether or not to honor the volume expiration date. (No=0; Yes=1)

Job Name: Specifies a job name mask used to select only those datasets which match the VRS and which were created by certain jobs.

Location: Where vital records managed by this VRS are retained.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Name: Name of the VRS. Depending on the VRS type could be a dataset name, GDG base name, VOLSER or VRS name.

Next VRS: Name of the subsequent VRS in the chain.

Next VRS Type: Type of the subsequent VRS in the chain.

- Undefined=0
- Next=1
- And=2

Owner: Owner of the VRS.

Priority: Specifies a value to override the priority set for the location on a LOCDEF command.

Retain Until Expired: VRS will retain datasets until their expiration dates. (No=0; Yes=1)

Retain While Cataloged: Retain to the specified count value while the dataset is cataloged. (No=0; Yes=1)

Retention Type: How datasets are retained by this VRS.

- Cycles=1
- Days=2
- RefDays=3
- Volumes=4
- ExtraDays=5
- ByDaysCycle=6

Scratch Immediate: Pending release volume goes to scratch after expiration processing. (No=0; Yes=1)

Timestamp: Time table row was created in CT® timestamp format.

Type: The type of this VRS.

- Name=4
- Volume=5
- DSName=6
- PseudoGDG=7
- GDG=8

S3 PPRC Connection

The S3 PPRC Connection attributes provide information about the S3 PPRC connection.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

VOLSER: Volume serial number.

Device Number: Device number.

Subchannel Set: Subchannel Set.

Primary Storage Subsystem ID: Primary storage subsystem ID.

Primary Logical Storage Subsys: Primary logical storage subsystem ID.

Primary CCA: Primary CCA number.

Primary Serial Number: Primary physical unit serial number.

Secondary SSID: Secondary Storage Subsystem ID.

Secondary LSS: Secondary Logical Storage Subsystem ID.

Secondary Serial Number: Secondary physical unit serial number.

Path Count: Path count within a single primary-secondary pair.

Primary Adapter: ID within the Storage Subsystem.

Secondary Adapter: ID within the Storage Subsystem.

Secondary Status: Secondary Status.

Secondary WWNN: Secondary World Wide Node Name.

S3 PPRC Volume

The S3 PPRC Volume attributes provide information about the S3 PPRC volume.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

VOLSER: Volume serial number.

Storage Group: Storage Group.

Device Number: Device number.

Subchannel Set: Subchannel Set.

Primary Storage Subsystem ID: Primary storage subsystem ID.

Primary Logical Storage Subsys: Primary logical storage subsystem ID.

Primary CCA: Primary CCA number.

Primary Serial Number: Primary physical unit serial number.

First Secondary SSID: First Secondary SSID.

First Secondary LSS: First Secondary LSS.

First Secondary CCA: First Secondary CCA.

First Secondary SSN: First Secondary SSN.

Copy State: Copy State.

Suspended XRC: Suspended XRC.

Primary Volume: Primary Volume.

Secondary Volume: Secondary Volume.

Consistency Group Member: Volume is part of a consistency group.

Critical: Critical.

VPD: VPD.

Incremental Resync: Specifies whether the incremental resync change recording mechanism for Metro/Global Mirror should be enabled.

PPRC XRC Hyperswap: PPRC XRC Hyperswap.

XRC Session: XRC Session.

CopyCross Session: CopyCross Session.

Auto Resync Enabled: Auto Resync Enabled.

Withdrawn Preserve Mirror: This volume is a target of a withdrawn preserve mirror.

Suspended Time: Suspended Time.

First Trk Group Out of Sync: First Trk Group Out of Sync.

Last Trk Group Out of Sync: Last Trk Group Out of Sync.

Num of Trks Out of Sync: Num of Trks Out of Sync.

Num of Trks On Volume: Num of Trks On Volume.

Percent of Tracks Copied: Percent of Tracks Copied.

Primary WWNN: Primary World Wide Node Name.

S3 RLS Cache Structure Summary

The RLS Cache Structure Summary attribute group provides performance summary for each active RLS Cache Structure.

Cache Structure Name: DFSMS Cache Structure name.

Cast-out: The Cast-out counter.

Data Table Entry Reclaim: The Data table entry reclaim counter.

Directory Entry: The Directory entry counter.

Directory Entry Reclaim: The Directory entry reclaim counter.

Managed System: An internal node name used to identify the system.

Read-hit: The Read-hit counter.

Read-miss Assignment Suppressed: The Read-miss assignment suppressed counter.

Read-miss Directory-hit: The Read-miss directory-hit counter.

Read-miss Target Storage Class Full: The Read-miss target storage class full counter.

System Sysplex Name: The system or sysplex name from which the data is derived. If a sysplex name, the data is derived from all systems in the sysplex.

Target Storage Class Full: The Target storage class full counter.

XI Complement Invalidation: The XI for complement invalidation counter.

XI Directory Reclaim: The XI for directory reclaim counter.

XI Name Invalidation: The XI for name invalidation counter.

XI Write: The XI for write counter.

S3 XRC Session

The S3 XRC Session attributes provide information about the S3 XRC session.

Name: Session name.

HLQ: HLQ.

Address Space Name: Address Space Name.

Error Level: Error Level.

Last Record Processed Time: Last Record Processed Time.

State: State.

Utility Volume Config: Utility Volume Config.

Timeout: Configured timeout in milliseconds.

Total Volume Pairs: Total Volume Pairs.

Duplex Volume Pairs: Duplex Volume Pairs.

Copying Volume Pairs: Copying Volume Pairs.

Pending Volume Pairs: Volume pair that XRC has not yet synchronized or resynchronized.

Sequence Check Volume Pairs: A volume pair that contains non-timestamped writes. This temporary condition indicates that data consistency is uncertain.

Suspended Volume Pairs: An established volume pair in an XRC session where XRC is no longer copying data from the primary volume to the secondary volume. The volume pair remains in the XRC session.

Utility Volume Pairs: Utility Volume Pairs.

PAV Volume Pairs: The count of pairs using Parallel Access Volumes.

Fix Utility Enabled: Fix Utility Enabled.

Device Block Enabled: Device Block Enabled.

Write Pacing Enabled: Write Pacing Enabled.

Suspend On Long Busy Pending: Suspend On Long Busy Pending.

Enhanced Reader Support Enabled: Enhanced Reader Support Enabled.

Incremental Resync Enabled: Incremental Resync Enabled.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

S3 XRC Volume Pair

The S3 XRC Volume Pair attributes provide information about the S3 XRC volume pair.

Blocked: Blocked.

Command: Command.

Error Level: Error Level.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Options: Options.

Pace: Indicates the average pacing milliseconds injected per write channel program during the monitor interval.

Primary CCA: Primary CCA.

Primary Device Number: Primary device number.

Primary SCID: Storage controller ID.

Primary SCSN: Storage controller serial number.

Primary Storage Subsystem ID: Primary storage subsystem ID.

Primary VOLSER: Primary VOLSER.

Primary Volume Cylinders: Primary Volume Cylinders.

Record Sets per Second: The number of record sets created per second for this volume during the monitor interval.

Residual Count: The number of updated records for this device that the data mover has yet to read from the primary storage control for this volume.

Secondary CCA: Secondary CCA.

Secondary Device Number: Secondary device number.

Secondary Storage Subsystem ID: Secondary storage subsystem ID.

Secondary VOLSER: Secondary VOLSER.

Session Name: Session Name.

State: State.

Synchronization Percent: Synchronization Percent.

Threshold Count: The current blocking threshold count or write pacing level for the primary volume.

Timestamp: Timestamp.

Write Pacing Level: Write pacing level.

Writes per Second: The number of write channel programs completed per second for this volume during the monitor interval.

S3 Flash Copy Relation

The S3 Flash Copy Relation attributes provide information about flash copy relationships.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

VOLSER: Volume serial number.

Device Number: Device number.

Storage Subsystem ID: Storage subsystem ID.

Logical Storage Subsystem: Logical Storage Subsystem ID.

CCA Number: CCA number.

Storage Subsystem ID: Partner Subsystem ID.

Logical Storage Subsystem: Partner LSS number.

CCA Number: Partner CCA number.

Target: This is a target relationship.

Full Volume: Relationship is for all extents on volume.

Background Copy: Background copy will be (or is being) done on relation.

Background Copy Active: Background copy is running.

Persistent: Relation is persistent.

Target Write Protected: Target is write protected.

Validation Required: Relation data may be inaccurate.

Incremental: AKA change recording active.

Revertible: Revertible.

Source Write Inhibited: Source is write inhibited.

Target Space Efficient: Target volume is not pre-allocated from extent pool.

Failing On Out Of Space: Failing on out of space.

Failed: Failed.

SWI Due To Out Of Space: SWI due to out of space.

Preserve Mirror Configuration: Preserve mirror configuration.

Contiguous Extent Count: Contiguous extent count.

Source CCHH: Source CCHH.

Target CCHH: Target CCHH.

Track To Copy: Track to copy.

Established Time: Established Time.

Latest Activity Time: Latest Activity Time.

Sequence Number: Consistency Group Sequence Number.

S3 Flash Copy Volume

The S3 Flash Copy Volume attributes provide information about the S3 Flash Copy Volume.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

VOLSER: Volume serial number.

Storage Group: Storage Group.

Device Number: Device number.

Storage Subsystem ID: Storage Subsystem ID.

Logical Storage Subsystem: LSS number.

CCA Number: CCA number.

Control Unit Model: Storage Control Unit Model Type, e.g. "2107", "2105", "3990".

Serial Number: Serial Number.

Remote Serial Number: Remote Serial Number.

Active Relation Count: Count of mapped extents between volumes.

Max Relation Count: Maximum supported mapped extents for this volume.

Continuation Track: Continuation Track.

Flash Copy Capable: Can run Flash Copy on this volume.

Incapability Reason: Incapability Reason.

P-to-P Remote Copy Session: Point-to-Point Remote Copy Session (MetroMirror).

Extended Remote Copy Session: Extended Remote Copy Session.

Concurrent Copy Session: Concurrent Copy Session.

Flash Copy Session: Flash Copy Session.

Flash Copy Phase1 Active: Flash Copy Phase1 Active.

Target Volume Fixed Block: Target Volume Fixed Block.

Consistency Group Wait: Consistency Group Wait.

Relation Info Inaccessible: Relation Info Inaccessible.

Relation Info Truncated: Relation Info Truncated.

S3 Link Statistics

The S3 Link Statistics attributes provide information about the S3 links.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Storage Subsystem ID: Storage Subsystem ID.

Device Number: Device Number.

Serial Number: Serial Number.

Interface ID: Interface ID.

Interface Type: Interface Type.

Bits Per Second: Bits Per Second.

MB Per Second: MB Per Second.

GB Per Second: GB Per Second.

ECKD Read Bytes: ECKD Read Bytes.

ECKD Read Rate Bytes: ECKD Read Rate Bytes.

ECKD Read MB: ECKD Read MB.

ECKD Read Rate MB: ECKD Read Rate MB.

ECKD Read GB: ECKD Read GB.

ECKD Read Rate GB: ECKD Read Rate GB.

ECKD Write Bytes: ECKD Write Bytes.

ECKD Write Rate Bytes: ECKD Write Rate Bytes.

ECKD Write MB: ECKD Write MB.

ECKD Write Rate MB: ECKD Write Rate MB.

ECKD Write GB: ECKD Write GB.

ECKD Write Rate GB: ECKD Write Rate GB.

ECKD Read Operations: ECKD Read Operations.

ECKD Read Operations Rate: ECKD Read Operations Rate.

ECKD Write Operations: ECKD Write Operations.

ECKD Write Operations Rate: ECKD Write Operations Rate.

ECKD Read Accumulated Time: ECKD Read Accumulated Time.

ECKD Read Accum Time Rate: ECKD Read Accum Time Rate.

ECKD Write Accumulated Time: ECKD Write Accumulated Time.

ECKD Write Accum Time Rate: ECKD Write Accum Time Rate.

PPRC Read Bytes: PPRC Read Bytes.

PPRC Read Rate Bytes: PPRC Read Rate Bytes.

PPRC Read MB: PPRC Read MB.

PPRC Read Rate MB: PPRC Read Rate MB.

PPRC Read GB: PPRC Read GB.

PPRC Read Rate GB: PPRC Read Rate GB.

PPRC Write Bytes: PPRC Write Bytes.

PPRC Write Rate Bytes: PPRC Write Rate Bytes.

PPRC Write MB: PPRC Write MB.

PPRC Write Rate MB: PPRC Write Rate MB.

PPRC Write GB: PPRC Write GB.

PPRC Write Rate GB: PPRC Write Rate GB.

PPRC Read Operations: PPRC Read Operations.

PPRC Read Operations Rate: PPRC Read Operations Rate.

PPRC Write Operations: PPRC Write Operations.

PPRC Write Operations Rate: PPRC Write Operations Rate.

PPRC Read Accumulated Time: PPRC Read Accumulated Time.

PPRC Read Accum Time Rate: PPRC Read Accum Time Rate.

PPRC Write Accumulated Time: PPRC Write Accumulated Time.

PPRC Write Accum Time Rate: PPRC Write Accum Time Rate.

SCSI Read Bytes: SCSI Read Bytes.

SCSI Read Rate Bytes: SCSI Read Rate Bytes.

SCSI Read MB: SCSI Read MB.

SCSI Read Rate MB: SCSI Read Rate MB.

SCSI Read GB: SCSI Read GB.

SCSI Read Rate GB: SCSI Read Rate GB.

SCSI Write Bytes: SCSI Write Bytes.

SCSI Write Rate Bytes: SCSI Write Rate Bytes.

SCSI Write MB: SCSI Write MB.

SCSI Write Rate MB: SCSI Write Rate MB.

SCSI Write GB: SCSI Write GB.

SCSI Write Rate GB: SCSI Write Rate GB.

SCSI Read Operations: SCSI Read Operations.

SCSI Read Operations Rate: SCSI Read Operations Rate.

SCSI Write Operations: SCSI Write Operations.

SCSI Write Operations Rate: SCSI Write Operations Rate.

SCSI Read Accumulated Time: SCSI Read Accumulated Time.

SCSI Read Accum Time Rate: SCSI Read Accum Time Rate.

SCSI Write Accumulated Time: SCSI Write Accumulated Time.

SCSI Write Accum Time Rate: SCSI Write Accum Time Rate.

S3 VTS Cluster Worst Statistics Summary

The S3 VTS Cluster Worst Statistics Summary attribute group provides worst cluster statistics for the last 24 hours.

Active GB: Maximum number of GB of the active data for the last 24 hours.

Active Logical Volumes: Maximum number of active logical volumes for the last 24 hours.

Avg Copy Throttle: The maximum average copy throttle value during the interval for the last 24 hours.

Avg Deferred Copy Throttle: The maximum average deferred copy throttle value during the interval for the last 24 hours.

Avg Overall Throttle: The maximum average of all throttle values during the interval for the last 24 hours.

Average Virtual Mnt Pend Time: The minimum average amount of time in seconds that a z/OS system waited for a virtual tape volume to be mounted on this cluster of the selected TS7700 grid for the last 24 hours.

Avg Write Overrun Throttle: The maximum average throttle value during the interval for the last 24 hours.

Base Deferred Copy Throttle: The maximum base deferred copy throttle value which would be used if deferred copy throttle were to be applied for the last 24 hours.

CLIDMSER: Cluster ID and Machine Serial Number.

Cluster ID: This hexidecimal field indicates the Cluster ID.

Cluster Name: Name of this TS7700 Cluster.

Copy Throttle Reasons: The maximum reason(s) for copy throttling during the interval for the last 24 hours.

Deferred Copy Throttle Reasons: The maximum reason(s) for deferred copy throttling during the interval for the last 24 hours.

GB Read: Maximum GB Read by the Channel for the last 24 hours.

GB Written: Maximum GB Written by the Channel for the last 24 hours.

Grid Library Sequence Number: This 5 character field contains the Sequence Number of the Grid (Composite) Library.

Host Channel Activity: The maximum total number of GB (read + write) processed by the z/OS host channels to the virtual tape volumes associated with this cluster of the selected TS7700 grid for the last 24 hours.

Host Write Throttle Reasons: The maximum reason(s) for write overrun (host) throttling during the interval for the last 24 hours.

Machine Model: Indicates the node's model type.

Machine Type: Indicates the node's machine type.

Machine Serial Number: This 8 character field contains the serial number of this node.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

MB Read: Maximum MB Read by the Channel for the last 24 hours.

MB Written: Maximum MB Written by the Channel for the last 24 hours.

Node ID: This hexidecimal field indicates the Node ID.

Percent Copy Throttle: The maximum percentage of 30-second periods where there was at least one throttling value greater than zero and that copy was the predominant reason for the throttling for the last 24 hours.

Percent Deferred Copy Throttle: The maximum percentage of 30-second periods where deferred copy throttle from cache partition 0 of this Hnode was active (non-zero) for the last 24 hours.

Percent Write Overrun Throttle: The maximum percentage of 30-second periods where there was at least one throttling value greater than zero and that write overrun was the predominant reason for that throttling for the last 24 hours.

Physical Drives Available: Minimum number of physcial drives available on this cluster of the select TS7700 grid for the last 24 hours.

Reserved: Reserved.

Scratch Mounts: Maximum number of scratch mounts for the last 24 hours.

Scratch Volumes: The minimum number of empty physical cartridges for use in this cluster of the selected TS700 grid for the last 24 hours.

Tape Volume Cache Size: The maximum size of the Cluster's TVC in increments of GB (1024x1024x1000) for the last 24 hours.

Timestamp: Time table row was created in CT timestamp format.

Total GB Transferred: Maximum sum of GB Read and GB Written for the last 24 hours.

Total MB Transferred: Maximum sum of MB Read and MB Written for the last 24 hours.

Total Mounts: Maximum number of total mounts for the last 24 hours.

VE Code Level: The code level of the Virtualization Engine.

Volume Recall Percent: The maximum percentage of the virtual mounts that result in a physical mount in this cluster of the selected TS7700 grid for the last 24 hours.

SMS Cache Sets

The SMS Cache Sets attributes provide the name of the cache sets and the structures that they contain.

Cache Set Name: The name of this cache set. An SMS base configuration can contain a maximum of 256 Coupling Facility (CF) cache sets. Each set can contain up to eight Coupling Facility cache structures, which are defined by MVS™ policies.

Managed System: The SMS name of this system.

Structures: The structure names contained in this cache set.

SMS Configuration

The SMS Configuration attributes provide information about the base SMS installation.

Default Management Class: The name of the default management class. Management class is an optional SMS construct containing attributes that determine migration, backup and retention objectives for datasets.

Default Track Size: The tracks-size component of the default device geometry. SMS uses default device geometry when placing datasets, so that the amount of space is consistent and most efficiently used no matter what device type is chosen. Default device geometry isolates the user from the actual physical devices.

Default Tracks per Cylinder: The tracks-per-cylinder component of the default device geometry. SMS uses default device geometry when placing datasets, so that the amount of space is consistent and most efficiently used no matter what device type is chosen. Default device geometry isolates the user from the actual physical devices.

Default Unit Type: The unit-type component of the default device geometry.

Last Update User ID: The ID of the user that performed the last update.

Last Updated: The date and time of the last update to a storage group within the SMS configuration.

Managed System: The SMS name of this system.

MVS™ Level: The version number associated with the MVS™ system in use.

SCDS Name: The name of the Source Control Dataset. A linear VSAM dataset containing an SMS configuration.

SMS Level: The version number associated with the DFP Product in use.

SMS Status: The status of the SMS subsystem (active or inactive) on the local system.

SMS Data Class

The SMS Data Class attributes provide information about the SMS data classes that have been defined.

Average Record Length: The average record length used for space allocation.

Compaction: Specifies whether the datasets allocated under this data class should be compressed when stored on tape or hard disk drive.

Control Area Free Space Percent: The percentage of free space to be left in the control area when it is loaded or after a CA split.

Control Area Reclaim: The CA Reclaim indicates whether the DASD space for empty CAs may be reused on z/OS 1.12 or later systems.

Control Interval Free Space Percent: The percent of free space to be left in the control interval when it is loaded or after a CI or CA split.

Control Interval Size: The size of the control interval.

Cross Region Share Option: The amount of sharing of clusters allowed by tasks in the same system.

Cross System Share Option: The level of sharing of the datasets allocated under this data class among multiple systems.

Data Class Name: The name of the data class. A data class is a collection of allocation and space attributes, defined by the storage administrator, that are used to create a dataset.

Dataset Type: The format of the datasets allocated under this data class; either PDS, PDSE, HFS, or extended.

Description: A description of the data class.

Directory Blocks: The number of directory blocks in a persistent data store.

Expiration Date: The date the datasets allocated under this data class will be expired.

Extended Addressability: Allows a VSAM dataset to grow beyond four GB when allocated in the extended format.

Extent Relief: The EXTENT CONSTRAINT REMOVAL indicates whether or not a VSAM dataset is allowed to go beyond the 255 extents limit.

Index Options: The options for a VSAM index.

Key Length: The length of the key in bytes.

Key Offset: The displacement, in bytes, to the key field of a KSDS.

Last Update User ID: The ID of the user who last updated this data class definition.

Last Updated: The date and time that the data class definition was updated.

Logical Record Length: The logical record length of datasets allocated under this data class.

Managed System: The SMS name of this system.

Maximum Volumes: The maximum number of volumes datasets allocated under this data class can span.

Primary Space Allocation: The primary space allocation.

Record Bias: The system specifies whether to let VSAM determine how many and which type of buffers to use when accessing VSAM extended format datasets by batch processing.

Record Format: The record format of datasets allocated under this data class.

Record Organization: The organization of datasets allocated under this data class.

Record Scaling Factor: The scaling factor used to scale the average record length.

Reduce Space Percent: The percentage by which the allocation is reduced when the first attempt fails due to space being unavailable.

Retention Days: The number of days before datasets allocated under this data class are expired.

RLS Above Bar: Indicates if the buffers associated with a VSAM cluster are to be released at close of dataset. If no, then the buffers can be reused if the dataset is re-opened in an appropriate interval.

RLS Cache: The RLS CF CACHE VALUE field shows that VSAM RLS data with greater than 4K CI's defined to DFSMS CF cache structures can be cached.

Secondary Space: The number of units to allocate for secondary extents.

SMS Management Class

The SMS Management Class attributes provide information about the SMS management classes that have been defined.

Automatic Backup: Specifies whether objects (datasets) of this class are automatically backed up by DFHSM.

Automatic Migration: The attributes which specify whether a dataset is eligible for migration by command (command) automigration (auto) both, or none.

Backup Frequency: The backup frequency specifies how many days must elapse before DFHSM can back up datasets.

Backup Versions Dataset Deleted: Specifies the number of backup versions to retain for a dataset that has been deleted from primary.

Backup Versions: Specifies the number of backup versions to retain for a dataset.

Command Backup: Specifies the authority level necessary to issue backup commands for datasets in the management class. Values are as follows:

ADMIN: the storage administrator can issue backup commands
USER: user can issue backup commands
NONE: neither storage administrators nor users can issue backup commands

Copy Technique: Copy technique specifies the method to be used when data in this class is backed-up by DFSMSdss™ and DFSMShsm™.

- Standard
- Concurrent_Preferred
- Concurrent_Required
- Virtual_Preferred
- Virtual_Required
- Cache_Preferred
- Cache_Required

Description: A storage administrator supplied description of this management class.

Expiration NonUsage: Specifies the criteria to determine when an unreferenced dataset is eligible for deletion during automatic space management.

Expiration: The number of days or date when the dataset or object will be deleted.

GDS Elements on Primary: Specifies how many versions of a generation dataset group use the normal primary days non-usage attribute for auto space management processing.

Last Update User ID: Specifies the user ID that last updated this management class.

Last Updated: The date and time when this management class was last updated.

Managed System: The SMS name of this system.

Management Class Name: The management class name. A named collection of management attributes describing the retention, backup and class transition characteristics for a group of objects in an object storage hierarchy.

Minimum Days Level 1: Specifies the minimum number of days since the dataset was referenced on primary before a dataset can be migrated from level1 to level2 migration.

Minimum Days Primary: Specifies the minimum number of days since a dataset on primary was last referenced before it is eligible for migration.

Partial Release: The partial release specifies the conditions under which unused allocated space is released. Values are as follows:

NO: no release of unused space occurs.
YES: space is released at space management cycle time.
CONDITIONAL: if a non-zero secondary space allocation has been specified, release unused space at space management time.

Retain Extra Backup Version: Specifies the number of days to retain backup version in addition to the most recent one.

Retain Only Backup Version: Specifies the number of days to retain the most recent backup version of a deleted dataset. This attribute applies only to dataset no longer resident on primary.

Retention Limit: The retention limit value is a required value that limits the use of retention period (RETPD) and expiration date (EXPDT) values. These values are explicitly specified in JCL, are derived from data class definitions, or are explicitly specified in the OSREQ store macro. NO LIMIT effectively causes this value to be ignored.

Rolled off GDS: The rolled-off GDS action value indicates whether to expire rolled-off GDSs or make them eligible for migration.

SMS Storage Class

The SMS Storage Class attributes provide information about the SMS storage classes that have been defined.

Availability: Specifies whether dataset access should continue in the event of a single device failure. This determines whether the dataset is placed on RAID or duplexed volumes.

Cache Set Name: Specifies a name that is associated with a coupling facility structure that is used for datasets that are defined for VSAM record level sharing.

Description: The description of this storage class as entered by a storage administrator.

Direct Bias: Bias determines which volumes MSR (millisecond response time) performance numbers (READ, WRITE, or BOTH) to consider during volume selection. If you specify a READ (R) bias, a cache storage control should be available to allow caching. If you specify a WRITE (W) bias, the DASD fast write feather of an IBM® 90 storage controller with cache should be available to allow the use of DASD fast write. If you do not specify a value for bias (BLANK), the MSR time determines whether caching or DASD fast write are used.

Direct MSR Objective: Specifies the desired millisecond response time for direct access I/O operations. This is used at allocation time to select a device capable of satisfying this objective.

Disconnect Sphere: Indicates if the buffers associated with a VSAM cluster are to be released at close of dataset. If no, then the buffers can be reused if the dataset is re-opened in an appropriate interval.

Guaranteed Space: Specifies whether this storage class has the guaranteed space attribute. You can allocate space for single volume and multi-volume datasets by specifying a storage class with the guaranteed space attribute. The request fails when space is insufficient.

The guaranteed space attribute does not apply to objects, but it does apply to both VSAM and non-VSAM-managed datasets. For a multi-volume system-managed dataset, primary space is preallocated on all volumes. The first volume becomes the primary volume. The remaining volumes become candidate volumes with preallocated space. When the primary extent on the current volume becomes full, VSAM attempts to create secondary extents on the current volume. If there is not enough space left on the current volume, VSAM uses the preallocated primary extent on the next volume. The next volume is converted from a candidate volume to a primary volume.

Guaranteed Synchronous Writes: Specifies whether guaranteed synchronous write is defined for this storage class. Ensures that data transfer occurs and is complete between the channel and the device when write operations are issued to a PDS.

Last Update User ID: The last user ID to update this storage class definition.

Last Updated: The date and time the storage class definition was last updated.

Managed System: The SMS name of this system.

Sequential Bias: Bias determines which volumes MSR (millisecond response time) performance numbers to consider during the volume selection. Values are READ WRITE, or BOTH.

Sequential MSR Objective: Specifies the desired millisecond response time for sequential access I/O operations. This is used at allocation time to select a device capable of satisfying this objective.

Storage Class Name: The name that identifies this storage class.

Sysplex Name: The system or sysplex name from which the data is derived.

SMS Storage Group

The SMS Storage Group attributes provide information about the SMS storage groups that have been defined.

Automatic Backup System/Group: Specifies the name of the system or system group that is eligible to perform automatic backup processing.

Automatic Backup: Specifies whether the DASD volumes in this pool storage group are eligible for automatic backup processing.

Automatic Dump System/Group: Specifies the name of the system or system group that is eligible to perform automatic dump processing.

Automatic Dump: Specifies whether the DASD volume in this pool storage group are eligible for automatic dump processing.

Automatic Migration System/Group: Specifies the name of the system, or system group, that is eligible to perform an automatic migration processing.

Automatic Migration: Specifies whether this storage group is eligible for interval migration, primary space management migration, both, or no automatic migration.

Breakpoint Value: For Extended Address Volumes (EAV), when a disk space request is this size or more, expressed in cylinders, the system prefers to use the cylinder-managed space for that extent. This applies to each request for primary or secondary space for datasets that are eligible for the cylinder-managed space. If not enough cylinder-managed space is available, then the system uses the track-managed space or uses both areas. When the size of a disk space request is less than the breakpoint value, the system prefers to use the track-managed area and if enough space is not available there, then the system uses the cylinder-managed space or uses both areas. A value of 0 means that all EAS eligible datasets is allocated in the cylinder-managed space of a EAV by DADSM as long as space is available.

Copy Pool Storage Group: The copy pool backup storage group associated with this primary storage group.

Description: The description of this storage group as entered by a storage administrator.

Dump Classes: Specifies up to five unique dump class names. When DFSMSHSM dumps DASD volumes that belong to the storage group it directs their contents to the dump classes.

Extend Storage Group: The extend storage group associated with this primary storage group.

Guaranteed Backup Frequency: Specifies the number days within the last backup period in which the backup process should have a copy of each of the datasets within the storage group. NOLIMIT specifies that datasets in the storage group are backed up according to management class specifications.

High Allocation Migration Threshold %: Specifies the upper limit, as a percentage of volume occupancy, beyond which new allocations to this volume are avoided by SMS, if possible. In addition, this threshold is used by HSM to determine whether datasets should be migrated off a DASD volume.

Last Update User ID: Specifies the user ID of the storage administrator who last updated this storage group definition.

Last Updated: Specifies the date and time this storage group definition was last updated by a storage administrator.

Low Allocation Migration Threshold %: The value that is used as the threshold goal in reducing the amount of space occupied on a DASD volume in the storage group during interval migration or daily space management. A value of zero causes all datasets in this storage group to be migrated if interval migration is also selected for this storage group.

Managed System: The SMS name of this system.

Overflow Storage Group: The storage group has been defined as an overflow storage group.

Storage Group Name: The name of this storage group as defined by a storage administrator.

Storage Group Type: Specifies the type of storage group defined by this row.

Tape Library1: Specifies the name defined by a storage administrator for one of up to eight tape libraries that can own volumes within this storage group.

Tape Library2: Specifies the name defined by a storage administrator for one of up to eight tape libraries that can own volumes within this storage group.

Tape Library3: Specifies the name defined by a storage administrator for one of up to eight tape libraries that can own volumes within this storage group.

Tape Library4: Specifies the name defined by a storage administrator for one of up to eight tape libraries that can own volumes within this storage group.

Tape Library5: Specifies the name defined by a storage administrator for one of up to eight tape libraries that can own volumes within this storage group.

Tape Library6: Specifies the name defined by a storage administrator for one of up to eight tape libraries that can own volumes within this storage group.

Tape Library7: Specifies the name defined by a storage administrator for one of up to eight tape libraries that can own volumes within this storage group.

Tape Library8: Specifies the name defined by a storage administrator for one of up to eight tape libraries that can own volumes within this storage group.

Tape Select: A tape library name used as an index value to find a storage group that can contain tape volumes owned by that library.

Track Managed Low Threshold %: Used as the threshold goal for the track managed area in reducing the amount of space occupied on a DASD volume in the storage group during interval migration or daily space management. This attribute specifies a lower limit, as a percentage of volume occupancy. A value of zero causes all datasets in this storage group to be migrated if interval migration is also selected for this storage group.

Track Managed High Threshold %: Specifies the upper limit, as a percentage of volume occupancy. When this value is exceeded, new allocations to this volume in the track managed area are avoided by SMS, if possible. In addition, this threshold is used by HSM to determine whether datasets should be migrated off a DASD volume.

VIO Maximum Size (KB): Specifies the maximum size, in KB, of datasets to be allocated to VIO.

VIO Unit: Specifies the generic DASD device type that this storage group simulates.

Volser Select: A volume used as an index value to find the storage group that contains that volume.

SMS Storage Group Status

The SMS Storage Group Status attributes provide the SMS status of the storage group on each z/OS® image in the SMSplex.

Managed System: The SMS name of this system.

Storage Group Name: Specifies the name of this storage group as defined in the base configuration.

Storage Group Status: The status for this storage group for this system or system group.

DisAll = Does not permit datasets in this storage group to be accessed.
DisNew = Does not permit the allocation of new datasets in this storage group.
Enable = Permits access to datasets in this storage group.
Notcon = The storage group is not defined.
QuiAll = Does not schedule more jobs that access datasets in the storage group.
QuiNew = Does not schedule any jobs that create new datasets in this storage group.

System ID: Specifies the name of the system or system group whose status is being displayed.

SMS Systems

The SMS Systems attributes provide information about the z/OS® images that are in the same SMSplex.

Managed System: The SMS name of this system.

Name: Specifies the name of this system or system group as defined in the base configuration.

Type: Specifies whether this name is an individual system name or a system group name (name of a parallel sysplex).

SMSplex Navigation

The SMSplex Navigation attributes provide internal navigation information for the SMSplex.

CRQPlex: The target CRQPlex.

HSMPlex: The target HSMPlex.

Input Attribute: The input attribute to be evaluated.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Output Origin Node: The originnode of the target Storage agent.

SMF ID: The SMF ID of the target system.

Storage Group: The target storage group.

SMSplex Summary

The SMSplex Summary attributes provide information on storage in the SMSplex.

BMF Panic Mode Pct: The percent of LRU intervals during which BMF was in panic mode in tenths of a percent across the SYSplex.

CRQ Maximum Requests: The maximum number of requests on a common recall queue waiting for execution.

Disabled VTOC Index: Identifies if there are any disabled VTOC indices in the SMSplex.

Element Percent Full: The highest percentage of elements in tenths of a percent on the Common Recall Queue that are currently in use for any CRQPlex in the SMSplex.

Entry Percent Full: The highest percentage of entries in tenths of a percent on the Common Recall Queue that are currently in use for any CRQPlex in the SMSplex.

High Volume Fragmentation Index: The value of the highest fragmentation index within the SMSplex. This value does not include volumes which are unavailable for new dataset allocation.

High Volume Response Time: The highest volume response time for any volume in the SMSplex.

HSM Pct Free Space Data Component: The lowest ratio of the free space divided by the total allocated space for any DFSMShsm control dataset in the SMSplex.

HSM Status: An indicator if HSM is inactive on any system in the SYSplex.

Low Track Managed Free Space Cyl: The lowest amount of free track managed space within any storage group in the SMSplex in cylinders. This value does not include the space on volumes in an SMS storage group that have not been converted.

Low Track Managed Free Space GB: The lowest amount of free track managed space within any storage group in the SMSplex in GB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Low Track Managed Free Space MB: The lowest amount of free track managed space within any storage group in the SMSplex in MB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Low Track Managed Free Space Pct: The lowest percent of track managed area free space in any storage group in the SMSplex. This value does not include volumes which are unavailable for new dataset allocation.

Low Track Managed Free Space Trk: The lowest amount of free track managed space within any storage group in the SMSplex in tracks. This value does not include the space on volumes in an SMS storage group that have not been converted.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Oldest Request Age: The number of minutes the oldest request has waited to be completed. A large number of queued requests or a large completion time can be caused by an insufficient number of tasks started of this function type, a high rate of requests, or this function type being held.

RLS MAX Resp Time: The maximum RLS average response time for any storage class with RLS activity across the SYSplex.

RMM Percent Used: The highest percent used of any control dataset by RMM in the SMSplex.

RMM Subsystem Status: An indicator if RMM is inactive on any system in the SYSplex.

SMS CDS SYSplex Name: SMSplex CDS dataset name or SYSplex name if the SMSplex and SYSplex contain the same z/OS images.

Storage Group Used Space Pct: The highest percentage of used space in the SMSplex for any storage group. This includes both the track and cylinder managed areas of the volumes.

Storage Grp Low Free Space Cyl: The minimum free space in cylinders in any storage group in the SMSplex. This value does not include volumes which are unavailable for new dataset allocation.

Storage Grp Low Free Space GB: The minimum free space in GB in any storage group in the SMSplex. This value does not include volumes which are unavailable for new dataset allocation.

Storage Grp Low Free Space MB: The minimum free space in MB in any storage group in the SMSplex. This value does not include volumes which are unavailable for new dataset allocation.

Storage Grp Low Free Space Pct: The minimum percent of free space in any storage group in the SMSplex. This value does not include volumes which are unavailable for new dataset allocation.

Storage Grp Low Free Space Trk: The minimum free space in tracks in any storage group in the SMSplex. This value does not include volumes which are unavailable for new dataset allocation.

SYSplex Name: SYSplex name or the z/OS system name if the z/OS system is not a member of a SYSplex.

Track Managed High Volume Frag Index: Value of the highest track managed area fragmentation index for any volume within the SMSplex. This value does not include volumes which are unavailable for new dataset allocation.

Track Managed Used Space Percent: The highest percentage of used track managed space in the SMSplex for any storage group. This includes only the track managed areas of the volumes.

Storage Toolkit Requests

The Storage Toolkit Requests attributes provide details regarding actions that you have created and about the status of those actions.

Action Name: Name of the request as given by user.

Associated With Situation: The request is currently associated with the situation, if applicable.

Command: The command that the Storage Toolkit runs.

Command Property: The command properties to use for command resubmission.

Creator Name: Name of user who created action request.

Date Created: Date/time that the request was created.

Description: Description of the request as given by user.

Last Completed: Date/time that the most recent action completed.

Last Return Code: Return code for the most recently completed action.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Request ID: An ID number for the request.

Run Count: Number of times that the request has run.

Result Count: The number of results for the action.

Schedule: Dates and times when the action request runs. The value is a character string and specifies either a time or an interval. An **AT®**, **IN**, or **EV** prefix in a schedule specification identifies the type of schedule for running action requests:

- The **AT®** prefix is followed by a specific date and time, as in these examples. Each of the following example lines represents a separate example:
  AT®:*hh*:*mm*
  AT®:*YY.DDD.hh*:*mm*
  AT®:*YY.MM.DD*:*hh*: *mm*
  AT®:D*1234567*:*hh*: *mm*
  (Go to the end of this bullet list for descriptions of the preceding variables.)
- The **IN** prefix is followed by an interval value, as in these examples. Each of the following example lines represents a separate example:
  IN:*hh*:*mm*
  IN:*DDD*:*hh*: *mm*
  (Go to the end of this bullet list for descriptions of the preceding variables.)
- The **EV** (Every) prefix is followed by specific days and times of the week, as in this example: `EV:D1234567 :hh:mm`where:
  hh:*mm* is the hour and minute time or interval
  YY is the two-digit year (00 to 99)
  MM is the two-digit month (01 to 12)
  DD is the two-digit day of the month (01 to 31)
  DDD is the three-digit day of the year (001 to 366)

D*1234567* indicates the days of the week (1=Monday to 7=Sunday). You can specify from 1 to 7 digits (days) after the letter D.

Situation Name: The situation name associated with the toolkit command, if applicable.

Status: The status for the action request.

Type: Type of request.

Request types divided into three columns for readability		
Column 1	**Column 2**	**Column 3**
• DFSMShsm™ Volume Migration • DFSMShsm Volume Backup • DFSMShsm Dataset Migrate to Level 1 • DFSMShsm Dataset Migrate to Level 2 • DFSMShsm Dataset Backup • DFSMShsm Dataset Recover • DFSMShsm Dataset Recall • DFSMSdss Volume Release • DFSMSdss Volume Convert Test • DFSMSdss Volume Convert Prepare • DFSMSdss Volume Convert • DFSMSdss Dataset Copy or Move • DFSMSdss Dataset Release • DFSMSdss Dataset Compress • DFSMSdss Volume Defrag • DFSMSdss Release • DFSMShsm Dataset Migrate	• DFSMShsm Cancel Request • DFSMSdss Compress • DFSMShsm • DFSMSdss • DFSMSrmm • DFSMSrmm AddDataset • DFSMSrmm ChangeDataset • DFSMSrmm DeleteDataset • DFSMSrmm AddOwner • DFSMSrmm ChangeOwner • DFSMSrmm DeleteOwner • DFSMSrmm AddVolume • DFSMSrmm ChangeVolume Confirm • DFSMSrmm ChangeVolume • DFSMSrmm DeleteVolume • DFSMSrmm AddVRS	• DFSMSrmm ChangeVRS • DFSMSrmm DeleteVRS • IDCAMS • ICKDSF • TSO • Mainframe • Batch Job • IDCAMS Dataset Delete • IDCAMS Dataset Alter • IDCAMS Dataset Catalog • IDCAMS Dataset Allocate • IDCAMS Dataset Define Cluster • IDCAMS Dataset ListCat • IDCAMS Dataset Print • IDCAMS Dataset Repro • IDCAMS Dataset Verify

Storage Toolkit Result Summary

The Storage Toolkit Result Summary attributes provide a summary of the results from action requests that are either running or have completed.

Action Name: The name of the action.

Completed: Date/time that the action completed.

Execution Start: Date/time that the action request started.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Output Line Count: Number of lines of output.

Request ID: An ID number for the request.

Resource Count: Number of resources (VOLSERs, DSNs, and so on) that this action affected.

Result ID: An ID number for the result.

Return Code: Return code for the completed action.

Situation Name: The situation name associated with the toolkit command, if applicable.

Status: The status of the action. The status values are as follows:

Submitted
Executing
Completed
Cancelled
Scheduled
Held
Pending
BadToken
InvalidJCL
AuthorizationFailure
DataSetError
BadGroups
EmptyGroups
NonexistentGroups

Submitted: Date/time that the request was submitted.

Storage Toolkit Result Detail

The Storage Toolkit Result Detail attributes provide details regarding the results from action requests that are complete.

The result for action requests can contain output from one or more steps or locations. For this reason, the **Step Name** and **Dataset or DD Name** attributes identify the source of the details.

- When you run a command from a command dialog or issue a command from the Issue Command dialog, the values for the **Step Name** and **Dataset or DD Name** are predefined by the Storage Toolkit.

Command type	Step Name	Dataset or DD Name
DFSMSdss	KS3DSS1	SYSPRINT
DFSMShsm	KS3HSM1	KS3OUT
IDCAMS	KS3AMS1	SYSPRINT
ICKDSF	KS3DSF1	SYSPRINT
DFSMSrmm	KS3RMM1	KS3OUT
TSO	KS3TSO1	KS3OUT

- When you create JCL using the Create Batch Job dialog box, the values for the **Step Name** and **Dataset or DD Name** reflect the values that you specified in the **Files whose contents should be copied for later viewing** fields on the **Options** tab.
- When you select **Copy JCL and JES logs for later viewing** on the **JCL** tab, the **Step Name** values are predefined by the Storage Toolkit. The **Dataset or DD Name** values are blank.

Step Name	Description
*JCL	The JCL that was submitted
*JESMSGLG	The contents of the JESMSGLG log
*JESJCL	The contents of the JESJCL log
*JESYSMSG	The contents of the JESYSMSG log

No output from Mainframe commands is returned, because execution of the command is not synchronized with the execution of the batch job. Other reasons why you might not receive output include the following:

- The command that you are running might not have output that can be returned.
- A file or log to be copied for later viewing is empty.

Step Name: The name of a step in the JCL, the JCL log or JES log that identifies the source of this line of output.

Dataset or DD Name: The dataset or DD name that identifies the source of this line of output.

File Name: The file name associated with this action result.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Output: The lines of output that the action produced.

Request ID: An ID number for the request.

Resource Name: Name of the resource (VOLSER, DSN, and so on) for which this output was generated.

Result ID: An ID number for the result.

Storage Toolkit Result Resources

The Storage Toolkit Result Resources attributes provide data to track the identifiers for action names, requests, and results.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Resource Name: Name of the resource (VOLSER, DSN, and so on) for which this output was generated.

Request ID: An ID number for the request.

Result ID: An ID number for the result.

Symmetrix Configuration

The Symmetrix Configuration attributes provide status and configuration information about an EMC Symmetrix storage facility.

Cache Memory: The cache size in MB.

Cache Write Pending Ceiling Percent: The percentage of cache that can be dedicated to write I/O.

Code Check Sum: The microcode check sum.

Code Date: The date of the microcode (mdyy) in Hexadecimal format.

Code Version: The microcode version.

Logical Volumes: The number of logical units defined.

Managed System: The SMS name of this system.

Number of Directors: The number of directors.

Number of SSIDs: Specifies the number of SSIDs associated with this Symmetrix configuration.

Operating Hours: Specifies the number of hours since power-on.

Physical Volumes: Specifies the number of physical disks installed.

Serial Number Index: The eight low order characters of serial number which serve as the control unit index.

Serial Number: Specifies the serial number assigned.

SSID Index: One of up to 64 SSIDs that are associated with this Symmetrix configuration.

SSID_FWRD: The subsystem ID of the Symmetrix control unit that is controlling this device.

Time Stamp: Displays the date and time this data was collected from the virtual tape server (VTS).

Symmetrix Devices

The Symmetrix Devices attributes provide information about the physical disks in an EMC Symmetrix storage facility.

Buffer Ratio: The ratio of number of full buffers to number of empty buffers.

Cylinders: The number of cylinders mapped to the logical volume.

Director Number: The ID of the disk director managing this device.

Managed System: The SMS name of this system.

Mirror Number: The number of mirror volumes for this logical volume.

Not Ready Mirrors: The number of mirrors in a NOT-READY state.

Raid Type: The raid type, either raid 5 or raid 6.

RDF Device: The type of RDF device.

RDF Partner Symmetrix Device Number: The Remote Data Facility partner device number.

Remote Adapter Group: The remote adapter group.

SCSI Interface: The ID of the physical volume for which the logical volumes are being displayed.

Serial Number Index: The eight low order characters of serial number which serve as the control unit index.

SSID: The subsystem ID of the Symmetrix control unit that is controlling this device.

Symmetrix Device Number: The Symmetrix Logical device number.

Time Stamp: Displays the date and time this data was collected from the virtual tape server (VTS).

Valid Mirrors: The valid mirrors for this device.

Volser: The volume serial of a logical drive located on the physical drive.

Write Disabled Mirrors: The number of mirrors that are disabled for write.

Symmetrix Director Summary

The Symmetrix Director Summary attributes provide overview information about the directors in an EMC Symmetrix storage facility.

Director ID: The ID of the disk director managing these devices.

Director Index: The ID of the disk director managing these devices.

Disk IDX: The ID of the physical disk on which the logical devices are located.

Logical Volume Count: The number of logical volumes located on the physical drive.

MANAGED SYSTEM: The SMS name of this system.

SCSI Interface: The ID of the physical disk on which the logical devices are located.

Serial Number Index: The eight low order characters of serial number which serve as the control unit index.

SSID Index: The subsystem ID of the Symmetrix control unit that is controlling this device. To view the disk interfaces of a director controlled by another Symmetrix control unit, enter the four-character subsystem ID in this field to identify the control unit.

Subsystem ID: The subsystem ID of this Symmetrix disk director.

Time Stamp: Displays the date and time this data was collected from the virtual tape server (VTS).

Symmetrix Directors

The Symmetrix Directors attributes provide information about the disk directors on an EMC Symmetrix storage facility.

Cache Hit Pct: The percent of I/O activity that was satisfied without accessing the data from the disk drive itself over the interval. The interval corresponds to the cache interval.

Cache Hit Pct Long Interval: The percentage of prefetched tracks found in cache during a long interval. A long interval is cache hits since IPL or STATUS CLEAR command.

Cache Hit Pct Short Interval: The percentage of prefetched tracks found in cache during a short interval. A short interval can be up to 10 minutes.

Cache Hits Long Interval: The count of cache hits since IPL or STATUS CLEAR command.

Cache Hits Short Interval: The count of cache hits during a short interval. A short interval can be up to 10 minutes.

Cache Misses Long: The number of cache misses during a long interval.

Cache Misses Short: The number of cache misses during a short interval.

CCW Count: Total CCW chains over the interval. The interval corresponds to the cache interval.

CCWs per Second: The rate of CCW chains per second over the interval. The interval corresponds to the cache interval.

Director Function: The type of director.

Director Index: The ID of this director.

Director Integer: Director ID as an integer.

Director Number: Identifies the disk director.

Director Type: The type of director.

I/O Count: Total I/O operations over the interval. The interval corresponds to the cache interval.

I/O per Second: The rate of I/O operations per second over the interval. The interval corresponds to the cache interval.

I/O Requests Long Interval: The number of I/O requests for a long interval.

I/O Requests Per Second Long Interval: The average number of I/O requests per second during a long interval. A long interval is since IPL or the STATUS CLEAR command.

I/O Requests Per Second Short Interval: The average number of I/O requests per second during a short interval. A short interval can be up to 10 minutes.

I/O Requests Short Interval: The number of I/O requests for a short interval.

I/Os Per Second Long Interval: The average number of I/O operations per second during a long interval. A long interval is since IPL or STATUS CLEAR command.

I/Os Per Second Short Interval: The average number of I/O operations per second during a short interval. A short interval can be up to 10 minutes.

Long Cache Miss: The number of cache misses that were not satisfied by tracks that were in the process of being prefetched over the interval. The interval corresponds to the cache interval.

Long Cache Miss Pct: The percent of cache misses that were not satisfied by tracks that were in the process of being prefetched over the interval. The interval corresponds to the cache interval.

Long Interval: A long measurement interval expressed as hours, minutes, and seconds.

Managed System: The SMS name of this system.

Prefetched Tracks: The number of prefetched tracks over the interval. The interval corresponds to the cache interval.

Prefetched Tracks Unused: The number of prefetched tracks over the interval that were not referenced. The interval corresponds to the cache interval.

Prefetched Tracks Used: The number of prefetched tracks over the interval that were referenced. The interval corresponds to the cache interval.

Prefetched Tracks Used Pct: The percent of prefetched tracks over the interval that were referenced. The interval corresponds to the cache interval.

Prefetched Tracks Percent Used: The percentage of prefetched tracks used.

Serial Number Index: The eight low order characters of serial number which serve as the control unit index.

Short Cache Miss: The number of cache misses that were satisfied by tracks that were in the process of being prefetched over the interval. The interval corresponds to the cache interval.

Short Cache Miss Pct: The percent of cache misses that were satisfied by tracks that were in the process of being prefetched over the interval. The interval corresponds to the cache interval.

Short Interval: A short measurement interval expressed as hours, minutes, and seconds.

SSID Index: The SSID to which this director is attached.

SSID: The subsystem ID of the Symmetrix control unit controlling this disk director.

Time Stamp: Displays the date and time this data was collected from the virtual tape server (VTS).

Total CCW Chains: The total number of I/O requests handled by the director.

Total I/Os Long Interval: The total I/Os (reads and writes) for a long interval.

Total I/Os Short Interval: The total I/Os (reads and writes) for a short interval.

Total I/Os: The total number of I/Os.

Total Prefetch Cache Hits: The total number of prefetch cache hits.

Total Prefetch Tasks: The number of prefetch tasks.

Total Tracks Prefetched: The total number of prefetched tracks.

Total Write requests: Total write requests.

Unused Prefetch Tracks: The number of unused prefetch tracks.

Used Prefetch Tracks: The number of used prefetch tracks.

Write Pct Long Interval: The percentage of write commands during a long interval. The value is `(Write channel commands) / (Reads and writes)` A long interval is since IPL or STATUS CLEAR command.

Write Pct Short Interval: The percentage of write commands during a short interval. Write channel commands divided by reads and writes. A short interval can be up to 10 minutes.

Write Percent: The percent of write requests over the interval in tenths of a percent. The interval corresponds to the cache interval.

Write Requests: The number of write requests over the interval. The interval corresponds to the cache interval.

Write Requests Long Interval: The count of write requests during a long interval.

Write Requests Short Interval: The count of write requests during a short interval.

Symmetrix SSIDs

The Symmetrix SSIDs attributes provide a list of the SSIDs associated with an EMC Symmetrix storage facility.

Managed System: The SMS name of this system.

SSID Index: The SSID used as an index attribute.

Subsystem IDs: One of up to 64 SSIDs associated with this Symmetrix configuration.

System Summary

The System Summary attributes provide information on storage in the system.

BMF Panic Mode Pct: The percent of LRU intervals during which BMF was in panic mode in tenths of a percent across the SYSplex.

CRQ Maximum Requests: The maximum number of requests on a common recall queue waiting for execution.

CRQPlex Name: The name of the CRQPlex associated with the z/OS image.

Disabled VTOC Index: Identifies if there are any disabled VTOC indices in the SMSplex.

Element Percent Full: The highest percentage of elements in tenths of a percent on the Common Recall Queue that are currently in use for any CRQPlex in the SMSplex.

Entry Percent Full: The highest percentage of entries in tenths of a percent on the Common Recall Queue that are currently in use for any CRQPlex in the SMSplex.

High Fragment Index St Group: The storage group with the highest fragmentation index.

High Fragment Index St Group TA: The storage group with the highest fragmentation index in the track area.

High Volume Fragmentation Index: The value of the highest fragmentation index within the SMSplex. This value does not include volumes which are unavailable for new dataset allocation.

High Volume Response Time: The highest volume response time for any volume in the SMSplex.

High Vol Resp Time St Group: The storage group with the high MSR.

HSM Pct Free Space Data Component: The lowest ratio of the free space divided by the total allocated space for any DFSMShsm control dataset in the SMSplex.

HSMPlex Name: The name of the HSMPlex associated with the z/OS image.

HSM Status: An indicator if HSM is inactive on any system in the SYSplex.

Low Free Space Cyl St Group: The storage group with the low space in cylinders.

Low Free Space Cyl St Group TA: The storage group with the low space in cylinders in the track area.

Low Free Space GB St Group: The storage group with the low space in GB.

Low Free Space GB St Group TA: The storage group with the low space in GB in the track area.

Low Free Space MB St Group: The storage group with the low space in MB.

Low Free Space MB St Group TA: The storage group with the low space in MB in the track area.

Low Free Space Pct St Group: The storage group with the low free space percent.

Low Free Space Pct St Group TA: The storage group with the low free space percent in the track area.

Low Free Space Trk St Group: The storage group with the low space in tracks.

Low Free Space Trk St Group TA: The storage group with the low space in tracks in the track area.

Low Track Managed Free Space Cyl: The lowest amount of free track managed space within any storage group in the SMSplex in cylinders. This value does not include the space on volumes in an SMS storage group that have not been converted.

Low Track Managed Free Space GB: The lowest amount of free track managed space within any storage group in the SMSplex in GB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Low Track Managed Free Space MB: The lowest amount of free track managed space within any storage group in the SMSplex in MB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Low Track Managed Free Space Pct: The lowest percent of track managed area free space in any storage group in the SMSplex. This value does not include volumes which are unavailable for new dataset allocation.

Low Track Managed Free Space Trk: The lowest amount of free track managed space within any storage group in the SMSplex in tracks. This value does not include the space on volumes in an SMS storage group that have not been converted.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Oldest Request Age: Number of minutes the oldest request has waited to be completed. A large number of queued requests or a large completion time can be caused by an insufficient number of tasks started of this function type, a high rate of requests, or this function type being held.

RLS MAX Resp Time: The maximum RLS average response time for any storage class with RLS activity across the sysplex.

RMM Percent Used: The highest percent used of any control dataset by RMM in the SMSplex.

RMM Subsystem Status: An indicator of the RMM subsystem status.

SMF ID: SMSplex CDS dataset name or SYSplex name if the SMSplex and SYSplex contain the same z/OS images.

SMS CDS SYSplex Name: SMSplex CDS dataset name or SYSplex name if the SMSplex and SYSplex contain the same z/OS images.

Storage Grp Low Free Space Cyl: The minimum free space in cylinders in any storage group in the SMSplex. This value does not include volumes which are unavailable for new dataset allocation.

Storage Grp Low Free Space GB: The minimum free space in GB in any storage group in the SMSplex. This value does not include volumes which are unavailable for new dataset allocation.

Storage Grp Low Free Space MB: The minimum free space in MB in any storage group in the SMSplex. This value does not include volumes which are unavailable for new dataset allocation.

Storage Grp Low Free Space Pct: The minimum percent of free space in any storage group in the SMSplex. This value does not include volumes which are unavailable for new dataset allocation.

Storage Grp Low Free Space Trk: The minimum free space in tracks in any storage group in the SMSplex. This value does not include volumes which are unavailable for new dataset allocation.

Storage Group Used Space Pct: The highest percentage of used space in the SMSplex for any storage group. This includes both the track and cylinder managed areas of the volumes.

SYSplex Name: SYSplex name or the z/OS system name if the z/OS system is not a member of a SYSplex.

Track Managed High Volume Frag Index: Value of the highest track managed area fragmentation index for any volume within the SMSplex. This value does not include volumes which are unavailable for new dataset allocation.

Track Managed Used Space Percent: The highest percentage of used track managed space in the SMSplex for any storage group. This includes only the track managed areas of the volumes.

Used Space Pct St Group: The storage group with the highest percent of used space.

Used Space Pct St Group TA: The storage group with the highest percent of used space in the track area.

Tape Device

The Tape Device attributes provide status information about the tape devices in a tape group. A tape group is the set of devices like the IBM® ATL or VTS, the STK LSM or VTS, or the generic device type.

Device Address: The four hexadecimal digits that define a physical device. Also known as unit or unit address.

Device Status: The device status, including ONLINE, OFFLINE, ALLOCATED, or MOUNT PENDING.

Device Type: The generic tape drive type. Examples include 3480 and 3490.

Group Name: An IBM® Z OMEGAMON® AI for Storage® convention for identifying logical groupings of tapes by library of device. Tape subsystems that group volumes and devices using a library dataserver are considered libraries. All other devices are grouped by device considered non-library or NO LIB. Each tape library is considered to be a tape group whose name is the same as the library name.

I/O Count: The count of I/O operations by this drive during this interval. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

I/O Per Second: The number of I/O operations by the device during this interval. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Jobname: The name of the job to which the tape drive is currently allocated.

Managed System: The SMS name of this system.

Mount Pending Time: The amount of time in minutes that a device has been in MOUNT PENDING status. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Permanent Errors: The number of permanent read or write errors that the device has incurred. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Storage Group: The name of the SMS tape storage group to which the tape volume belongs.

Temporary Errors: The number of temporary read and write errors that the device has incurred. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Volume: The volume serial number of the tape mounted or being mounted on the drive.

Tape Group

The Tape Group attributes provide status information for tape devices like the IBM® ATL or VTS, the STK LSM or VTS, or the generic device type.

Available Drives: The number of tape drives that are online and have not been allocated to an application. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Available Slots: The number of slots in a tape library that are not in use. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Configured Drives: The number of tape drives configured to a tape group. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Console Name: The console associated with the tape library.

Drives Online: The number of tape drives that have been configured and are online. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

ID Number: The hardware identification of the tape library.

IO Per Second: The lowest number of I/Os, in I/Os per second, to a device currently active in the tape group. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Library Description: The description of the tape library.

Library Device: The tape library device model number.

Library Name: Shows the name of the tape library for which the information is displayed.

Library Type: The type of the tape library. Valid types are: automatic for automated tape libraries (ATLDS), and manual for manual tape libraries (MTLDS).

Managed System: The SMS name of this system.

Media1 Scratch Volumes: The number of scratch volumes of this type. Media1, for magnetic cartridges (3480 compatible). You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Media2 Scratch Volumes: The number of scratch volumes of this type. Media2, for enhanced capacity magnetic cartridges (3490E compatible). You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Media3 Scratch Volumes: The number of scratch volumes of this type. Media3, for magnetic cartridges (3590 compatible). You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Media4 Scratch Volumes: The number of scratch volumes of this type. Media4. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Media5 Scratch Volumes: The number of scratch volumes of this type. Media5. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Media6 Scratch Volumes: The number of scratch volumes of this type. Media6. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Media7 Scratch Volumes: The number of scratch volumes of this type. Media7. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Media8 Scratch Volumes: The number of scratch volumes of this type. Media8. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Media9 Scratch Volumes: The number of scratch volumes of this type. Media9. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Media10 Scratch Volumes: The number of scratch volumes of this type. Media10. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Media11 Scratch Volumes: The number of scratch volumes of this type. Media11. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Media12 Scratch Volumes: The number of scratch volumes of this type. Media12. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Mount Pend Time: The longest time in minutes that a device in the group has been waiting for a mount. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Mounts Pending: The number of tape drives waiting for mounts. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Permanent Errors: The highest number of permanent errors incurred by a tape drive in the group. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Temporary Errors: The highest number of temporary errors incurred by a tape drive in the group. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

Total Slots: The number of slots in a tape library. You can enable aggregation for this attribute in Tivoli® Data Warehouse.

TotalStorageDS Array Configuration

The TotalStorage Array Configuration attributes provide configuration information about the arrays in a TotalStorageDS storage facility.

Array ID: The array Identifier.

Array Type: The RAID type of the array.

Array Width: The width of the array (number of Disk Drive Modules in the array).

DDM Capacity: The capacity in GB of each Disk Drive Module in the array.

DDM Speed: The speed, in 1000 RPM, of the Disk Drive Modules in the array.

DDM Throttling: Indicates if a near-line DDM in the array is throttling performance due to temperature or workload.

Extent Pool ID: The identifier of the extent pool.

Managed System: The internal node name used to identify the system from which this row of data was obtained.

Parity Group: The name of the parity group. This is for HDS storage facilities only.

RAID Degraded: Indicates if one or more DDMs in the array need rebuilding.

Rank ID: The rank identifier.

RPM Exception: Indicates if a DDM with a slower RPM than the normal array DDMs is a member of the array as a result of a sparing action.

SSD Array: Indicates if the array is a solid state device.

Storage Facility ID: The storage facility identifier of the device (Type-Model-Serial Number).

Time Stamp: The date and time associated with this row of data.

TotalStorageDS Configuration

The TotalStorage® Configuration attributes provide configuration information about a TotalStorageDS storage facility.

Available Cache: The available cache in MB.

Configured NVS: The configured Non-Volatile Storage in MB.

Logical Subsystems: The total number of logical subsystems defined.

Logical Volumes: The total number of logical volumes defined.

Managed System: The internal node name used to identify the system from which this row of data was obtained.

Number of arrays with problems: The number of arrays with error conditions of Raid Degraded (one or more array members need rebuilding), DDM Throttling (a near-line DDM in the array is throttling performance due to temperature or workload), or RPM Exception (a DDM with an slower RPM than the normal array DDMs is a member of the array as a result of a sparing action).

Pinned NVS: The pinned NonVolatile Storage in MB.

Storage Facility ID: The storage facility identifier of the device (Type-Model-Serial Number).

Storage Facility Serial Number: The storage facility serial number of the device.

Time Stamp: The date and time associated with this row of data.

Total Arrays: The total number of arrays defined.

Total Cache: The total cache in MB.

Total Extent Pools: The total number of extent pools defined.

Total Parity Groups: The number of parity groups configured. This is for HDS storage facilities only.

Total Ranks: The total number of ranks defined.

Total SSD Arrays: The total number of solid state device arrays.

TotalStorageDS Extent Pool

The TotalStorage® Extent Pool attributes provide information about the extent pools in a TotalStorageDS storage facility.

Allocated Extents: The total number of real allocated extents in the pool.

Avg Read Op Rate: The average read operations per second to all ranks in the extent pool.

Avg Read Resp Time: The average read response time (in milliseconds) to all ranks in the extent pool.

Avg Write Op Rate: The average number of write operations per second to all ranks in the extent pool.

Avg Write Resp Time: The average write response time (in milliseconds) to all ranks in the extent pool.

Dyn Reloc Src Extents: The total number of extents used as source for dynamic relocation.

Dyn Reloc Tgt Extents: The total number of extents used as the target for dynamic relocation.

Dyn Reloc Tgt Extents: The total number of extents used as the target for dynamic relocation.

Extent Pool ID: The identifier of the extent pool.

Managed System: The internal node name used to identify the system from which this row of data was obtained.

Max Read Op Rate: The maximum read operations per second to any rank in the extent pool.

Max Read Resp Time: The maximum read response time (in milliseconds) to any rank in the extent pool.

Max Write Op Rate: The maximum number of write operations per second to any rank in the extent pool.

Max Write Resp Time: The maximum write response time (in milliseconds) to any rank in the extent pool.

Number of arrays with problems: The number of arrays with error conditions of Raid Degraded (one or more array members need rebuilding), DDM Throttling (a near-line DDM in the array is throttling performance due to temperature or workload), or RPM Exception (a DDM with an slower RPM than the normal array DDMs is a member of the array as a result of a sparing action).

Pool Type: The extent pool type, either CKD (Count Key Data) or FB (Fixed Block).

Pool Utilization: The percentage of the total extents in the pool that are currently allocated.

Real Extents Converted: The total number of real extents converted in the extent pool.

Real Extents: The total number of real extents in the extent pool.

Real Pool Capacity: The total real capacity in GB of the extent pool.

SSD Ranks: The status of solid state devices in the extent pool, including All, None, and Mixed (both SSD and non-SSD ranks).

Storage Facility ID: The storage facility identifier for the device (Type-Model-Serial Number).

Storage Facility Serial Number: The storage facility serial number of the device.

Time Stamp: The date and time associated with this row of data.

Total Read Operations: The total read operations to all ranks in the extent pool.

Total Write Operations: The total number of write operations to all ranks in the extent pool.

Virt Extents Converted: The total number of virtual extents converted in the extent pool.

Virtual Extents: The total number of virtual extents in the extent pool.

Virtual Pool Capacity: The total virtual capacity in GB of the extent pool.

TotalStorageDS Rank

The TotalStorage® Ranks attributes provide performance information about the ranks in a TotalStorageDS storage facility.

Average Read Resp: The average read response time to the rank over the interval, in milliseconds.

Average Write Resp: The average write response time to the rank over the interval, in milliseconds.

Avg Bytes/ Read Op: The average number of bytes per read operation.

Avg Bytes/ Write Op: The average number of bytes per write operation.

Bytes Read Per Sec: The rate per second at which bytes were read from the rank over the interval.

Bytes Written Per Sec: The rate per second at which bytes were written to the rank over the interval.

Extent Pool ID: The extent pool identifier.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Number of Arrays: The total number of arrays that make up the rank.

Number of arrays with problems: The number of arrays with error conditions of Raid Degraded (one or more array members need rebuilding), DDM Throttling (a near-line DDM in the array is throttling performance due to temperature or workload), or RPM Exception (a DDM with an slower RPM than the normal array DDMs is a member of the array as a result of a sparing action).

Parity Group: The name of the parity group. This is for HDS storage facilities only.

Rank ID: The identification number of the rank.

Read Ops Per Sec: The rate per second that read operations were issued to the rank over the interval.

SSD Rank: Indicates that the rank is backed by SSD arrays.

Storage Facility ID: The storage facility identifier for the device (Type-Model-Serial Number).

Time Stamp: The date and time associated with this row of data.

Total Bytes Read: The total number of bytes read from the rank over the interval.

Total Bytes Written: The total number of bytes written to the rank over the interval.

Total Read Ops: The total number of read operations issued to the rank over the interval.

Total Read Resp: The total read response time to the rank over the interval, in milliseconds.

Total Write Ops: The total number of write operations issued to the rank over the interval.

Total Write Resp: The total write response time to the rank over the interval, in milliseconds.

Write Ops Per Sec: The rate per second that write operations were issued to the rank over the interval.

TotalStorageDS SSIDs

The TotalStorage® SSIDs attributes provide a list of the SSIDs associated with a TotalStorageDS storage facility.

Managed System: The internal node name used to identify the system from which this row of data was obtained.

Storage Facility ID: The storage facility identifier for the device (Type-Model-Serial Number).

Subsystem IDs: A list of Subsystem IDs that are defined on this device.

UDG Cache Devices

The UDG Cache Devices attributes provide information about the cache and the associated metrics for UDG devices.

Avg Logical KB Read: The average number of KB read from the logical volume per second over the cache interval.

Avg Logical KB Written: The average number of KB written to the logical volume per second over the cache interval.

Avg Logical Read Time: The average time spent reading from the logical volume per I/O over the cache interval.

Avg Logical Write Time: The average time spent writing to the logical volume per I/O over the cache interval.

Avg Phys Read Response Time: The average time to read from the physical disk per operation in milliseconds over the cache interval.

Avg Phys Write Response Time: The average time to write to the physical disk per operation over the cache interval.

Bypass Cache Percent: The percentage of I/O's that bypassed cache this interval. These I/O requests are sent directly to DASD. Tracks are neither read from cache nor promoted to cache and their entries in the LRU list are not updated.

Cache Status: The caching status for this subsystem. *Active* means caching is enabled for this subsystem. *Inactive* means caching is disabled for this subsystem.

Cache to DASD Tracks per Sec: The rate per second of tracks written from cache to physical disk over the cache interval asynchronous of the channel to cache transfer of data.

Cache to DASD Total Tracks: The number of tracks written from cache to physical disk over the cache interval asynchronous of the channel to cache transfer of data.

CFW Read Percent: The percentage of READ requests that were satisfied from cache due to a previous cache fast write.

CFW Write Percent: The percentage of I/O operations that were Cache Fast Write requests. Cache Fast Write allows data to be written to cache only. The data is later written to DASD. CFW differs from DASD Fast Write in that the data is not written to Non-Volatile Storage (NVS).

Channel Index: The channel path identifier defining the path on which this device resides.

Control Unit Model: The model number of the control unit to which this device is attached.

DASD to Cache Non Promotes: The number of operations that did not promote a track to cache because free cache segments were not available.

DASD to Cache Non Promotes per Second: The rate per second of operations that did not promote a track to cache because free cache segments were not available.

DASD to Cache Tracks per Sec: The rate per second of tracks read from physical disk to cache over the cache interval for I/O operations.

DASD to Cache Tracks per Sec Non-sequential: The rate per second of tracks read from physical disk to cache over the cache interval for non-sequential I/O operations.

DASD to Cache Tracks per Sec Sequential: The rate per second of tracks read from physical disk to cache over the cache interval for sequential I/O operations.

DASD to Cache Total Tracks: The number of tracks read from physical disk to cache over the cache interval for I/O operations.

DASD to Cache Total Tracks Non-sequential: The number of tracks read from physical disk to cache over the cache interval for non-sequential I/O operations.

DASD to Cache Total Tracks Sequential: The number of tracks read from physical disk to cache over the cache interval for sequential I/O operations.

Device Address: Four hexadecimal digits that define the physical address of this device. Also known as unit or unit address.

Device Hex Address: The device address (in hexadecimal).

DFW Hit Percent: The percentage of DASD Fast Writes that were satisfied from cache without requiring access to DASD. DASD Fast Writes are written to both cache and Non-Volatile Storage (NVS).

DFW Retry Percent: The percentage of DASD Fast Write operations that could not complete due to Non-Volatile Storage (NVS) constraints.

DFW Status: Reports whether DASD Fast Write is enabled for this volume (the value can be Active, Inactive, Unknown, or Timeout).

Extent Pool ID: The extent pool identifier.

I/O Count: The number of I/O requests for this device in the current collection interval.

Inhibit Cache Percent: The percentage of I/O requests which set the "Inhibit Cache Load" mode in the channel program to prevent the data being cached. If the data is already in cache, the request is satisfied from cache.

Logical Control Unit Index: The system assigned subchannel address of the control unit to which this device is attached. An LCU is the logical representation of either a single control unit (with or without attached devices) or a group of control units that share devices.

Logical KB Read Total: The total number of KB read from the logical volume over the cache interval.

Logical KB Written Total: The total number of KB written to the logical volume over the cache interval.

Logical Read Time: The amount of time spent reading from the logical volume over the cache interval.

Logical Write Time: The amount of time spent writing to the logical volume over the cache interval.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Physical Device: Defines the type of cache controller on which this volume resides, such as 2105, TDS or Symmetrix.

Physical KB Read Rate: The KB read from the physical disk per second over the cache interval.

Physical KB Read Total: The total KB read from the physical disk over the cache interval.

Physical KB Write Rate: The KB written to the physical disk per second over the cache interval.

Physical KB Write Total: The total KB written to the physical disk over the cache interval.

Physical Read Rate: The rate of I/O read operations from the physical disk per second over the cache interval.

Physical Read Total: The total number of I/O read operations from the physical disk over the cache interval.

Physical Write Rate: The rate of I/O write operations to the physical disk per second over the cache interval.

Physical Write Total: The total number of I/O write operations to the physical disk over the cache interval.

Read Hit Percent: The percentage of READ requests that were satisfied from cache.

Read I/O Count: The number of READ requests for data on this device.

Read I/O Percent: The percentage of all I/O requests which were READ requests.

Solid State Device: An indication of whether the logical volume is backed by a solid state physical device.

Storage Facility ID: The name of the storage group to which this volume belongs. The storage group is an SMS construct used to group DASD devices in order to separate the physical requirements of storing data from the logical requirements.

Storage Group Index: The storage facility identifier for the device (type-model-serial number).

SSID Index: The full word representation of the SSID used for indexing in the product.

Subsystem ID: A two-byte unique value used to identify a 3990 or equivalent storage control unit.

Subsystem ID Hex: The subsystem ID associated with this device.

Time Stamp: The date and time this data was reported. Primarily for use in historical reporting.

Total Phys Read Response Time: The total time to read from the physical disk in milliseconds over the cache interval.

Total Phys Write Response Time: The total time to write to the physical disk in milliseconds over the cache interval.

User Group: A user DASD group is an OMEGAMON construct that enables you to define groups of DASD volumes by VOLSER or device address. You may also specify generic volsers and devices (patterns) by using a wildcard character.

Volume: This is the volume serial number of a disk device currently in use by this application.

Write Hit Percent: The percent of I/O WRITE operations that were satisfied without having to access the physical drive. This type of situation occurs as the result of DASD Fast Write or Cache Fast Write features. DASD Fast Write is a default in controllers that support it. Cache Fast Write may be requested in the Define Extent command of the channel program.

Write I/O Count: The number of WRITE requests for this device.

Volume Group Summary

The Volume Group Summary attributes provide aggregate performance and aggregate space information for SMS storage groups and user DASD groups.

Average Response Time: The average response time weighted by the unfiltered number of I/Os per volume. Unlike High Response Time, this value is calculated using all volumes in group, including ones with very low I/O count.

```
AvgMSR = <sum of (device response time * device I/O count)> / <sum of devices I/Os>
```

Busy: The percentage of time a resource is busy processing I/O requests. For PAV volumes, this value is normalized by dividing by the number of PAV exposures.

Description: This is a user-defined description used to identify an User DASD Group.

Device MPL: The highest MPL of all volumes associated with this storage group. This value is limited to those volumes that have met the minimum I/O requirements for exception processing.

Free Space Cylinders: The total amount of free space within the storage group, in cylinders. This value does not include the space on volumes in an SMS storage group that have not been converted.

Free Space GB: The total amount of free space within the storage group, in GB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Free Space MB: The total amount of free space within the storage group, in MB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Free Space Percent: The free space summed across all volumes within the group, expressed as a percentage. This value does not include volumes which are unavailable for new dataset allocation.

Free Space Tracks: The total amount of free space within the storage group, in tracks. This value does not include the space on volumes in an SMS storage group that have not been converted.

Group Name: SMS storage group name or OMEGAMON User DASD Group name.

Group Type: The type of storage group described. Group type can be either of the following:

- **SMSGROUP:** Volumes owned by an SMS storage group of this name.
- **USERGROUP:** Volumes that are part of a user-defined group, which can contain SMS-controlled volumes and non-SMS-controlled volumes.

High Busy Percent: The highest percentage of time busy processing I/O requests. For PAV volumes, this value is normalized by dividing by the number of PAV exposures.

High DFW Retry Percent: The highest percentage of DFW Retries. A DASD Fast Write retry occurs when a DASD fast write operation cannot complete due to NVS (nonvolatile storage) constraints.

High Response Time: The highest MSR (millisecond response time) of any volume in the storage group.

High Volume Fragmentation Index: The value of the highest fragmentation index within the group. This value does not include volumes which are unavailable for new dataset allocation.

Inactive Freespace Cylinders: Total amount of free space on inactive volumes within the storage group, in cylinders. This value is the sum of free space not available for allocation in the storage group across all volumes.

Inactive Freespace GB: Total amount of free space on inactive volumes within the storage group, in GB. This value is the sum of free space not available for allocation in the storage group across all volumes.

Inactive Freespace MB: Total amount of free space on inactive volumes within the storage group, in MB. This value is the sum of free space not available for allocation in the storage group across all volumes.

Inactive Freespace Tracks: Total amount of free space on inactive volumes within the storage group, in tracks. This value is the sum of free space not available for allocation in the storage group across all volumes.

Integer ID: This column is used internally to process user group toolkit requests originating as a result of take action from a situation. The value does not persist across TEMS restarts. This column is not kept historically.

Largest Free Extent Cylinders: The size of the largest free extent within the storage group, in cylinders. This value does not include volumes which are unavailable for new dataset allocation.

Largest Free Extent GB: The size of the largest free extent within the storage group, in GB. This value does not include volumes which are unavailable for new dataset allocation.

Largest Free Extent MB: The size of the largest free extent within the storage group, in MB. This value does not include volumes which are unavailable for new dataset allocation.

Largest Free Extent Tracks: The size of the largest free extent within the storage group, in tracks. This value does not include volumes which are unavailable for new dataset allocation.

Low Read Hit Percent: The lowest read hit percentage of any volume in this group. A Read Hit occurs when a read request for data is made by the processor and resolved by cache. As a result, the request does not require DASD access.

Low Volume Free Space %: The lowest free space value for a volume within the group, expressed as a percentage. This value does not include volumes which are unavailable for new dataset allocation.

Low Write Hit Percent: The lowest write hit percentage of any volume in this group. This value is limited to those volumes that have met the minimum I/O requirements for exception processing.

Managed System: The SMS name of this system.

Non-SMS Volumes: The number of non-SMS volumes in this storage group.

Non-Enabled Volumes: The indicator that this group contains SMS volumes with a status other than ENABLED.

Persistence Flag: The user defined Persistence Flag for User DASD Groups only, including (0) No persist after monitoring server restart and (1) Persist after monitoring server restart.

SMS Volumes: Represents the number of SMS volumes in this storage group.

Storage Group Status: Represents the SMS storage group status for the system being monitored.

Storage Group Type: The SMS storage group type.

Time Stamp: The date and time this data was collected.

Total Space Cylinders: The total amount of space within the storage group, in cylinders. This value does not include the space on volumes in an SMS storage group that have not been converted.

Total Space GB: The total amount of space within the storage group, in GB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Total Space MB: The total amount of space within the storage group, in MB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Total Space Tracks: The total amount of space within the storage group, in tracks. This value does not include the space on volumes in an SMS storage group that have not been converted.

Total Volumes: The number of volumes in this storage group.

Track Managed Free Space Cylinders: The total amount of track managed free space within the storage group, in cylinders. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Free Space GB: The total amount of track managed free space within the storage group, in GB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Free Space MB: The total amount of track managed free space within the storage group, in MB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Free Space Percent: The percent of track managed free space in the storage group. This value does not include volumes which are unavailable for new dataset allocation.

Track Managed Free Space Tracks: The total amount of track managed free space within the storage group, in tracks. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed High Volume Frag Index: Value of the highest track managed area fragmentation index within the group. This value does not include volumes which are unavailable for new dataset allocation.

Track Mngd Inact Freespace Cyls: Total amount of track managed free space on inactive volumes within the storage group, in cylinders. This value is the sum of free space not available for allocation in the storage group across all volumes.

Track Managed Inact Freespace GB: Total amount of track managed free space on inactive volumes within the storage group, in GB. This value is the sum of free space not available for allocation in the storage group across all volumes.

Track Managed Inact Freespace MB: Total amount of track managed free space on inactive volumes within the storage group, in MB. This value is the sum of free space not available for allocation in the storage group across all volumes.

Track Mngd Inact Freespace Trks: Total amount of track managed free space on inactive volumes within the storage group, in tracks. This value is the sum of free space not available for allocation in the storage group across all volumes.

Track Managed Largest Free Extent Cylinders: Size of the largest free track managed area extent within the storage group, in cylinders. This value does not include volumes which are unavailable for new dataset allocation.

Track Managed Largest Free Extent GB: Size of the largest free track managed area extent within the storage group, in GB. This value does not include volumes which are unavailable for new dataset allocation.

Track Managed Largest Free Extent MB: Size of the largest free track managed area extent within the storage group, in MB. This value does not include volumes which are unavailable for new dataset allocation.

Track Managed Largest Free Extent Tracks: Size of the largest free track managed area extent within the storage group, in tracks. This value does not include volumes which are unavailable for new dataset allocation.

Track Managed Low Volume Free Space %: Lowest track managed area free space value for a volume within the group, expressed as a percentage. This value does not include volumes which are unavailable for new dataset allocation.

Track Managed Total Space Cylinders: The total amount of track managed space within the storage group, in cylinders. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Total Space GB: The total amount of track managed space within the storage group, in GB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Total Space MB: The total amount of track managed space within the storage group, in MB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Total Space Tracks: The total amount of track managed space within the storage group, in tracks. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Used Space Cylinders: The total amount of track managed used space within the storage group, in cylinders. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Used Space GB: The total amount of track managed used space within the storage group, in GB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Used Space MB: The total amount of track managed used space within the storage group, in MB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Used Space Percent: The percentage of track managed used space in the storage group or the User DASD group. This includes both the track and cylinder managed areas of the volumes.

Track Managed Used Space Tracks: The total amount of track managed used space within the storage group, in tracks. This value does not include the space on volumes in an SMS storage group that have not been converted.

User Group Name: This is a user-defined name used to identify a User DASD group.

Used Space Cylinders: The total amount of used space within the storage group, in cylinders. This value does not include the space on volumes in an SMS storage group that have not been converted.

Used Space GB: The total amount of used space within the storage group, in GB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Used Space MB: The total amount of used space within the storage group, in MB. This value does not include the space on volumes in an SMS storage group that have not been converted.

Used Space Percent: The percentage of used space in the storage group or the User DASD group. This includes both the track and cylinder managed areas of the volumes.

Used Space Tracks: The total amount of used space within the storage group, in tracks. This value does not include the space on volumes in an SMS storage group that have not been converted.

VTOC Index Status: The worst status of any VTOC index in the group.

Volume Status

The Volume Status attributes provide the z/OS® and SMS status of a volume across the z/OS® images in an SMSplex.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

MVS™ Status: The z/OS® status of this device in the named system or group.

SMS Status: The SMS status of this device in the named system or group.

System ID: The SMF ID or group ID of the system from which this information was obtained.

Volser: The volume serial number of a disk device.

Volume Users

The Volume Users attributes provide data that assists a system administrator with identifying the applications that are allocated to a logical volume and the datasets that are being used on the volume by each application.

Application: The name of the application (STC, TSO user, or batch job) using the volume.

ASID: ASID of the application using the volume, displayed as 4 hex digits.

Dataset Name: The name of the dataset on the volume that is allocated.

DD Name: The DDNAME associated with the allocation.

Managed System: The name identifying the monitored system.

Open Count: Number of opens issued by the job against the dataset.

VTS Cache

The VTS Cache attributes provide performance and capacity information about the cache for a virtual tape server (VTS).

Average Recall Throttling Value: The average of the NON-ZERO throttling values over the last hour for the VTS where recall was the predominate reason.

Average Throttling Value: The average of both the ZERO and the NON-ZERO throttling values over the last hour for the VTS.

Average TVC Volume Age Minutes: The percentage of throttling over the last hour for the VTS where write overrun was the predominate reason.

Average Write Throttling Value: The average of the NON-ZERO throttling values over the last hour for the VTS where write overrun was the predominate reason.

Cache Hits: The count of cache hits a mount requests for a virtual tape volume in AVTS for which the virtual tape volume is already located in the TVC. In this case, there is no physical tape mount associated with the virtual mount request.

Fast Ready Mounts: The count of fast ready mounts. A mount attribute that allows a scratch mount request in a virtual tape server to be satisfied by the tape header information kept in the TVC for a previously-initialized virtual tape volume.

GB TVC Cache: The size of the cache in the TVC of the VTS in units of GB.

Managed System: The SMS name of this system.

Maximum TVC Volume Age Minutes: The maximum age in minutes, of all virtual tape volumes in the TVC at the end of the reporting interval.

Pre-migration Mounts: The number of processes used to copy virtual volumes to logical volumes. For example, from the TVC to a physical tape, by the VTS. Normally, this number is relatively high after high write activity.

Recall Throttling Percent: The percentage of throttling over the last hour for the VTS where recall was the predominate reason.

Stage-In Mounts: The count of stage-in mounts. A stage-in mount is a mount request for a virtual tape volume in a virtual tape system for which the virtual tape volume is not residing in the TVC. In this case, there is a physical tape mount associated with the virtual mount request to access the logical volume from a physical tape.

Time Stamp: Displays the date and time this data was collected from the VTS.

TVC Virtual Volume Count: The number of virtual volumes in the TVC at the end of the reporting interval.

TVC Virtual Volume Size MB: The average size in MB of the virtual volumes in the TVC at the end of the reporting interval.

VTS Library Type: Specifies the library type of this VTS. Values for library type are:

- STANDALONE VTS: a VTS that is not part of a peer-to-peer VTS environment.
- DISTRIBUTED VTS: a VTS that is part of a peer-to-peer VTS environment, but has not been designated for user interface commands.
- DISTRIBUTED USER INTERFACE VTS: a VTS that is part of a peer-to-peer VTS environment, and has been designated for user interface commands.
- COMPOSITE VTS: the logical view of a peer-to-peer VTS environment.

VTS Name: Specifies the name of this virtual tape server.

Write Throttling Percent: The percentage of 30-second intervals where there was a non-zero throttling value and write overrun was the predominate reason.

VTS Capacity

The VTS Capacity attributes provide utilization information about the physical tapes utilized by a virtual tape server (VTS).

Active Data TB: The number of bytes, in TB, of data being managed by this VTS.

Active Logical Volumes(1000): The number of logical volumes managed by the VTS that contains active data. This data is reported in units of 1000.

Empty Physical Volumes: The number of empty physical cartridges available for use by the VTS.

Empty Volume Free Space GB: The estimated storage capacity in GB, available on the empty 3590 cartridges managed by the VTS.

ESCON Channels Connected: The number of ESCON channels connected to this VTS.

Host Channel Activity GB: The total, READ + WRITE, number of GB (1,073,741,824) processed by the MVS host channels to the virtual tape volumes associated with the VTS.

Host Channel GB Read: The number of GB (1,073,741,824) read by the MVS host channels from the virtual tape volumes associated with the VTS.

Host Channel GB Written: The number of GB (1,073,741,824) written by the MVS host channels to the virtual tape volumes associated with the VTS.

Host Channel Read Percent: The percentage of all host channel byte movement that was the result of a read request by the MVS host channels from the virtual tape volumes associated with the VTS.

Managed System: The SMS name of this system.

Physical Tape GB Read: The number of GB (1,073,741,824) read by the MVS host channels from the virtual tape volumes associated with the VTS.

Physical Tape GB Written: The number of GB (1,073,741,824) written by the virtual tape server to the physical tape drives.

Physical Tape Read Percent: The percentage of all physical tape byte movement that was the result of a read request by the VTS.

Reclaim Threshold Percentage: The percentage identified when a physical cartridge is made available for reclamation. When the percentage of active data falls below the reclaim threshold, the physical volume is available to go through reclamation.

SCSI Channels Connected: The number of SCSI channels connected to this VTS.

Time Stamp: Displays the date and time this data was collected from the VTS.

Volumes 0-5% Utilized: The number of volumes containing 0 to 5% active data. This field is 0 for a composite library.

Volumes 10-15% Utilized: The number of volumes containing 10 to 15% active data. This field is 0 for a composite library.

Volumes 15-20% Utilized: The number of volumes containing 15 to 20% active data. This field is 0 for a composite library.

Volumes 20-25% Utilized: The number of volumes containing 20 to 25% active data. This field is 0 for a composite library.

Volumes 25-30% Utilized: The number of volumes containing 25 to 30% active data. This field is 0 for a composite library.

Volumes 30-35% Utilized: The number of volumes containing 30 to 35% active data. This field is 0 for a composite library.

Volumes 35-40% Utilized: The number of volumes containing 35 to 40% active data. This field is 0 for a composite library.

Volumes 40-45% Utilized: The number of volumes containing 40 to 45% active data. This field is 0 for a composite library.

Volumes 45-50% Utilized: The number of volumes containing 45 to 50% active data. This field is 0 for a composite library.

Volumes 50-55% Utilized: The number of volumes containing 50 to 55% active data. This field is 0 for a composite library.

Volumes 5-10% Utilized: The number of volumes containing 5 to 10% active data. This field is 0 for a composite library.

Volumes 55-60% Utilized: The number of volumes containing 55 to 60% active data. This field is 0 for a composite library.

Volumes 60-65% Utilized: The number of volumes containing 60 to 65% active data. This field is 0 for a composite library.

Volumes 65-70% Utilized: The number of volumes containing 65 to 70% active data. This field is 0 for a composite library.

Volumes 70-75% Utilized: The number of volumes containing 70 to 75% active data. This field is 0 for a composite library.

Volumes 75-80% Utilized: The number of volumes containing 75 to 80% active data. This field is 0 for a composite library.

Volumes 80-85% Utilized: The number of volumes containing 80 to 85% active data. This field is 0 for a composite library.

Volumes 85-90% Utilized: The number of volumes containing 85 to 90% active data. This field is 0 for a composite library.

Volumes 90-95% Utilized: The number of volumes containing 95 to 100% active data. This field is 0 for a composite library.

Volumes 95-100% Utilized: The number of volumes containing 95 to 100% active data. This field is 0 for a composite library.

VTS Library Type: Specifies the library type of this VTS. Values for library type are:

STANDALONE VTS: A VTS that is not part of a peer-to-peer VTS environment.
DISTRIBUTED VTS: A VTS that is part of a peer-to-peer VTS environment, but has not been designated for user interface commands.
DISTRIBUTED USER INTERFACE VTS: A VTS this is part of a peer-to-peer VTS environment, and has been designated for user interface commands.
COMPOSITE VTS : The logical view of a peer-to-peer VTS environment.
NOT APPLICABLE: The library type attribute does not apply to TS7700 resources.

VTS Name: The name of the VTS as assigned via the ISMF panels.

VTS Composite Library

The VTS Composite Library attributes provide performance statistics for a composite virtual tape server (VTS).

Cache Hit Mounts: The number of mount requests for a virtual tape volume in a VTS for which the virtual tape volume is already located in the TVC. In this case, there is no physical tape mount associated with the virtual mount request.

Category Mounts: The count of category or scratch mounts for this composite server.

Controller ID: The controller ID associated with this composite server.

GB Data Copied: The number of GB of data copied by this composite server.

Managed System: The SMS name of this system.

Read GB: The amount of data in GB read by the composite server.

Recall Mounts: The number of virtual mounts that result in a physical mount in the VTS.

Time Stamp: Displays the date and time this data was collected from the VTS.

Volumes Copied: The number of volumes copied by this composite server.

VTS Name: The name of the VTS as assigned through the ISMF panels.

Write GB: The amount of data in GB written by this composite server.

VTS Overview

The VTS Overview attributes provide an overview of the performance and configuration of a virtual tape server (VTS).

Average Virtual Mount Pend Time: The average amount of time in seconds that an MVS™ system waited for a virtual tape volume to be mounted on a virtual tape server.

Grid Library Sequence Number This five-character field contains the Library Sequence Number of the Grid (Composite) library.

Host Channel Activity GB: The total number of GB (read + write), processed by the MVS™ host channels to the virtual tape volumes associated with the VTS.

Library Serial Number: The manufacturer's serial number for the VTS.

Library Type: The physical tape library type.

Managed System: The SMS name of this system.

Maximum Virtual Mount Pend Time: The maximum average mount pending time for any cache partition within this TS7700.

Model number: The manufacturer's model number for the VTS.

Physical Drives Available: The number of physical tape drives available on the VTS.

Scratch Volumes: The number of empty physical cartridges available for use by the virtual tape server.

Time Stamp: Displays the date and time this data was collected from the VTS.

Virtual Devices Online: The number of virtual tape drives on line for the VTS.

Virtual Disconnect Time: The time that I/O spends, free from the channel, searching for or requesting data. This time includes the time spent waiting for the VTS to service the virtual device.

Volume Recall Percent: The percentage of virtual mounts that result in a physical mount in the VTS.

VTS Library Type: Specifies the library type of this VTS. Valid library types are:

- **STANDALONE VTS:** A VTS that is not part of a peer-to-peer VTS environment.
- **DISTRIBUTED VTS:** A VTS that is part of a peer-to-peer VTS environment, but has not been designated for user interface commands.
- **DISTRIBUTED USER INTERFACE VTS:** A VTS that is part of a peer-to-peer VTS environment, and has been designated for user interface commands.
- **COMPOSITE VTS:** The logical view of a peer-to-peer VTS environment.
- **NOT APPLICABLE:** The library type attribute does not apply to TS7700 resources.

VTS Name: The name of this virtual tape server.

VTS Physical Devices

The VTS Physical Devices attributes provide detailed performance and configuration information about the physical devices in a virtual tape server (VTS).

Average Concurrent Mounts: The average number of physical tape drives available to a VTS that were used at the same time.

Average Mount Pend Time: The average amount of time in seconds that this VTS waited for a tape volume to be mounted on a physical tape drive.

Managed System: The SMS name of this system.

Maximum Concurrent Mounts: The maximum number of physical tape drives available to VTS that were used at the same time.

Maximum Mount Pend Time: The maximum amount of time in seconds that this VTS waited for a tape volume to be mounted on a physical tape drive.

Migration Mounts: The number of physical tape drive mounts required to move new data from the TVC to the physical tape volumes associated with this VTS.

Minimum Concurrent Mounts: The minimum number of physical tape drives available to a VTS that were used at the same time.

Minimum Mount Pend Time: The minimum amount of time in seconds that this VTS waited for a tape volume to be mounted on a physical tape drive.

Physical Drives Available: The number of physical tape drives available to this VTS to move data between the TVC and the physical tape volumes.

Physical Drives Installed: The number of physical tape drives installed in this VTS.

Reclamation Mounts: The number of physical tape drive mounts required to remove obsolete data from the physical tape volumes and to consolidate active data on physical tape volumes associated with this VTS.

Stage Mounts: The number of mount requests in this VTS for a virtual tape volume for which the virtual tape volume is not located in the TVC. In this case, there is a physical tape mount associated with the virtual mount request to access the logical volume from a physical tape.

Time Stamp: Displays the date and time this data was collected from the VTS.

VTS LIB TY: Specifies the library type of this VTS.

Stand Alone: VTS that is not part of a peer-to-peer VTS environment.
Distributed: VTS that is part of a peer-to-peer VTS environment, but has not been designated for user interface commands.
Distributed UI: VTS that is part of a peer-to-peer VTS environment, and has been designated for user interface commands.
Composite: Logical view of a peer-to-peer VTS environment.
Not Applicable: The library type attribute does not apply to TS7700 resources.

VTS Name: Specifies the name of this virtual tape server.

VTS TS7700-series Cache Container

The VTS TS7700-series Cache Container attributes provides basic information about a cache in a TS7700 tape virtualization system. All attributes in this table are derived from the "HSM - Cache Container" record of the "Hnode HSM Historical Record."

Active Hnode: Specifies whether this Hnode was the active node in charge of Tape Volume Cache (TVC) management at the end of the interval being reported.

Average Copy Throttle: Average copy throttle value during the interval. The value presented is the average of the non-zero throttling values where copy was the predominant reason for throttling. The value is reported in thousandths of a second.

Average Overall Throttle: Average of all throttling values during the interval. The calculation includes samples for periods where throttling was both zero and non-zero. The value is reported in thousandths of a second.

Average Write Overrun Throttle: Average write overrun throttle value during the interval. The value presented is the average of the non-zero throttling values where write overrun was the predominant reason for throttling. The value is reported in thousandths of a second.

Cluster ID: Hexadecimal value indicating the cluster ID. Values are 0 to 7.

Grid Library Sequence Number: Library Sequence Number of the Grid (Composite) library.

Machine Model: Machine model of the node. Initially this field is set to "V06".

Machine Serial Number: Serial number of the node. This field is left justified and padded with blanks. The format is XX-YYYYY where XX is the plant of manufacture and the YYYYY is the sequence number of the node's machine. The dash character (-) is fixed. It should be noted that there can only be one machine per-cluster.

Machine Type: Machine type of the node. Initially this field is set to "3957".

Managed System: The SMS name of this system.

Node ID: A hexidecimal field that indicates the ID of the node.

Partition Container Number: An integer value 0-7 that indicates which cache partition you are viewing.

Percent Copy Throttle: Percentage of 30-second periods where there was at least one throttling value greater than zero and that copy was the predominant reason for throttling.

Percent Write Overrun Throttle: Percentage of 30-second periods where there was at least one throttling value greater than zero and that write overrun was the predominant reason for throttling.

Timestamp: Time table row was created in CT timestamp format.

Tape Volume Cache Size: Current size of the cluster's Tape Volume Cache (TVC) in increments of 1GB (1024 x 1024 x 1024). A TVC that is 1.7 TB in size is reported as x000006A4 (1700 decimal).

VE Code Level: Code level of the TS7700 Virtualization Engine™ (VE). The VE code level is expressed as `Version.Release.Modification.Fix` in a decimal form.

CLIDMSER: Cluster ID and Machine Serial Number

Total Mounts: Number of total mounts.

Scratch Mounts: Number of scratch mounts.

Percent Deferred Copy Throttle: The percentage of 30-second periods where deferred copy throttle from cache partition 0 of this Hnode was active (non-zero).

Average Deferred Copy Throttle: The average deferred copy throttle value during the interval. The value presented is the average of 30-second intervals of the deferred copy throttle value over the historical record interval. The value is reported in thousandths of a second.

Base Deferred Copy Throttle: The base deferred copy throttle value that would be used if deferred copy throttle were applied. The value is reported in thousandths of a second.

Host Write Throttle Reason(s): The reason(s) for write overrun (host) throttling during the interval.

Copy Throttle Reason(s): The reason(s) for copy throttling during the interval.

Deferred Copy Throttle Reason(s): The reason(s) for deferred copy throttling during the interval.

VTS TS7700-series Cache Partition Container

The VTS TS7700-series Cache Partition Container attributes provide basic information about a cache partition in a TS7700 tape virtualization system. All attributes in this table are derived from the "HSM - Cache - Partition Container" record of the "Hnode HSM Historical Record."

Average Cache Hit Mount Time: Average time, in milliseconds, taken to complete Cache Hit mounts during the interval. Mount time is accrued from the time the mount request is accepted by the system (PLF received, DE returned) until the x20 message is received from the Library Manager. The mount time is averaged into the interval time when the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.

Average Cache Miss Mount Time: Average time, in milliseconds, taken to complete Cache Miss mounts during the interval. Mount time is accrued from the time the mount request is accepted by the system (PLF received, DE returned) until the x20 message is received from the Library Manager. The mount time is averaged into the interval time when the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.

Average Fast Ready Mount Time: Average time, in milliseconds, taken to complete Fast-Ready mounts during the interval. Mount time is accrued from the time the mount request is accepted by the system (PLF received, DE returned) until the x20 message is received from the Library Manager. The mount time is averaged into the interval time when the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.

Cache Hit Mounts: Number of mount requests completed that the data was resident in the Tape Volume Cache (TVC) during this interval. A mount is accredited to the interval when the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.

Cache Miss Mounts: Number of mount requests completed that required recall from a stacked volume during this interval. A mount is accredited to the interval when the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.

Cluster ID: Hexadecimal value indicating the cluster ID. Values are 0 to 7.

Fast Ready Mounts: Number of mount requests completed using the Fast Ready method during this interval. A mount is accredited to the interval when the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.

Grid Library Sequence Number: Library Sequence Number of the Grid (Composite) library.

Machine Model: Machine model of the node. Initially this field is set to "V06".

Machine Serial Number: Serial number of the node. This field is left justified and padded with blanks. The format is XX-YYYYY where XX is the plant of manufacture and the YYYYY is the sequence number of the node's machine. The dash character (-) is fixed. It should be noted that there can only be one machine per-cluster.

Machine Type: Machine type of the node. Initially this field is set to "3957".

Managed System: The SMS name of this system.

Node ID: A hexidecimal field that indicates the ID of the node.

Partition Container Number: An integer value 0-7 that indicates which cache partition you are viewing.

Partition Size: Amount of cache assigned to this partition. The value is reported in increments of 1 GB.

Timestamp: Time table row was created in CT timestamp format.

VE Code Level: Code level of the TS7700 Virtualization Engine™ (VE). The VE code level is expressed as `Version.Release.Modification.Fix` in a decimal form.

VTS TS7700-series Cache Preference Group

The VTS TS7700-series Cache Preference Group attributes provide basic information about a cache preference group in a TS7700 tape virtualization system. All attributes are derived from the HSM HSM - Cache - Partition - Preference Group Container record of the "Hnode HSM Historical Record."

35 Day Average Cache Age: 35-day rolling average of cache age, in minutes, of the virtual volumes migrated out of the cache partition that were assigned to the preference group this data is for. Cache age is measured from when a volume is created or recalled into cache until it has been migrated from cache. The cache age of each volume is rounded up to the nearest minute. This data is calculated once an hour, on the hour. The data for this field is calculated at the end of the interval.

4 Hour Average Cache Age: Four-hour rolling average of cache age, in minutes, of the virtual volumes migrated out of the cache partition that were assigned to the preference group this data is for. Cache age is measured from when a volume is created or recalled into cache until it has been migrated from cache. The cache age of each volume is rounded up to the nearest minute. This data is calculated once an hour, on the hour. The data for this field is calculated at the end of the interval.

48 Hour Average Cache Age: 48-hour rolling average of cache age, in minutes, of the virtual volumes migrated out of the cache partition that were assigned to the preference group this data is for. Cache age is measured from when a volume is created or recalled into cache until it has been migrated from cache. The cache age of each volume is rounded up to the nearest minute. This data is calculated once an hour, on the hour. The data for this field is calculated at the end of the interval.

Cluster ID: Hexadecimal value indicating the cluster ID. Values are 0 to 7.

Data Resident in Cache: Amount of data in the TVC partition whose volumes are assigned to the preference this data is for. The value is reported in increments of 1 MB. Any residual data causes the value to be rounded up to the next higher value. This metric is the value at the end of the interval.

Grid Library Sequence Number: Library Sequence Number of the Grid (Composite) library.

Machine Model: Machine model of the node. Initially this field is set to "V06".

Machine Serial Number: Serial number of the node. This field is left justified and padded with blanks. The format is XX-YYYYY where XX is the plant of manufacture and the YYYYY is the sequence number of the node's machine. The dash character (-) is fixed. It should be noted that there can only be one machine per-cluster.

Machine Type: Machine type of the node. Initially this field is set to "3957".

Managed System: The SMS name of this system.

Node ID: A hexidecimal field that indicates the ID of the node.

Partition Container Number: Integer value 0-7 that indicates which cache partition you are viewing.

Preference Group Number: Integer value (0 or 1) that denotes the preference group of this particular row.

Timestamp: Time table row was created in CT timestamp format.

VE Code Level: Code level of the TS7700 Virtualization Engine™ (VE). The VE code level is expressed as `Version.Release.Modification.Fix` in a decimal form.

Virtual Volumes in Cache: Number of virtual volumes in the Tape Volume Cache (TVC) partition that are assigned to the preference group this data is for. This is the value at the end of the interval.

Volumes Migrated Last 35 Days: Number of virtual volumes migrated from the cache partition over the past 35 days that are assigned to the preference group this data is for. This data is calculated once an hour, on the hour. The data for this field is calculated at the end of the interval.

Volumes Migrated Last 4 Hours: Number of virtual volumes migrated from the cache partition over the past 4 hours that are assigned to the preference group this data is for. This data is calculated once an hour, on the hour. The data for this field is calculated at the end of the interval.

Volumes Migrated Last 48 Hours: Number of virtual volumes migrated from the cache partition over the past 48 hours that are assigned to the preference group this data is for. This data is calculated once an hour, on the hour. The data for this field is calculated at the end of the interval.

VTS TS7700-series Cluster

The VTS TS7700-series Cluster attributes provide basic information about a cluster in a TS7700 tape virtualization system. The attributes include metrics that are derived from standard sources, such as the virtualization (vNode) or hierarchical data storage management (hNode). These nodes are part of the TS7700 Virtualization Engine™ distributed-node architecture.

Average Virtual Mount Pend Time: Average amount of time in seconds that an MVS™ system waited for a virtual tape volume to be mounted on this cluster of the selected TS7700 grid.

Cache Percent Free: Cache Utilization Free Percent.

Cache Utilization Percent: The percent of cache utilization

Cluster ID: Hexadecimal value indicating the cluster ID. Values are 0 to 7.

Cluster Name: Name of this TS7700 Cluster.

Grid Library Sequence Number: Library Sequence Number of the Grid (Composite) library.

Host Channel Activity GB: Total number of GB (read + write), processed by the z/OS® host channels to the virtual tape volumes associated this cluster of the selected TS7700 grid.

Machine Model: Machine model of the node. Initially this field is set to "V06".

Machine Serial Number: Serial number of the node. This field is left justified and padded with blanks. The format is XX-YYYYY where XX is the plant of manufacture and the YYYYY is the sequence number of the node's machine. The dash character (-) is fixed. It should be noted that there can only be one machine per cluster.

Machine Type: Machine type of the node. Initially this field is set to "3957".

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Node ID: This hexadecimal field indicates the Node ID. Valid values are 0-1.

Physical Drives Available: Number of physical tape drives available on this cluster of the selected TS7700 grid.

Scratch Volumes: Number of empty physical cartridges available for use in this cluster of the selected TS7700 grid.

Timestamp: Time table row was created in CT timestamp format.

VE Code Level: Code level of the TS7700 Virtualization Engine™ (VE). The VE code level is expressed as Version.Release.Modification.Fix in a decimal form.

Volume Recall Percent: Percentage of virtual mounts that result in a physical mount in this cluster of the selected TS7700 grid.

CLIDMSER: Cluster ID and Machine Serial Number.

Tape Volume Cache Size: The current size of the Cluster's TVC in increments of GB (1024x1024x1000). A TVC that is 1.7TB in size will be reported as x00006A4 (1700 decimal).

Active Logical Volumes: Number of active logical volumes.

Active GB: Number of GB of active data.

Total Mounts: Number of total mounts.

Scratch Mounts: Number of scratch mounts.

MB Read: MB read by the channel.

GB Read: GB read by the channel.

MB Written: MB written by the channel.

GB Written: GB written by the channel.

Total MB Transferred: Sum of MB read and MB written.

Total GB Transferred: Sum of GB read and GB written.

VTS TS7700-series Cluster Replication

The VTS TS7700-series Cluster Replication attributes provide information about data transferred from the cache of a cluster to other clusters as part of a copy operation.

Timestamp: Time table row was created in CT timestamp format.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

MB By Copy: Data transferred from the cache of a cluster to other clusters as part of a copy operation.

GB By Copy: Data transferred from the cache of a cluster to other clusters as part of a copy operation.

Source Grid Library Sequence Number: This 5 character field contains the Source Sequence Number of the Grid (Composite) Library.

Source CLIDMSER: Source Cluster ID and Machine Serial Number.

Destination Grid Library Sequence Number: This 5 character field contains the Destination Sequence Number of the Grid (Composite) Library.

Destination CLIDMSER: Destination Cluster ID and Machine Serial Number.

VTS Virtual Devices

The VTS Virtual Devices attributes provide detailed internal performance information about the virtual devices associated with a virtual tape server (VTS).

Average Cache Hit Mount Time: The average time required for a mount request for a virtual tape volume in a TVS for which the virtual tape volume is already located in the TVC. In this case, there is no physical tape mount associated with the virtual mount request.

Average Concurrent Mounts: The average number of virtual tape drives available to this VTS that were used at the same time.

Average Fast Ready Mount Time: The average fast-ready mount time for tape drives in this VTS. Requests in a virtual tape server are satisfied by the tape header information kept in the TVC for a previously-initialized virtual tape volume. This approach avoids a physical tape mount for volume verification and tape management system integrity checking.

Average Mount Pend Time: The average amount of time in seconds that an MVS™ system waited for a virtual tape volume to be mounted on a virtual tape drive.

Average Recall Mount Time: The average time required to perform a recall mount on this VTS. A recall mount is a virtual mount that results in a physical mount.

Average Time Mount Used: The average time in seconds that a virtual tape drive in this VTS was mounted This time is accrued from the completion of the mount until a demount request is issued.

Data To Copy: The amount of data to be copied in GB.

LCU Disconnect Time: The average disconnect time in milliseconds for each I/O to each virtual drive in the logical control unit of this virtual tape server.

Managed System: The SMS name of this system.

Maximum Cache Hit Mount Time: The maximum time required for a mount request for a virtual tape volume in a TVS for which the virtual tape volume is already located in the TVC. In this case, there is no physical tape mount associated with the virtual mount request.

Maximum Concurrent Mounts: The maximum number of virtual tape drives available to this VTS that were used at the same time.

Maximum Fast Ready Mount Time: The maximum fast-ready mount time for tape drives in this VTS. Requests in a virtual tape server are satisfied by the tape header information kept in the TVC for a previously-initialized virtual tape volume. This approach avoids a physical tape mount for volume verification and tape management system integrity checking.

Maximum Mount Pend Time: The maximum amount of time in seconds that an MVS™ system waited for a virtual tape volume to be mounted on a virtual tape drive.

Maximum Recall Mount Time: The maximum time required to perform a recall mount on this VTS. A recall mount is a virtual mount that results in a physical mount.

Maximum Time Mount Used: The maximum time in seconds that a virtual tape drive in this VTS was mounted This time is accrued from the completion of the mount until a demount request is issued.

Minimum Cache Hit Mount Time: The minimum time required for a mount request for a virtual tape volume in a TVS for which the virtual tape volume is already located in the TVC. In this case, there is no physical tape mount associated with the virtual mount request.

Minimum Concurrent Mounts: The minimum number of virtual tape drives available to this VTS that were used at the same time.

Minimum Fast Ready Mount Time: The minimum fast-ready mount time for tape drives in this VTS. Requests in a virtual tape server are satisfied by the tape header information kept in the TVC for a previously-initialized virtual tape volume. This approach avoids a physical tape mount for volume verification and tape management system integrity checking.

Minimum Mount Pend Time: The minimum amount of time in seconds that an MVS™ system waited for a virtual tape volume to be mounted on a virtual tape drive.

Minimum Recall Mount Time: The minimum time required to perform a recall mount on this VTS. A recall mount is a virtual mount that results in a physical mount.

Minimum Time Mount Used: The minimum time in seconds that a virtual tape drive in this VTS was mounted This time is accrued from the completion of the mount until a demount request is issued.

Time Stamp: Displays the date and time this data was collected from the VTS.

Unit B180 Mounts: The mounts for unit B18-0.

Unit B181 Mounts: The mounts for unit B18-1.

Virtual Drives Configured: The number of virtual tape drives configured for this VTS.

Virtual Drives Online: The number of virtual tape drives online for this virtual tape server.

Volumes To Copy: The number of logical volumes that can be copied.

VTS LIB TY: Specifies the library type of this VTS.

Stand Alone: VTS that is not part of a peer-to-peer VTS environment.
Distributed: VTS that is part of a peer-to-peer VTS environment, but has not been designated for user interface commands.
Distributed UI: VTS that is part of a peer-to-peer VTS environment, and has been designated for user interface commands.
Composite: Logical view of a peer-to-peer VTS environment.

VTS Name: The name of this virtual tape server.

VTS Virtual Devices Summary

The VTS Virtual Devices Summary attributes provide response time information for the virtual devices assigned to a virtual tape server (VTS).

Busy Percent: The percentage of time this virtual drive was busy processing I/O requests.

Connect Time: The connect time for this virtual drive. Connect time is the time that I/O spends searching for or requesting data.

Device Address: The MVS™ address of this virtual drive.

Device Model: The generic tape drive type, such as 3590-1.

Device Status: The device status of this virtual drive including Online, Offline, Allocated, or Mount Pending.

Disconnect Time: The disconnect time for this virtual drive. Disconnect time is the time that I/O spends, free from the channel, searching for or requesting data.

I/O Per Second: The number of I/O operations completed per second for this virtual drive.

Managed System: The SMS name of this system.

Pend Time: The total pending for this virtual drive. Pending time is time that I/O is delayed in the path to the device. Pending time can be attributable to the channel or control unit path being busy.

Response Time: The average response time for this virtual drive in milliseconds. The amount of time (in milliseconds) it takes for an I/O request to complete.

Time Stamp: Displays the date and time this data was collected from the VTS.

VTS Name: The name of this virtual tape server.

VTSS Channels

The VTSS Channels attributes provide performance information for the VTSS channels.

Addresses Trapped: The number of unit addresses trapped by the channel interface.

Channel Name: The name of the channel.

Channel State: The state of the channel, either enabled or not enabled.

Channel Type: The type of channel. Either ESCON, FICON, or an IP link.

Control Unit_Busy: The percent of time that the control unit was busy.

I/O per Second: The number of I/Os per second over the previous recording interval.

Link Type: The type of link. Host is a link from z/OS to the VTSS. Real tape drive indicates a link from the VTSS to a physical tape device.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

TapePlex Name: The name of the tapeplex that contains this virtual tape subsystem.

VTSS_Name: Name of the VTSS (virtual tape subsystem).

VTSS Configuration

The VTSS Configuration attributes provide configuration information for the Virtual Tape Subsystems.

Active Migration Tasks: The current number of auto migration tasks.

Auto Migration Threshold: The threshold for auto migration.

Default ACS ID: The default ACS associated with the VTSS.

Disk Buffer % Utilization: The percentage of the cache that is currently utilized.

High Auto Migration Level: Specifies the level at which auto migration starts.

Host Accessible: Identifies if the VTSS is accessible from this z/OS host.

Host Supports Auto Migration: Identifies if this z/OS host supports auto migration.

Host Supports Migration: Identifies if this z/OS host supports migration of virtual tape volumes.

Host Supports Reclaim: Identifies if this z/OS host supports volume reclaim.

Immediate Migration Wait: The immediate migration wait time in minutes.

Low Auto Migration Level: Specifies the level at which auto migration stops.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Maximum Auto Migrate tasks: The maximum number of tasks available for auto migration.

Maximum VTV Size: The maximum size of a virtual tape volume.

Minimum Auto Migrate tasks: The minimum number of tasks available for auto migration.

Model: The VTSS model type.

Real Drives: The number of real tape drives installed.

Real Drives Online: The number of real tape drives online.

Scratch VTVs: The number of scratch virtual tape volumes.

TapePlex Name: The name of the tapeplex that contains this virtual tape subsystem.

Virtual Drives: The number of installed virtual tape drives.

Virtual Drives Online: The number of virtual tape drives online.

VTSS Capacity: The total capacity of the VTSS.

VTSS Name: The name of the VTSS.

VTSS State: The state of the virtual tape subsystem.

VTV Page Size: The size of a virtual tape volume page in cache.

VTVs Currently Resident: The number of virtual tape volumes currently in the cache.

VTSS RTD Queue

The VTSS RTD Queue Report attributes provide information for the Virtual Tape Subsystems Real Tape Devices Queue.

ACS ID: The ID of the Automated Cartridge System associated with this request.

Function: The request function, for example, migrate.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Mount Request: Identifies if waiting for a mount.

Parent ID: The numerical identifier of the parent process that issued this request.

Process ID: The numerical identifier for the request.

Storage Class: The storage class associated with the tape, or the name of the VTSS.

TapePlex Name: The name of the tapeplex that contains this virtual tape subsystem.

Task Type: The task type that is queued or active.

Virtual Tape Volume Movement: Identifies if there is movement of a virtual tape volume.

VTSS Name: The name of the VTSS.

Wait Reason: Identifies the reason if there is a wait.

VTSS RTD Status

The VTSS RTD Status attributes provide status information for the Virtual Tape Subsystem Real Tape Drives.

ACS ID: The ID of the Automated Cartridge System associated with this request.

Channel ID: The ID of the channel associated with the device.

Device Address: The address of the device.

Device Initial Selects: The number of initial selects made to the device.

Device Name: The name of the device.

Device Type: The device type.

LSM ID: The ID of the Library Storage Module associated with this device.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Multi Vol Cart Allocated: Whether or not a multi-volume cartridge has been allocated.

Multi Vol Cart Mounted: Whether or not a multi-volume cartridge has been mounted.

Multi Vol Cart Volser: The volser of the multi volume cartridge on the device.

Percent Device Available: The percentage of time that the device was available to be used.

Percent Device Connected: The percentage of time that the real tape drive was transmitting data.

Percent Device Utilized: The percentage of time that the device was utilized.

Read MB Second: The rate of read activity to the device in MB per second.

Requests Waiting: The number of requests waiting for this device.

RTD Operational Status: The operational status of the drive.

SMF ID: The SMF ID of the z/OS system that is being served by the drive.

TapePlex Name: The name of the tapeplex that contains this virtual tape subsystem.

VTSS Name: The name of the VTSS.

Write MB Second: The rate of write activity to the device in MB per second.

VTSS Status

The VTSS Status attributes provide information on the storage in the Virtual Tape Subsystems.

Available Cache: The amount of cache available for customer data.

Back End Capacity: The capacity of the cache buffers in non-compressed VTVs.

Base Cache: The amount of cache utilized by the system and not available for customer data.

Collected Free Back End Capcty: The amount of free space collected from the cache buffer.

Defined Standard Capacity: The capacity of the cache for VTVs that are compressed.

ECAM-T Channel Programs: The number of ECAM-T channel programs processed.

ECAM-T Messages Processed: The number of ECAM-T messages processed.

ECAM-T Msgs Bypassed Buffer: The number of ECAM-T messages bypassed because no buffer space was available.

ECAM-T Msgs Bypassed Busy: The number of ECAM-T messages bypassed because the configuration was busy.

Free Back End Capacity: The amount of free space in the cache buffer.

Free Space Collect Bytes Read: The number of bytes read during free space collection for the cache buffer.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Non-Volatile Storage: The amount of non-volatile storage.

Offline Cache: The amount of cache that is offline.

Pinned Cache: The amount of cache that is pinned.

Redundancy Group Count: The count of redundancy groups.

TapePlex Name: The name of the tapeplex that contains this virtual tape subsystem.

Total Free Space Collection: The total amount of free space collected from the cache buffer.

VTSS Name: The name of the VTSS.

VTSS Virtual Volumes

The VTSS Virtual Volumes attributes provide information on subpool storage in the Virtual Tape Subsystem.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Recording Media: The recording media used in the tape subpool.

Recording Tecnique: The recording technique used in the subpool.

Scratch Volumes: The number of scratch volumes in the subpool.

Subpool Location: The location of the subpool. Either VSM for a virtual tape subpool or the LSM address for a multi-volume cartridge subpool.

Subpool Name: The name of the subpool.

Tape Label Type: The type of tape label used in the subpool.

TapePlex Name: The name of the tapeplex that contains this virtual tape subsystem.

VTSS Name: The name of the VTSS.

Situations

In this section, you can find the information about predefined situations of OMEGAMON AI for Storage.

Overview of situations

You can use the predefined situations shipped with IBM® Z OMEGAMON® AI for Storage® as they are or you can modify them to meet your requirements.

If you choose to modify a predefined situation, first make a copy to ensure a fallback, if necessary. You can also create your own situations using the attributes provided by IBM® Z OMEGAMON® AI for Storage®.

> **Note:** Do not modify the product-provided situations. If you want to modify a product-provided situation, copy the situation, modify the copy, and rename the copy.

Definition of a predefined situation

A situation is a logical expression involving one or more system conditions. IBM® Z OMEGAMON® AI for Storage® uses situations to monitor the systems in your network. To improve the speed with which you begin using IBM® Z OMEGAMON® AI for Storage®, the product provides situations that check for system conditions common to many enterprises. You can examine and if necessary, change the conditions or values being monitored to those best suited to your enterprise. Be sure to start the situations that you want to run in your environment.

Using situations

You manage situations from the Tivoli® management portal using the Situation editor. Using the Situation editor you can perform the following tasks:

- Create a situation
- Save a situation
- Display a situation
- Edit a situation
- Start, stop, or delete a situation
- Investigate the situation event workspace for a situation

When you open the Situation editor, the left frame initially lists the situations associated with the Navigator item you selected. When you click a situation name or create a new situation, the right frame of the Situation editor opens to provide the following information about the situation and allow you to further define that situation:

Condition

View, add to, and edit the condition being tested.

Distribution

View the systems to which the situation is assigned and assign the situation to systems.

Expert advice

Write comments or instructions to be read in the situation event workspace.

Action

Specify a command to be sent to the system.
You can also specify a Storage Toolkit request to be run when a situation becomes true if IBM® Z OMEGAMON® AI for Storage® is installed and a storage table is enabled for Storage Toolkit commands.

Until

Reset a true situation when another situation becomes true or a specified time interval elapses.

Predefined situations descriptions

The following predefined situations are included in the IBM® Z OMEGAMON® AI for Storage® product.

KS3_Applic_Resp_Time_Critical

```
If VALUE S3_Application_Monitoring.High_Dataset_MSR GE 50
```

Monitors the response time components to determine the reason for a poor response time when an application is accessing a dataset and the response time is greater than the critical threshold. Also examine the volume for over-utilization, cache settings, and the response time components at the volume level.

KS3_Applic_Resp_Time_Warning

```
If VALUE S3_Application_Monitoring.High_Dataset_MSR GE 40 AND
VALUE S3_Application_Monitoring.High_Dataset_MSR LT 50
```

Monitors the response time components to determine the reason for a poor response time when an application is accessing a dataset and the response time is greater than the warning threshold. Also examine the volume for over-utilization, cache settings, and the response time components at the volume level.

KS3_Cachecu_Cache_Stat_Critical

```
If VALUE S3_Cache_Control_Unit.Cache_Status NE Active
```

Monitors for the condition where caching is not active for the control unit. Use the SETCACHE command to activate caching, if appropriate.

KS3_Cachecu_DFW_Retry_Critical

```
If VALUE S3_Cache_Control_Unit.DFW_Retry_Percent GE 2
```

Monitors for the condition where the percent of DASD fast write attempts that cannot be satisfied because a shortage of available nonvolatile storage (NVS) space exceeds the critical threshold. Check for pinned NVS and correct the problem if NVS is pinned. Otherwise, if the impact on performance is not acceptable, you need to move a volume or dataset to another cache control unit or add NVS to this control unit.

KS3_Cachecu_DFW_Retry_Warning

```
If VALUE S3_Cache_Control_Unit.DFW_Retry_Percent GE 1 AND
VALUE S3_Cache_Control_Unit.DFW_Retry_Percent LT 2
```

Monitors for the condition where the percent of DASD fast write attempts that cannot be satisfied because a shortage of available nonvolatile storage (NVS) space has exceeded the warning threshold. Check for pinned NVS and correct the problem if NVS is pinned. Otherwise, if the impact on performance is not acceptable, move a volume or dataset to another cache control unit or add NVS to this control unit.

KS3_Cachecu_Inact_Vols_Critical

```
If VALUE S3_Cache_Control_Unit.Deactivated_Volumes GE 15
```

Monitors for the condition where the number of deactivated volumes on the control unit exceeds the critical threshold. You can use the SETCACHE command to activate caching on the volumes, if necessary.

KS3_Cachecu_Inact_Vols_Warning

```
If VALUE S3_Cache_Control_Unit.Deactivated_Volumes GE 10 AND
VALUE S3_Cache_Control_Unit.Deactivated_Volumes LT 15
```

Monitors for the condition where the number of deactivated volumes on the control unit exceeds the warning threshold. You can use the SETCACHE command to activate caching on the volumes, if necessary.

KS3_Cachecu_NVS_Stat_Critical

```
If Value_S3_Cache_Control_Unit.NVS_Status NE Active
```

Monitors for the condition where nonvolatile storage is not active for the control unit. All writes to volumes on the control unit are written directly to the hard disk drive. Use the SETCACHE command to activate NVS (nonvolatile storage), if appropriate.

KS3_Cachecu_Read_HitP_Critical

```
If VALUE S3_Cache_Control_Unit.Read_Hit_Percent LE 50 AND
VALUE S3_Cache_Control_Unit.Read_Hit_Percent GT 0
```

Monitors for the condition where the percent of read I/O requests resolved from cache has fallen below the critical threshold. If performance is a problem, look for volume with a low read hit percent and consider moving them to another control unit to balance the load. This condition can be caused by cache-unfriendly applications or a shortage of cache.

KS3_Cachecu_Read_HitP_Warning

```
If VALUE S3_Cache_Control_Unit.Read_Hit_Percent LE 60 AND
VALUE S3_Cache_Control_Unit.Read_Hit_Percent GT 50
```

Monitors for the condition where the percent of read I/O requests resolved from cache has fallen below the warning threshold. If performance is a problem, look for volume with a low read hit percent and consider moving them to another control unit to balance the load. This condition can be caused by cache-unfriendly applications or a shortage of cache.

KS3_Cachecu_Trk_Dstg_Critical

```
If VALUE S3_Cache_Control_Unit.Track_Destaging_Rate GE 70
```

Monitors for the condition where the rate at which tracks are being removed from cache and written to DASD exceeds the critical threshold. If performance is being impacted, you need to migrate datasets or volumes to another cache control unit. An alternative is to increase the cache capacity.

KS3_Cachecu_Trk_Dstg_Warning

```
If VALUE S3_Cache_Control_Unit.Track_Destaging_Rate GE 50 AND
VALUE S3_Cache_Control_Unit.Track_Destaging_Rate LT 70
```

Monitors for the condition where the rate at which tracks are being removed from cache and written to DASD exceeds the warning threshold. If performance is being impacted, you need to migrate datasets or volumes to another cache control unit. An alternative is to increase the cache capacity.

KS3_Cachecu_Trk_Stag_Critical

```
If VALUE S3_Cache_Control_Unit.Track_Staging_Rate GE 70
```

Monitors for the condition where the movement of tracks from the physical device to cache has exceeded the critical threshold. If performance is impacted, you might need to move the logical volume that is causing the excessive activity or to move datasets on the logical volume.

KS3_Cachecu_Trk_Stag_Warning

```
If VALUE S3_Cache_Control_Unit.Track_Staging_Rate GE 50 AND
VALUE S3_Cache_Control_Unit.Track_Staging_Rate LT 70
```

Monitors for the condition where the movement of tracks from the physical device to cache has exceeded the warning threshold. If performance is impacted, you might need to move the logical volume that is causing the excessive activity or to move datasets on the logical volume.

KS3_Cachecu_Write_HitP_Critical

```
If VALUE S3_Cache_Control_Unit.Write_Hit_Percent LE 45 AND
VALUE S3_Cache_Control_Unit.Write_Hit_Percent GE 0
```

Monitors for the condition where the percent of DASD/Cache fast write commands that were successfully processed without accessing the volume is below the critical threshold. If performance is impacted, you might need to move a volume or dataset to another control unit to balance the workload.

KS3_Cachecu_Write_HitP_Warning

```
If VALUE S3_Cache_Control_Unit.Write_Hit_Percent LE 50 AND
VALUE S3_Cache_Control_Unit.Write_Hit_Percent GT 45
```

Monitors for the condition where the percent of DASD/Cache fast write commands that were successfully processed without accessing the volume is below the warning level. If performance is impacted, you might need to move a volume or dataset to another control unit to balance the workload.

KS3_Channel_Busy_Pct_Critical

```
If VALUE S3_Channel_Path.Complex_Percent_Utilized GE 85
```

Monitors high response time for I/O requests to volumes being serviced by the channel due to over utilization of that channel. You might need to balance the workload between channels by moving volumes or datasets.

KS3_Channel_Busy_Pct_Warning

```
If VALUE S3_Channel_Path.Complex_Percent_Utilized GE 70 AND
VALUE S3_Channel_Path.Complex_Percent_Utilized LT 85
```

Monitors high response time for I/O requests to volumes being serviced by the channel due to over utilization of that channel. You might need to balance the workload between channels by moving volumes or datasets.

KS3_FICON_Port_Frame_Pacing_W

Situation is true if the port is not a switch (if port is connected to a CU or channel path) and the average frame pacing is greater than 100.

This situation is written against the S3_FICON_Director_Ports attribute group.

KS3_FICON_Switch_Frame_Pacing_W

Situation is true if the port is not a switch (if port is connected to a CU or channel path) and the average frame pacing is greater than 100.

This situation is written against the S3_FICON_Director_Ports attribute group.

KS3_HSM_Backup_Held_Critical

```
If VALUE S3_HSM_Function_Summary.Function_Status EQ Held AND
VALUE S3_HSM_Function_Summary.Function EQ Backup
```

Monitors the HSM backup function to see if it is being held. If the hold is inadvertent, issue the HSM RELEASE BACKUP command to allow the backup function to continue processing.

KS3_HSM_Backup_Queue_Critical

```
If VALUE S3_HSM_Function_Summary.Waiting_Requests GE 50 AND
VALUE S3_HSM_Function_Summary.Function EQ Backup
```

Monitors the HSM backup queue for a condition where the number of backup requests waiting exceeds the critical threshold. If the number of backup tasks is not at the maximum, issue the HSM SETSYS MAXBACKUPTASKS command to increase the number of backup tasks, thus increasing the processing rate. Keep in mind that the number of available backup volumes serves as a constraint on the number of active backup tasks.

KS3_HSM_Backup_Queue_Warning

```
If VALUE S3_HSM_Function_Summary.Waiting_Requests GE 15 AND
VALUE S3_HSM_Function_Summary.Waiting_Requests LT 50 AND
VALUE S3_HSM_Function_Summary.Function EQ Backup
```

Monitors the HSM backup queue for a condition where the number of backup requests waiting exceeds the warning threshold. If the number of backup tasks is not at the maximum, issue the HSM SETSYS MAXBACKUPTASKS command to increase the number of backup tasks, thus increasing the processing rate. Keep in mind that the number of available backup volumes serves as a constraint on the number of active backup tasks.

KS3_HSM_CRQ_Element_Full_Warn

```
If VALUE S3_HSM_CRQplex.Element_Percent_Full GT 80
```

Monitors the percentage of elements on the Common Recall Queue that are currently in use. HSM throttles the use of the CRQ when the percent used reaches 95%. To expand the CRQ structure, issue the SETXCF START,ALTER command.

KS3_HSM_CRQ_Entry_Full_Warning

```
If VALUE S3_HSM_Cross_System_CRQplex.Entry_Percent_Full GT 80
```

Monitors the percentage of entries on the Common Recall Queue that are currently in use. HSM throttles the use of the CRQ when the percent used reaches 95%. To expand the CRQ structure, issue the SETXCF START,ALTER command.

KS3_HSM_CRQ_Host_Critical

```
If VALUE S3_HSM_Cross_System_CRQ_Hosts.HSM_Host_CRQ_State NE Connected
AND VALUE S3_HSM_Cross_System_CRQ_Hosts.CRQplex_Base_Name NE n/a
```

Monitors the state of the host in regards to the Common Recall Queue. To connect an HSM host to the CRQ, issue the HSM SETSYS command.

KS3_HSM_CRQ_Host_Disconn_Crit

```
If VALUE S3_HSM_Cross_System_CRQplex.HSM_Hosts_Not_Connected GT 0
```

Monitors the number of HSM hosts currently not connected to the Common Recall Queue.

KS3_HSM_CRQ_Host_Held_Critical

```
If VALUE S3_Cross_System_HSM_CRQ_Hosts.Host_CRQ_Held EQ Yes
```

Monitors the commonqueue status for this host. This condition can occur if the HOLD COMMONQUEUE command has been issued. To resolve this condition, issue a RELEASE COMMONQUEUE command.

KS3_HSM_CRQ_Host_Place_Crit

```
If VALUE S3_Cross_System_HSM_CRQ_Hosts.Host_CRQ_Recall_Place_Held EQ Internal OR
VALUE S3_Cross_System_HSM_CRQ_Hosts.Host_CRQ_Recall_Place_Held EQ External OR
VALUE S3_Cross_System_HSM_CRQ_Hosts.Host_CRQ_Recall_Place_Held EQ Both
```

Monitors the commonqueue status for this host and whether requests can be placed on the common recall queue. This condition can occur if the HOLD COMMONQUEUE(RECALL(PLACEMENT)) command has been issued or inferred because a HOLD COMMONQUEUE or HOLD COMMONQUEUE(RECALL) was issued. To resolve this condition, issue a RELEASE COMMONQUEUE(RECALL(PLACEMENT)) command.

KS3_HSM_CRQ_Host_Recall_Crit

```
If VALUE S3_Cross_System_HSM_CRQ_Hosts.Host_CRQ_Recall_Held EQ Internal OR
VALUE S3_Cross_System_HSM_CRQ_Hosts.Host_CRQ_Recall_Held EQ External OR
VALUE S3_Cross_System_HSM_CRQ_Hosts.Host_CRQ_Recall_Held EQ Both
```

Monitors the commonqueue status for this host and whether requests can be recalled from the common recall queue. This condition can occur if the HOLD COMMONQUEUE(RECALL) command has been issued or inferred because a HOLD COMMONQUEUE was issued. To resolve this condition, issue a RELEASE COMMONQUEUE(RECALL) command.

KS3_HSM_CRQ_Host_Select_Crit

```
If VALUE S3_Cross_System_HSM_CRQ_Hosts.Host_CRQ_Recall_Select_Held EQ Internal OR
VALUE S3_Cross_System_HSM_CRQ_Hosts.Host_CRQ_Recall_Select_Held EQ External OR
VALUE S3_Cross_System_HSM_CRQ_Hosts.Host_CRQ_Recall_Select_Held EQ Both
```

Monitors the commonqueue status for this host and whether requests can be pulled from the common recall queue. This condition can occur if the HOLD COMMONQUEUE(RECALL(SELECT)) command has been issued or inferred because a HOLD COMMONQUEUE or HOLD COMMONQUEUE(RECALL) was issued. To resolve this condition, issue a RELEASE COMMONQUEUE(RECALL(SELECT)) command.

KS3_HSM_Dump_Held_Critical

```
If VALUE S3_HSM_Function_Summary.Function_Status EQ Held AND
VALUE S3_HSM_Function_Summary.Function EQ Dump
```

Monitors the HSM dump function to see if it is being held. If the hold is inadvertent, issue the HSM RELEASE DUMP command to allow dump processing to continue.

KS3_HSM_Dump_Queue_Critical

```
If VALUE S3_HSM_Function_Summary.Function_Status EQ Held AND
VALUE S3_HSM_Function_Summary.Function EQ Dump
```

Monitors the HSM dump queue for a condition where the number of dump requests waiting exceeds the critical threshold. If the number of dump tasks is not at the maximum, use the HSM SETSYS MAXDUMPTASKS command to increase the number of dump tasks, thus increasing the processing rate. Keep in mind that the number of available tape drives serves as a constraint on the number of active dump tasks.

KS3_HSM_Dump_Queue_Warning

```
If VALUE S3_HSM_Function_Summary.Waiting_Requests GE 15 AND
VALUE S3_HSM_Function_Summary.Function EQ Dump AND
VALUE S3_HSM_Function_Summary.Waiting_Requests LT 50
```

Monitors the HSM dump queue for a condition where the number of dump requests waiting exceeds the warning threshold. If the number of dump tasks is not at the maximum, use the HSM SETSYS MAXDUMPTASKS command to increase the number of dump tasks, thus increasing the processing rate. Keep in mind that the number of available tape drives serves as a constraint on the number of active dump tasks.

KS3_HSM_Inactive_Host_Warning

```
If VALUE S3_HSM_Status.Inactive_HSM_Hosts GT 0
```

Monitors when an inactive HSM host has been detected. The event workspace for this situation has a link to the DFSMShsm Host Details workspace.

KS3_HSM_Migrate_Held_Critical

```
If VALUE S3_HSM_Function_Summary.Function_Status EQ Held AND
VALUE S3_HSM_Function_Summary.Function EQ Migration
```

Monitors the migrate function to see if it is being held. If the hold on the function is inadvertent, issue the HSM RELEASE MIGRATION command to allow migration to continue.

KS3_HSM_Migrate_Queue_Critical

```
If VALUE S3_HSM_Function_Summary.Waiting_Requests GE 50 AND
VALUE S3_HSM_Function_Summary.Function EQ Migration
```

Monitors the HSM migration queue for a condition where the number of migration requests waiting exceeds the critical threshold. If the number of migrate tasks is not at the maximum, use the HSM SETSYS MAXMIGRATIONTASKS command to increase the number of migration tasks, thus increasing the processing rate. Note that this affects only those migrations requested by automatic functions. Only one task is available to process command migration requests.

KS3_HSM_Migrate_Queue_Warning

```
If VALUE S3_HSM_Function_Summary.Waiting_Requests GE 15 AND
VALUE S3_HSM_Function_Summary.Waiting_Requests LT 50 AND
VALUE S3_HSM_Function_Summary.Function EQ Migration
```

Monitors the HSM migration queue for a condition where the number of migration requests waiting exceeds the warning threshold. If the number of migrate tasks is not at the maximum, use the HSM SETSYS MAXMIGRATIONTASKS command to increase the number of migration tasks, thus increasing the processing rate. Note that this affects only those migrations requested by automatic functions. Only one task is available to process command migration requests.

KS3_HSM_Recall_Held_Critical

```
If VALUE S3_HSM_Function_Summary.Function_Status EQ Held AND
VALUE S3_HSM_Function_Summary.Function EQ Recall
```

Monitors the recall function to see if it is being held. If the hold on the function is inadvertent, issue the HSM RELEASE RECALL command to allow recalls to resume.

KS3_HSM_Recall_Queue_Critical

```
If VALUE S3_HSM_Function_Summary.Waiting_Requests GE 50 AND
VALUE S3_HSM_Function_Summary.Function EQ Recall
```

Monitors the HSM recall queue for a condition where the number of recall requests waiting exceeds the critical threshold. If the number of recall tasks is not at the maximum, use the HSM SETSYS MAXRECAL LTASKS command to increase the number of recall tasks, thus increasing the processing rate.

KS3_HSM_Recall_Queue_Warning

```
If VALUE S3_HSM_Function_Summary.Waiting_Requests GE 15 AND
VALUE S3_HSM_Function_Summary.Waiting_Requests LT 50 AND
VALUE S3_HSM_Function_Summary.Function EQ Recall
```

Monitors the HSM recall queue for a condition where the number of recall requests waiting exceeds the critical threshold. If the number of recall tasks is not at the maximum, use the HSM SETSYS MAXRECAL LTASKS command to increase the number of recall tasks, thus increasing the processing rate.

KS3_HSM_Recovery_Held_Critical

```
If VALUE S3_HSM_Function_Summary.Function_Status EQ Held AND
VALUE S3_HSM_Function_Summary.Function EQ Recovery
```

Monitors the recovery function to see if it is being held. If the hold on the function is inadvertent, issue the HSM RELEASE RECOVER command to allow recovery function to resume.

KS3_HSM_Recovery_Queue_Critical

```
If VALUE S3_HSM_Function_Summary.Waiting_Requests GE 50 AND
VALUE S3_HSM_Function_Summary.Function EQ Recovery
```

Monitors the HSM recovery queue for a condition where the number of recover requests waiting exceeds the critical threshold. If the number of recovery tasks is not at the maximum, use the HSM SETSYS MAXDSRECOVERTASKS command to increase the number of recover tasks, thus increasing the processing rate. Keep in mind that the number of backup tape cartridges serves as a constraint on the number of active recovery tasks.

KS3_HSM_Recovery_Queue_Warning

```
If VALUE S3_HSM_Function_Summary.Waiting_Requests GE 15 AND
VALUE S3_HSM_Function_Summary.Waiting_Requests LT 50 AND
VALUE S3_HSM_Function_Summary.Function EQ Recovery
```

Monitors the HSM recovery queue for a condition where the number of recover tasks waiting exceeds the warning threshold. If the number of recovery tasks is not at the maximum, use the HSM SETSYS MAXDSRECOVERTASKS command to increase the number of recover tasks, thus increasing the processing rate. Keep in mind that the number of backup tape cartridges serves as a constraint on the number of active recovery tasks.

KS3_HSM_Status_Inactive_Crit

```
If VALUE S3_HSM_Status.HSM_Status EQ InActive
```

Monitors the status of the HSM. If status is not active, restart HSM.

KS3_LCU_Av_Delay_Q_Critical

```
If VALUE S3_Logical_Control_Unit.Average_Delay_Queue GE 0.500
```

Monitors for the condition where the average number of requests queued to devices assigned to a logical control unit due to busy conditions on physical paths has exceeded the critical threshold. If performance is impacted, you might be able to balance the workload across multiple LCUs by moving a volume or dataset. Otherwise, you need to add physical paths to the LCU.

KS3_LCU_Av_Delay_Q_Warning

```
If VALUE S3_Logical_Control_Unit.Average_Delay_Queue GE 0.2 AND
VALUE S3_Logical_Control_Unit.Average_Delay_Queue LT 0.500
```

Monitors for the condition where the average number of requests queued to devices assigned to a logical control unit due to busy conditions on physical paths has exceeded the warning threshold. If performance is impacted, you might be able to balance the workload across multiple LCUs by moving a volume or dataset. Otherwise, you need to add physical paths to the LCU.

KS3_LCU_Cont_Rate_Critical

```
If VALUE S3_Logical_Control_Unit.Contention_Rate GE 1.001
```

Monitors for the condition where the rate at which I/O requests are being queued to devices on a logical control unit (LCU) due to busy conditions on physical paths has exceeded the critical threshold. If performance is impacted, you need to migrate volumes or datasets to another LCU, otherwise, you need to add physical paths to the LCU.

KS3_LCU_Cont_Rate_Warning

```
If VALUE S3_Logical_Control_Unit.Contention_Rate GE 0.2 AND
VALUE S3_Logical_Control_Unit.Contention_Rate LT 1.001
```

Monitors for the condition where the rate at which I/O requests are being queued to devices on a logical control unit (LCU) due to busy conditions on physical paths has exceeded the warning threshold. If performance is impacted, you need to migrate volumes or datasets to another LCU, otherwise, you need to add physical paths to the LCU.

KS3_LCU_IO_Rate_Sec_Critical

```
If VALUE S3_Logical_Control_Unit.Channel_Path_I/O_Rate GE 600
```

Monitors for the condition where the I/O rate per second to volumes in the logical control unit (LCU) has exceeded the critical threshold. If performance is impacted, you need to balance the workload across multiple LCUs by moving volumes or datasets.

KS3_LCU_IO_Rate_Sec_Warning

```
If VALUE S3_Logical_Control_Unit.Channel_Path_I/O_Rate GE 200 AND
VALUE S3_Logical_Control_Unit.Channel_Path_I/O_Rate LT 600
```

Monitors for the condition where the I/O rate per second to volumes in the logical control unit (LCU) has exceeded the warning threshold. If performance is impacted, you need to balance the workload across multiple LCUs by moving volumes or datasets.

KS3_RLS_Accelerated_Mode_Warn

```
If VALUE S3_RLS_Buffer_LSU_Summary.BMF_Accelerated_Mode_Pct GT 75
```

Monitors for the warning condition where the BMF Accelerated Mode Pct is greater than 75.

KS3_RLS_Dataset_Avg_Resp_Time_C

```
If VALUE S3_RLS_Dataset_Group_Details.Average_Response_Time GT 3
```

Monitors for the critical condition where the Average Response Time is greater than 3.

KS3_RLS_Dataset_Avg_Resp_Time_W

```
If VALUE S3_RLS_Dataset_Group_Details.Average_Response_Time GT 2 AND
VALUE S3_RLS_Dataset_Group_Details.Average_Response_Time LT 3
```

Monitors for the warning condition where the Average Response Time is between 2 and 3.

KS3_RLS_DSG_Max_Avg_Resp_Time_C

```
If VALUE S3_RLS_Dataset_Group_Summary.Max_Average_Response_Time GT 3
```

Monitors for the critical condition where the Max Average Response Time is greater than 3.

KS3_RLS_DSG_Max_Avg_Resp_Time_W

```
If VALUE S3_RLS_Dataset_Group_Summary.Max_Average_Response_Time GT 2 AND
VALUE S3_RLS_Dataset_Group_Summary.Max_Average_Response_Time LT 3
```

Monitors for the warning condition where the Max Average Response Time is between 2 and 3.

KS3_RLS_Panic_Mode_Critical

```
If VALUE S3_RLS_Buffer_LSU_Summary.BMF_Panic_Mode_Pct GT 75
```

Monitors for the critical condition where the BMF Panic Mode Pct is greater than 75.

KS3_RLS_Panic_Mode_Warning

```
If VALUE S3_RLS_Buffer_LSU_Summary.BMF_Panic_Mode_Pct GT 50 AND
VALUE S3_RLS_Buffer_LSU_Summary.BMF_Panic_Mode_Pct LT 75
```

Monitors for the warning condition where the BMF Panic Mode Pct is between 50 and 75.

KS3_RLS_StorCls_Avg_Resp_Time_C

```
If VALUE S3_RLS_Storage_Class.Average_Response_Time GT 3
```

Monitors for the critical condition where the Average Response Time is greater than 3.

KS3_RLS_StorCls_Avg_Resp_Time_W

```
If VALUE S3_RLS_Storage_Class.Average_Response_Time GT 2 AND
VALUE S3_RLS_Storage_Class.Average_Response_Time LT 3
```

Monitors for the warning condition where the Average Response Time is between 2 and 3.

KS3_RMM_CDS_Backup_Critical

```
If VALUE S3_RMM_Control_Dataset.Days_Since_Last_Backup GT 3
```

The number of days since the last backup of the DFSMSrmm™ CDS or Journal exceeded the critical threshold.

KS3_RMM_CDS_Backup_Warning

```
If VALUE S3_RMM_Control_Dataset.Days_Since_Last_Backup GT 1 AND
VALUE S3_RMM_Control_Dataset.Days_Since_Last_Backup LE 3
```

The number of days since the last backup of the DFSMSrmm™ CDS or Journal exceeded the warning threshold.

KS3_RMM_CDS_Space_Critical

```
If VALUE S3_RMM_Control_Dataset.RMM_Percent_Used GT 90
```

The percentage of space used by the DFSMSrmm™ CDS or Journal is greater than the critical threshold.

KS3_RMM_CDS_Space_Warning

```
If VALUE S3_RMM_Control_Dataset.RMM_Percent_Used GE 80 AND
VALUE S3_RMM_Control_Dataset.RMM_Percent_Used LE 90
```

The percentage of space used by the DFSMSrmm™ CDS or Journal is greater than the warning threshold.

KS3_RMM_Exit_Status_Critical

```
If ( ( VALUE S3_RMM_Config.EDGUX200_Status NE Enabled ) OR
( VALUE S3_RMM_Config.EDGUX100_Status NE Enabled ) )
```

The DFSMSrmm™ EDGUX100 or EDGUX200 exit is not Enabled.

KS3_RMM_Journal_Status_Critical

```
If VALUE S3_RMM_Config.Journal_Status NE Enabled
```

The DFSMSrmm™ Journal is either Disabled or Locked. DFSMSrmm™ does not allow further updates to the journal until BACKUP is run to back up the DFSMSrmm™ control dataset and to clear the journal. If the Journal is Locked, DFSMSrmm™ fails any requests that result in an update to the DFSMSrmm™ control dataset. Message EDG2103D might also have been issued to the DFSMSrmm™ operator console.

KS3_RMM_Operating_Mode_Warning

```
If VALUE S3_RMM_Config.Operating_Mode NE Protect
```

DFSMSrmm is not operating in Protect mode. Certain actions that should be rejected are permitted if DFSMSrmm™ is not operating in protect mode, for example attempting to read a scratch tape volume.

KS3_RMM_Scratch_Tape_Critical

```
If VALUE S3_RMM_Summary.Type EQ 0 AND
VALUE S3_RMM_Summary.Scratch_Volumes LT 100
```

The number of Scratch volumes is below the critical threshold.

KS3_RMM_Scratch_Tape_Warning

```
If VALUE S3_RMM_Summary.Type EQ 0 AND
VALUE S3_RMM_Summary.Scratch_Volumes LT 200 AND
VALUE S3_RMM_Summary.Scratch_Volumes GE 100
```

The number of Scratch volumes is below the warning threshold.

KS3_RMM_Inactive_Critical

```
If VALUE S3_RMM_Config.Subsystem_Status EQ Inactive
```

The DFSMSrmm™ subsystem is inactive.

KS3_Stg_Toolkit_Result_Critical

```
If VALUE S3_Storage_Toolkit_Result_Summary.Return_Code GT 4
```

The batch job submitted by the Storage Toolkit to execute a command or user-defined JCL returns a value greater than 4. Or the Storage Toolkit encountered an error while attempting to process a command or user-defined JCL. A value that is greater than 4, and is not specific to the Storage Toolkit, typically denotes that a command failed to complete. If you elected to save the results of the batch job, go to the Storage Toolkit Result Detail workspace to determine whether the error requires further attention.

Values set by the Storage Toolkit when it detects an error while processing a command or user-defined JCL are described in “Using the Storage Toolkit” on page 146 (see the description of return codes). If one of these values is returned, also consult the RKLVLOG for additional messages to help determine the cause of the failure.

KS3_Stg_Toolkit_Result_Warning

```
If VALUE S3_Storage_Toolkit_Result_Summary.Return_Code EQ 4
```

The batch job submitted by the Storage Toolkit to execute a command or user-defined JCL returns the value 4. A value of 4 typically denotes a warning. If you elected to save the results of the batch job, go to the Storage Toolkit Result Detail workspace to determine whether the warning requires further attention.

KS3_Storage_Gr_Pct_Free_Crit

```
If VALUE S3_Volume_Group_Summary.Free_Space_Percent LT 5.0 AND
VALUE S3_Volume_Group_Summary.Group_Type EQ SMSGROUP AND
VALUE S3_Vol ume_Group_Summary.Free_Space_Percent GE 0.0
```

Monitors the percentage of free space available for allocation in the storage group and detects when free space has dropped below the critical threshold. To prevent allocation failures, you might have to either add one or more logical volumes to the storage group, or to move datasets off of the logical volumes in the storage group.

KS3_Storage_Gr_Pct_Free_Warning

```
If VALUE S3_Volume_Group_Summary.Free_Space_Percent LT 10.0 AND
VALUE S3_Volume_Group_Summary.Group_Type EQ SMSGROUP AND
VALUE S3_Volume_Group_Summary.Free_Space_Percent GE 5.0",
```

Monitors the percentage of free space available for allocation in the storage group and detects when free space has dropped below the warning threshold. In order to prevent allocation failures, you might have to either add one or more logical volumes to the storage group, or to migrate datasets off of the logical volumes in the storage group.

KS3_TDS_Array_Degraded_Crit

```
If VALUE S3_TotalStorageDS_Array.RAID_Degraded EQ Yes
```

Monitors the arrays in a TotalStorageDS storage facility for a degraded condition where one or more arrays need rebuilding.

KS3_TDS_Array_Prob_Crit

```
If VALUE S3_TotalStorageDS_Configuration.Number_of_arrays_with_problems GT 0
```

Monitors for the condition where the number of arrays in the TotalStorageDS storage facility running degraded, throttled, or with an RPM exception exceeds the threshold. The RAID Degraded condition indicates that one or more DDMs in the array need rebuilding. The DDM Throttling condition indicates that a near-line DDM in the array is throttling performance due to temperature or workload. The RPM Exception condition indicates that a DDM with a slower RPM than the normal array DDMs is a member of the array as a result of a sparing action.

KS3_TDS_Array_RPM_Crit

```
If VALUE S3_TotalStorageDS_Array.RPM_Exception EQ Yes
```

Monitors the arrays in a TotalStorageDS for a condition where a DDM with a slower RPM than the normal array DDMs is a member of the array as a result of a sparing action.

KS3_TDS_Array_Throttled_Crit

```
If VALUE S3_TotalStorageDS_Array.DDM_Throttling EQ Yes
```

Monitors the arrays in a TotalStorageDS for a condition where the array is throttling performance due to overload or temperature.

KS3_TDS_ExtPool_Array_Prob_Crit

```
If VALUE S3_TotalStorageDS_Extent_Pool.Number_of_arrays_with_problems GT 0
```

Monitors for the condition where the number of arrays in the extent pool running degraded, throttled, or with an RPM exception exceeds the threshold. The RAID Degraded condition indicates that one or more DDMs in the array need rebuilding. The DDM Throttling condition indicates that a near-line DDM in the array is throttling performance due to temperature or workload. The RPM Exception condition indicates that a DDM with a slower RPM than the normal array DDMs is a member of the array as a result of a sparing action.

KS3_TDS_Rank_Array_Prob_Crit

```
If VALUE S3_TotalStorageDS_Rank.Number_of_arrays_with_problems GT 0
```

Monitors for the condition where the number of arrays in the rank running degraded, throttled, or with an RPM exception exceeds the threshold. The RAID Degraded condition indicates that one or more DDMs in the array need rebuilding. The DDM Throttling condition indicates that a near-line DDM in the array is throttling performance due to temperature or workload. The RPM Exception condition indicates that a DDM with a slower RPM than the normal array DDMs is a member of the array as a result of a sparing action.

KS3_Vol_Cache_DFW_Retry_Critical

```
If VALUE S3_Cache_Devices.DFW_Retry_Percent GE 2 AND
VALUE S3_Cache_Devices.I/O_Count GE 25
```

Monitors for the condition where the percentage of DASD fast write attempts for a volume that cannot be satisfied due to a shortage of available nonvolatile storage (NVS) space exceeded the critical threshold. Check for pinned NVS and correct the problem if NVS is pinned. Otherwise, if the impact on performance is not acceptable, move a volume or dataset to another cache control unit or add NVS to this control unit.

KS3_Vol_Cache_DFW_Retry_Warning

```
If VALUE S3_Cache_Devices.DFW_Retry_Percent GE 1 AND
VALUE S3_Cache_Devices.DFW_Retry_Percent LT 2 AND
VALUE S3_Cache_Devices.I/O_Count GE 25
```

Monitors for the condition where the percentage of DASD fast write attempts for a volume that cannot be satisfied due to a shortage of available nonvolatile storage (NVS) space exceeded the warning threshold. Check for pinned NVS and correct the problem if NVS is pinned. Otherwise, if the impact on performance is not acceptable, move a volume or dataset to another cache control unit or add NVS to this control unit.

KS3_Vol_Cache_Read_HitP_Critical

```
If VALUE S3_Cache_Devices.Read_Hit_Percent LE 45 AND
VALUE S3_Cache_Devices.Read_Hit_Percent GE 0 AND
VALUE S3_Cache_Devices.I/O_Count GE 25
```

Monitors for the condition where the cache read hit percent is below the critical threshold. If performance is impacted, determine the reason for the low read hit percent. Common problems are cache-unfriendly applications and over-utilization of the control unit.

KS3_Vol_Cache_Read_HitP_Warning

```
If VALUE S3_Cache_Devices.Read_Hit_Percent LE 55 AND
VALUE S3_Cache_Devices.Read_Hit_Percent GT 45 AND
VALUE S3_Cache_Devices.I/O_Count GE 25
```

Monitors for the condition where the cache read hit percent is below the warning threshold. If performance is impacted, determine the reason for the low read hit percent. Common problems are cache-unfriendly applications and over-utilization of the control unit.

KS3_Vol_Cache_Writ_HitP_Critical

```
If VALUE S3_Cache_Devices.Write_Hit_Percent LE 20 AND
VALUE S3_Cache_Devices.Write_Hit_Percent GE 0 AND
VALUE S3_Cache_Devices.I/O_Count GE 25
```

Monitors for the condition where the cache write hit percent for a volume is below the critical threshold. Check the status of the nonvolatile storage in the cache control unit. You can move volumes or datasets to balance the workload.

KS3_Vol_Cache_Writ_HitP_Warning

```
If VALUE S3_Cache_Devices.Write_Hit_Percent LE 30 AND
VALUE S3_Cache_Devices.Write_Hit_Percent GT 20 AND
VALUE S3_Cache_Devices.I/O_Count GE 25
```

Monitors for the condition where the cache write hit percent for a volume is below the warning threshold. Check the status of the nonvolatile storage in the cache control unit. You can move volumes or datasets to balance the workload.

KS3_Vol_Disabled_VTOC_Critical

```
If VALUE S3_DASD_Volume_Space.VTOC_Index_Status EQ Disabled
```

Monitors for the condition where a VTOC index has been disabled. This condition can degrade performance on the volume. Enable the VTOC index.

KS3_Vol_EAV_Fragment_Index_Crit

```
If VALUE S3_DASD_Volume_Space.Extended_Address_Volume EQ Yes AND
VALUE S3_DASD_Volume_Space.Track_Managed_Fragmentation_Index GE 850
```

The fragmentation index in the track managed area of an Extended Address Volume exceeds the critical threshold.

KS3_Vol_EAV_Fragment_Index_Warn

```
If VALUE S3_DASD_Volume_Space.Extended_Address_Volume EQ Yes AND
VALUE S3_DASD_Volume_Space.Track_Managed_Fragmentation_Index GE 650
AND VALUE S3_DASD_Volume_Space.Track_Managed_Fragmentation_Index LT 850
```

The fragmentation index in the track managed area of an Extended Address Volume exceeds the warning threshold.

KS3_Vol_EAV_Free_Space_Pct_Crit

```
If VALUE S3_DASD_Volume_Space.Track_Managed_Percent_Free LE 5.0 AND
VALUE S3_DASD_Volume_Space.Track_Managed_Percent_Free GE 0.0 AND
VALUE S3_DASD_Volume_Space.Extended_Address_Volume EQ Yes
```

The percentage of free space in the track managed area of an Extended Address Volume is below the critical threshold.

KS3_Vol_EAV_Free_Space_Pct_Warn

```
If VALUE S3_DASD_Volume_Space.Track_Managed_Percent_Free LE 10.0 AND
VALUE S3_DASD_Volume_Space.Track_Managed_Percent_Free GT 5.0 AND
VALUE S3_DASD_Volume_Space.Extended_Address_Volume EQ Yes
```

The percentage of free space in the track managed area of an Extended Address Volume is below the warning threshold.

KS3_Vol_Fragment_Index_Critical

```
If VALUE S3_DASD_Volume_Space.Fragmentation_Index GE 850
```

Monitors for the condition where a volume has a fragmentation index that exceeds the critical threshold. Defragment the volume so that free extents are combined to help prevent dataset allocation failures.

KS3_Vol_Fragment_Index_Warning

```
If VALUE S3_DASD_Volume_Space.Fragmentation_Index GE 650 AND
VALUE S3_DASD_Volume_Space.Fragmentation_Index LT 850
```

Monitors for the condition where a volume has a fragmentation index that exceeds the warning threshold. Defragment the volume so that free extents are combined to help prevent dataset allocation failures.

KS3_Vol_Free_Space_Pct_Critical

```
If VALUE S3_DASD_Volume_Space.Percent_Free_Space LE 5 AND
VALUE S3_DASD_Volume_Space.Percent_Free_Space GE 0
```

Monitors for the condition where the percentage of free space on a volume is below the critical threshold. If datasets on the volume require more space, then either migrate some datasets to another volume or release space from datasets that might be over-allocated.

KS3_Vol_Free_Space_Pct_Warning

```
If VALUE S3_DASD_Volume_Space.Percent_Free_Space LE 10 AND
VALUE S3_DASD_Volume_Space.Percent_Free_Space GT 5
```

Monitors for the condition where the percentage of free space on a volume is below the critical threshold. If datasets on the volume require more space, then either migrate some datasets to another volume or release space from datasets that might be over-allocated.

KS3_Vol_Perf_Resp_Time_Critical

```
If VALUE S3_DASD_Volume_Performance.Response_Time GE 55 AND
VALUE S3_DASD_Volume_Performance.I/O_Count GE 25
```

Monitors for the condition where response time for the volume exceeds the critical threshold. Look at the volume to see if high utilization is a problem. If so, it might be necessary to migrate datasets from the volume to reduce utilization. Also check the cache status of the volume. Look at the components of I/O to determine where the time is being spent and address the problem accordingly.

KS3_Vol_Perf_Resp_Time_Warning

```
If VALUE S3_DASD_Volume_Performance.Response_Time GE 35 AND
VALUE S3_DASD_Volume_Performance.Response_Time LT 55 AND
VALUE S3_DASD_Volume_Performance.I/O_Count GE 25
```

Monitors for the condition where response time for the volume exceeds the warning threshold. Look at the volume to see whether high utilization is a problem. If so, you can migrate datasets from the volume to reduce utilization. Also check the cache status of the volume. Look at the components of I/O to determine where the time is being spent and address the problem accordingly.

KS3_VTS_Disconnect_Time_Crit

```
If VALUE S3_VTS_Overview.Virtual_Disconnect_Time GE 500
```

Monitors for the condition where the logical control unit disconnect time for the virtual tape server exceeds the critical threshold. This condition is often an indication that the tape volume cache capacity is being exceeded.

KS3_VTS_Host_GB_Warning

```
If VALUE S3_VTS_Overview.Host_Channel_Activity_GB GE 18
```

Monitors for the condition where the activity between the MVS™ system and the virtual tape server on the host channels exceeds 19 GB over the hour interval. This condition can be an indication that the virtual tape server is being overloaded.

KS3_VTS_Pct_Copy_Throt_Warn

```
If VTSTPVOLC.PCTCPT GT 50
```

Monitors for the condition where copy is the predominant reason for throttling.

KS3_VTS_Pct_Wr_Over_Throt_Warn

```
If VTSTPVOLC.PCTWROT GT 50
```

Monitors for the condition where write overrun is the predominant reason for throttling.

KS3_VTS_Recall_Pct_Warning

```
If VALUE S3_VTS_Overview.Volume_Recall_Percent GE 20
```

Monitors for the condition where the percent of virtual tape mounts that required a physical tape mount to be satisfied exceeded the warning threshold. This condition can lead to unacceptably large virtual mount times. If so, then investigate the reason for the recalls. If rescheduling or removing the application workload is not possible, you need to increase the cache capacity of the VTS.

KS3_VTS_Virt_MtPend_Av_Warning

```
If VALUE S3_VTS_Overview.Average_Virtual_Mount_Pend_Time GE 300
```

Monitors for the condition where the average seconds required to satisfy a virtual mount in the virtual tape subsystem exceeded the warning threshold. If this condition persists, then further study is required to determine the cause for the elongated mount times. The condition might be due to VTS-hostile applications or to a shortage of VTS resources.

KS3_VTS_Virt_MtPend_Mx_Warning

```
If VALUE S3_VTS_Overview.Maximum_Virtual_Mount_Pend_Time EQ 900
```

Monitors for the condition where the maximum seconds required to satisfy a virtual mount in the virtual tape subsystem exceeded the warning threshold. If this condition persists, then further study is required to determine the cause for the elongated mount times. The condition might be due to VTS-hostile applications or to a shortage of VTS resources.

Take action commands

Take Action commands

Learn how to use the Take Action feature to issue commands to any system in your network where one or more monitoring agents are installed.

PREREQUISITE

You can create a Take Action command and run it immediately from a workspace, or you can run a saved Take Action command.
The Tivoli® Enterprise Portal is always hosted on a distributed computer system. When you submit a Take Action command from the Tivoli® Enterprise Portal to a z/OS® system, a return code of zero is displayed in the portal interface. This code indicates successful submission of the command but gives no indication of the result of the command. You can find the command output in the z/OS® SYSLOG.

> **Note:** "Storage Toolkit" on page 143 provides functionality similar to Take Action commands. You create action requests in Storage Toolkit that issue commands in target z/OS® systems.

Here is an example of using a saved Take Action command:

1. Right-click an item in the Navigator.
2. From the pop-up menu, select **Take Action** to display the Take Action dialog box.
3. In the Take Action dialog box, **Action** area, select a saved Take Action command. To edit the command argument values, select **Arguments** and click **OK**.
4. In the Take Action dialog box, **Destination Systems** area, select the target system or systems to which you want to send the Take Action command and click **OK**.
   This causes the command to be sent to the system or systems where it is run.

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.)

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.

Published date: --

© Copyright International Business Machines Corporation 2000, 2024
US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp