



IBM Software Group

Best Practices for System z Storage

Mona Hogue, Sr. Product Marketing Manager – Mainstar Software Corp.

Marti Lombardi, IBM Tivoli Training and Education – IBM Corp.

Tivoli software

A decorative horizontal bar at the bottom of the slide, composed of several colored segments: red, purple, cyan, green, yellow, a white starburst on a red background, a woman's face, a red and white pixelated pattern, a red and white circular pattern, a red and white abstract pattern, and a red and white abstract pattern.

ON DEMAND BUSINESS™

Agenda

- Business Resiliency and Storage
- Data Classification
- Choosing Adequate Data Protection
- Best Practices for Assured Recovery
- Proving Data Validity for Compliance and Audits



Business Resiliency and Storage

- What Is Business Resiliency?
 - ▶ Encarta Dictionary
 - "Speedy recovery from problems"
 - "The ability to recover quickly from setbacks"



"The ability to not fail or to quickly recover in the event of a failure"

IBM Redpaper: IT Resiliency by Steve Finnes, Bob Gintowt, Mike Snyder, and Nick Harris

More than 50 percent of the companies surveyed by Deloitte Touche admitted to having a data loss within the last 12 months, with roughly 33 percent of those incidents directly resulting in financial losses.

Source: z/Flash e-Newsletter July 11, 2006



Business Resiliency and Storage

- Be prepared
 - ▶ The benefits of business resiliency go far beyond preparing for a disaster
 - A resilient company can also withstand errors and local problems that normally cause most outages
 - Many local problems start with poor performance
 - If not handled, can cause a more serious issue
 - ▶ Prevention, planning, and automation are crucial to achieving business resiliency



Data Classification

- Classification of data by importance of recovery requirements
 - ▶ All data is not created equal
 - Based on importance to the business as well as RTO, RPO
 - Business critical
 - Very important
 - Medium importance
 - Low importance
 - ▶ This step is the most time consuming
 - ▶ No software or automation to assist




Choosing Adequate Data Protection

- Type of protection needed for each data class
 - ▶ Determined by what type of outages need to be protected against
 - Local and Disaster
 - Onsite and Offsite
 - ▶ Outages can be grouped into three categories
 - **M**other Nature
 - Floods; hurricanes; earthquakes; tornadoes
 - **M**aliciousness
 - Intentional tampering
 - **M**ishap
 - Inadvertent human errors, fat finger, programming error
 - Data loss or corruption
 - Hardware or software problems
 - Loss of access to data – metadata
 - Loss of power



Choosing Adequate Data Protection

- Choices for data protection
 - ▶ Storage hardware redundancy
 - Raid 1, Raid 5
 - ▶ Backups
 - Tape or disk
 - ▶ Data Replication
 - Storage hardware mirroring
 - Storage hardware “snap” or “pointers”
 - Point-In-Time Copies
 - Storage hardware or software based
 - ▶ Performance monitoring
 - Business critical applications



Best Practices for Assured Recovery

- Challenges for protection of ICF Catalog meta data
 - ▶ Catalogs are critical structures in z/OS
 - All current and migrated data is cataloged
 - Cataloged data is not accessible if the catalog is not available
 - ▶ Large percentage of data in one catalog
 - Single point of failure
 - Commercial Bank - 241,000 data sets, 22 catalogs, largest catalog has 89% of all data sets
 - ▶ If a catalog becomes out of service, business critical data will also be unavailable until the catalog is recovered
- Mainstar CR+ is needed
 - ▶ Outage prevention is key
 - Regularly examines and diagnoses
 - Resolves known problems
 - Backs up frequently
 - Reorganizes when needed



Best Practices for Assured Recovery

- Proactive catalog management with Mainstar CR+
 - ▶ Powerful features for safe, reliable, and easy ICF catalog backup and recovery
 - Necessary for day-to-day ICF catalog management
 - ▶ Protects a catalog's complex structural integrity
 - ▶ Provides alerts about potential errors
 - ▶ Reduces recovery time to minutes from hours or days
 - ▶ Provides the fastest and most accurate catalog and DASD data set synchronization
 - ▶ Diagnostic commands for preventative maintenance
 - Automatic creation of fix commands
 - ▶ Reorganize and repair while open
 - ▶ Forward recovery using SMF data



Best Practices for Assured Recovery

- Critical file identification challenges
 - ▶ If the data isn't identified and protected, it can't be recovered
 - ▶ Dozens or hundreds of applications to protect
 - ▶ Keeping the list of critical files current over time
 - ▶ Be able to recover critical data anywhere, anytime
- Mainstar ASAP is needed
 - ▶ Improves the accuracy of critical file identification
 - ▶ Simplifies the process of critical file identification
 - ▶ Validates that critical data has a backup
 - ▶ Eliminates missing critical data sets locally and at the DR site



Best Practices for Assured Recovery

- **Automated critical file identification with Mainstar ASAP**
 - ▶ Tracks jobs as they are executing
 - Batch, Started task, Submitted by a user ID
 - ▶ Collects SMF data in real-time
 - ▶ Collects JCL information about data sets after the job is submitted to the internal reader
 - ▶ Identifies data by data set name
 - Re-creates the list of critical data each time the jobs execute
 - Keeps it up to date with changes in the application
 - Daily, weekly, monthly, quarterly, annual, and so forth
 - Automatically creates a list of the critical files
 - Selection data set format for ABARS
 - DFDSS or FDR format
 - Sequential list for reporting
 - Compare critical lists towards tracked backups
 - Report on data sets without a backup
 - Prove that critical data has a backup



Best Practices for Assured Recovery

- Challenges when using Aggregate Backup & Recovery Support (ABARS) for backups
 - ▶ Intensely manual processes for most backup and restore activities
- Mainstar ABARS Manager is needed
 - ▶ Provides an inventory of all ABARS backups
 - ▶ Facilitates recovery of aggregates
 - Single or selective data set recovery
 - ▶ Has recovery features to prevent back-leveling more current data
 - ▶ Reduces redundant backups
 - ▶ Provides recovery automation for daily, weekly, monthly, quarterly, and annual processing cycles
 - ▶ ISPF interface for ease-of-use



Best Practices for Assured Recovery

- Automate use of ABARS with Mainstar ABARS Manager
 - ▶ Critical file list generated (each execution) from Mainstar ASAP is fed into ABARS Manager
 - ▶ ABARS Manager backs up all of the data sets in the list
 - All disk, tape and migrated data
 - At the same time (typically at the end of the application) so the backup is synchronized
 - Stacks all of the data sets onto a single set of removable media
 - No JCL to code/maintain
 - Recovery is initiated via ABARS Manager panels
 - Single or selected group
 - All of the data sets



Best Practices for Assured Recovery

- Challenges for managing backups
 - ▶ Backups are performed by many different departments
 - ▶ How to easily identify 'what is and what isn't being backed up'
 - ▶ How to find all available backups when needed
 - ▶ How to avoid backing up data sets multiple times
 - ▶ No easy way to ensure you are restoring from the most current backup
- Mainstar All/Star is needed
 - ▶ Successful recovery anywhere, anytime
 - From minor to catastrophic
 - ▶ Know what was backed up and when
 - Even when backed up by a different department
 - ▶ Be able to easily find the data needed for recovery
 - Search on any data set without knowledge of the backup



Best Practices for Assured Recovery

- Centralized backup management with Mainstar All/Star
 - ▶ Provides an inventory of all backups
 - Centralized inventory provides visibility to all backups
 - ▶ Compares critical file lists towards backup inventory
 - Reports on critical data with backup
 - Reports on critical data not backed up
 - Reports on data backed up redundantly
 - ▶ Recovery locally or for offsite recovery
 - ▶ Search and restore by data set name or application
 - ▶ Automatically builds recovery JCL



Best Practices for Assured Recovery

- DFSSMshsm (HSM) managed data challenges
 - ▶ Customers utilize HSM to manage millions of data sets
 - ▶ Assurance that HSM managed data is available
 - ▶ Validate data with retention requirements for industry or government regulations is properly stored
 - ▶ That data is readily available in the event of a disaster
 - ▶ Assurance that HSM is healthy and dependable
- Mainstar HSM *FastAudit* is needed
 - ▶ Fast and easy way to audit and resolve HSM control data sets
 - ▶ Helps customers comply with government and industry regulations
 - Data retention
 - Data availability
 - Resiliency to outages
 - ▶ Reduces the risk of data loss
 - Mitigates costly financial exposures



Best Practices for Assured Recovery

- Simplify management of HSM with Mainstar HSM *FastAudit*
 - ▶ Automation
 - Speed and flexibility
 - Corrective actions
 - ▶ All inclusive
 - Executes multiple audits in a single execution
 - Discovers and corrects all errors
 - ▶ Proactive
 - Alerts to critical problems
 - Audits are fast and non-intrusive; can execute anytime



Proving Data Validity for Compliance and Audits

- CR+
 - ▶ Prove data is properly cataloged and readily available; catalogs are structurally sound and resilient
- ASAP
 - ▶ Prove that critical data is accurately identified; automation eliminates manual tasks prone to error
- ABARS Manager
 - ▶ Prove critical application data is backed up, regardless of the storage media
- All/Star
 - ▶ Prove critical data is backed up, regardless of the backup method
- HSM *FastAudit*
 - ▶ Prove HSM is healthy and error free





IBM Software Group

Top Storage Issues That OMEGAMON XE for Storage on z/OS Can Help You Solve

Marti Lombardi

Marti.lombardi@us.ibm.com

Tivoli software



Topics

- Product Introduction and Features
- Top Storage Issues solved by IBM Tivoli OMEGAMON XE for Storage on z/OS V3.1.0
- Question and Answer Session



IBM Tivoli OMEGAMON XE for Storage on z/OS V3.1.0

Introduction

- IBM Tivoli OMEGAMON XE for Storage on z/OS is a mainframe (host systems) storage monitor with both real-time and historical analysis.
- The product provides a wide breadth of mainframe storage information.
- The extended edition (XE) user interface comes packaged with the Common User Access (CUA) user interface (3270) component.
- Data in OMEGAMON XE for Storage includes:
 - ▶ Space (storage groups or user groups, which you create)
 - ▶ Performance (storage groups or user groups, which you create)
 - ▶ Tape / Virtual Tape System (VTS)
 - ▶ Cache
 - ▶ Channels using fiber connection (FICON)
 - ▶ Control Units (CUs)
 - ▶ Hierarchical Storage Manager (HSM), so that you can issue commands through the product.
 - ▶ SMS (System Managed Storage) constructs
 - ▶ Ability to see all logical volumes on a physical disk
 - ▶ Powerful applications view

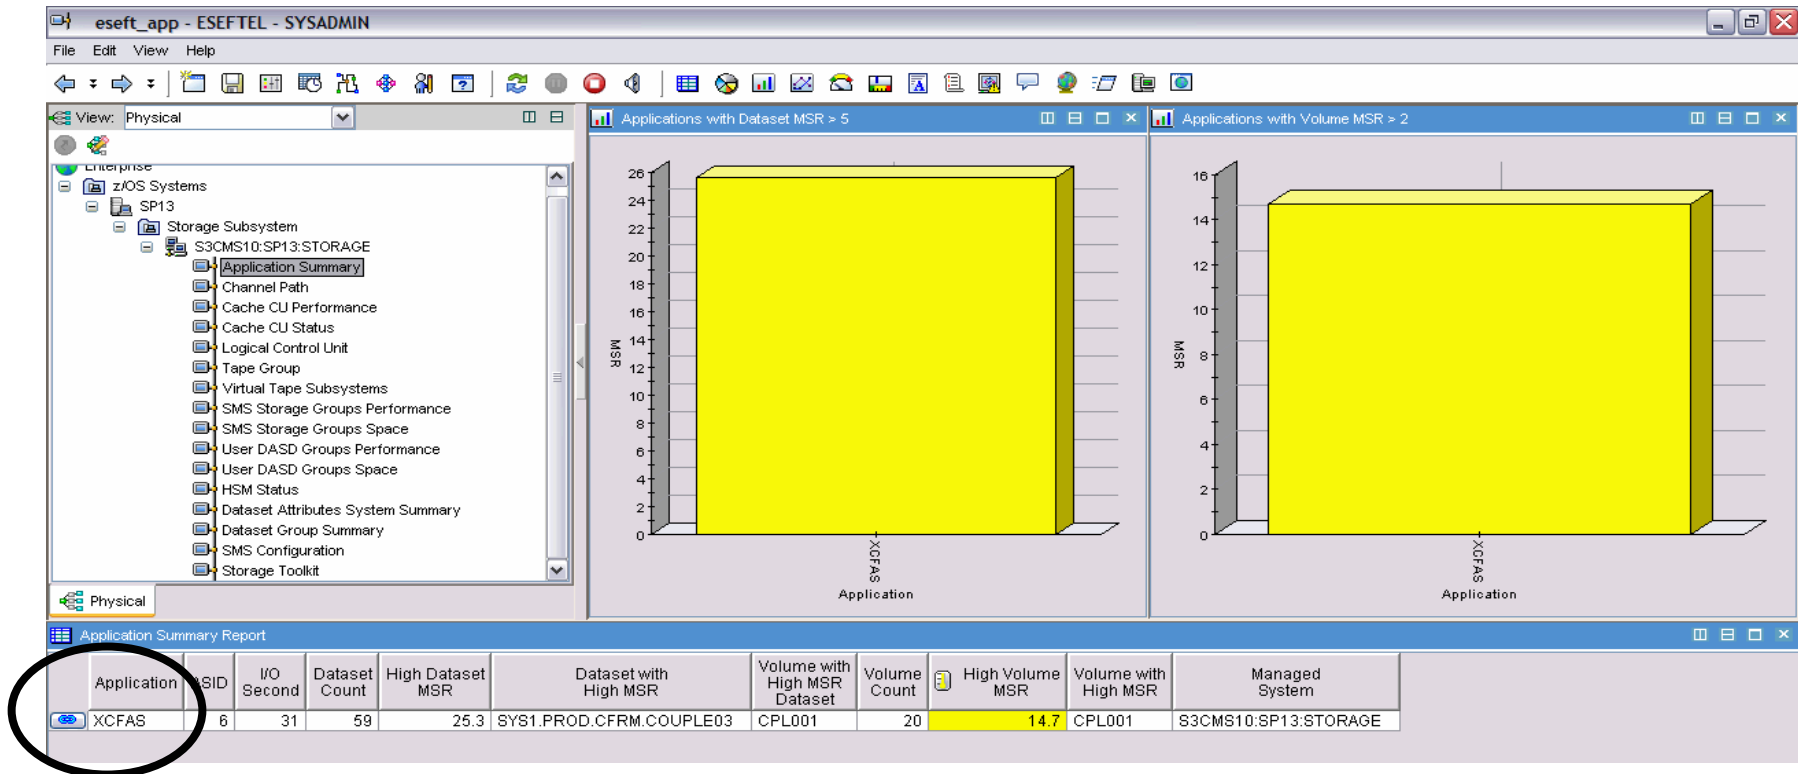


Issue #1: Slow Application Response

- You may monitor key applications for storage issues.
- You may select a report (workspace) that provides monitoring at the application level.
- The application level in IBM Tivoli OMEGAMON XE for Storage on z/OS V3.1.0 is provided by analysis of specific z/OS address spaces.
- You may define which of the applications (address spaces) you select for monitoring.
- A product-provided situation is available to help you get started defining address spaces for monitoring.

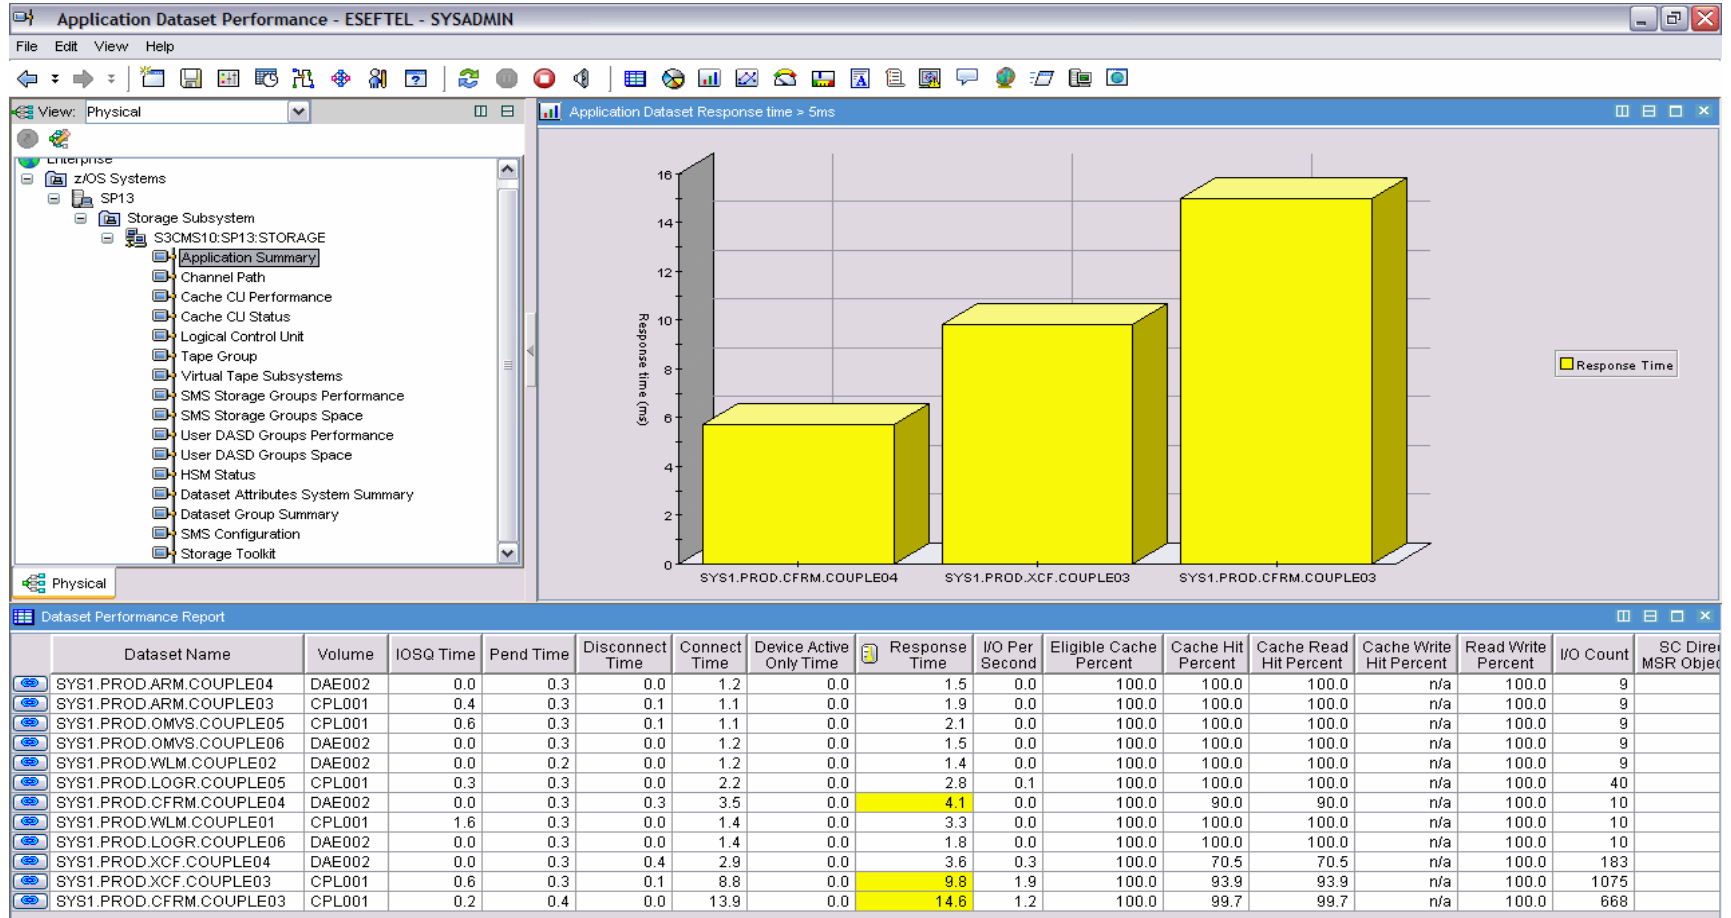


Issue #1: Slow Application Response



In this example, the application being monitoring is XCF (Cross Coupling Facility). The data set with the high response time is the Couple data set, which is a key resource in any environment.

Issue #1: Slow Application Response



Drilling down, you can see all the data sets exceeding specified application-specific thresholds, and therefore impacting critical business application performance.

Issue #1: Slow Application Response

The screenshot displays the 'Situation Editor' window. On the left, a tree view shows a list of R3 instances under 'Storage Subsystem', including various cache and HSM-related metrics. The main area is divided into several sections:

- Description:** An application is accessing a dataset with high response time.
- Formula:** A table with 3 rows and 2 columns. The first row is labeled 'High Dataset MSR' and contains the condition ≥ 50.0 . The second and third rows are empty.
- Situation Formula Capacity:** A progress bar showing 6% usage.
- Sampling interval:** A time selection interface with fields for ddd, hh, mm, and ss, and a checked 'Run at startup' option.

An inset window titled 'Text or Advice Location' is overlaid on the bottom right, showing the following text:

An application is accessing a dataset and the response time is greater than the critical threshold. Determine the reason for the poor response time by looking at the response time components for the dataset. Also examine the volume for over utilization, cache settings, and the response time components at the volume level.

Issue #2: HSM Status, HSM not running or queues are growing

- HSM should be active, where needed.
- Requests waiting for HSM can impact business applications.
- A product-provided situation is available to help you get started defining HSM for monitoring.
- Select the *HSM Status* workspace from the navigator, then link to the *HSM Function Statistics* workspace for details on the name of the application being affected.



Issue #3: RAID Rank Overload

- The DASD farm contains Enterprise Storage Server (Shark) devices, where logical volume space is defined in the same Raid rank.
- The z/OS operating system does not see any activity occurring in the storage facility for Shark DASD.
- The OMEGAMON XE for Storage on z/OS workspaces provide the ability to determine if a logical volume is impacting other logical volumes in the physical group.
- OMEGAMON XE for Storage on z/OS indicates if there is a significant amount of activity, including writes and copies, and if an overload is occurring.
- Rank metrics and logical volume metrics are available.
- First select the *Rand Rank Report*, then select *Raid Rank Trend*, to view all the logical volumes defined on a particular rank. In this manner you may identify the logical volume that may be impacting performance.



Issue #4: DS6000 and DS8000 Extent Overload

- TotalStorage DS6000 (control unit type 1750) and DS8000 (control unit type 2107) are the successor devices to the Shark (2105).
- The Cache CU Performance report has a column that indicates the storage facility type, and will show device types if appropriate.
- The DS6000/8000 physical architecture is comprised of a Storage Facility/Extent Pool/Rank/Array. OMEGAMON allows monitoring to identify problems or bottlenecks at any and all of these levels.
- Extent pools in the storage facility are different than extents associated with a dataset.
- You may navigate through the TotalStorage workspaces to identify where and why an extent pool overload may be occurring.
- Select the *Cache CU Performance* workspace, followed by the *TotalStorage Configuration* workspace to find read and write response times, both totals and averages, for extent pools.



Issue #5: Volume slow response due to cross system contention

- A specific volume is encountering slow response.
- You want to drill down on this volume to look at specific information.
- The volume shows cross system contention impacts.
- Select the *SMS Storage Group Performance* workspace, followed by the *Cross System Group Summary*, to get to the *Cross System Volume Performance* workspace.



Issue #6: Volume Slow Response Due To Data Set Contention

- A specific volume is encountering slow response.
- You want to drill down on this volume to look at specific information.
- The volume shows data set contention impacts from multiple applications running simultaneously.
- Select the *User DASD Groups Performance* workspace, followed by the *User DASD Group Volume Performance*, and finally the *Dataset Performance Summary* and *Dataset Performance Details* workspaces to see the impact of a dataset from one or more applications.



Issue #7: Storage group low on space

- Currently the business applications are operating within expected and required service levels.
- A single storage group, while performing well at this time, is close to running out of storage space.
- The business application could abruptly stop working altogether.
- Options include writing a situation to check the percentage of storage group space available (not in use).




Issue #7: Storage group low on space

The screenshot shows the Tivoli Enterprise Portal interface. The top navigation bar includes 'Welcome DNET241', 'Tivoli Enterprise Portal', and a 'Log out' button. Below the navigation bar is a menu with 'File', 'Edit', 'View', and 'Help'. The main content area is divided into several panes:

- Left Pane:** A tree view showing various reports such as 'Logical Control Unit Report', 'Tape Group Report', 'Virtual Tape Subsystems Report', 'SMS Storage Groups Performance Report', 'SMS Storage Groups Space Report' (highlighted), 'User DASD Groups Performance Report', 'User DASD Groups Space Report', 'HSM Status Report', and 'SMS Configuration Report'. Below the tree are tabs for 'Operations', 'Logical', 'DEMO MGMT', 'Physical', and 'Demo System'.
- Top Right Pane:** A 3D bar chart titled 'Storage Groups with Free Space % < 10'. The Y-axis is labeled '% Free Space' and ranges from -2 to 8. The X-axis is labeled 'Storage Group'. There are two bars: 'CANDLE' with a very low value (near 0) and 'STORAGE NON-SMS VOLUMES' with a value of approximately 6.5.
- Bottom Pane:** A table titled 'SMS Storage Groups Space Summary'. The table has columns for Group Name, Storage Group Status, Storage Group Type, Non-Enabled Volumes, Total Volumes, VTOC Index Status, Low Volume Free Space %, High Volume Fragmentation Index, Largest Free Extent MB, and Free Mega. The 'DLGROUP' row is highlighted in red, and its 'Low Volume Free Space %' value of 2.6 is circled in black.

The status bar at the bottom shows 'Hub Time: Mon, 02/13/2006 04:41 PM', 'Server Available', and 'SMS Storage Groups Space - hqdt2.demopkg.ibm.com - DNET241'.

Link to SMS Storage Group volume space workspace



Issue #8: Key Data Sets Running Out of Extents

- In IBM Tivoli OMEGAMON XE for Storage on z/OS 3.1.0, you may define the concept of a data set group.
- A data set group is the combined data sets selected by the data set masks that are assigned to the group. Support for masking is available.
- You may write situations against the data set group or individual data sets to check for the possibility of running out of extents.
- Running out of extents for data sets could cause business applications to stop executing as expected.
- Data set groups provide these benefits:
 - ▶ Ability to monitor dataset space and performance, based on dataset name.
 - ▶ Extended VSAM metrics for CI/CA splits, RBA (relative byte address) information, and record counts.
 - ▶ Full multi-volume support.
 - ▶ Space utilization, allocated and used.
 - ▶ Ability to set situations (alerts) against all attributes.
 - ▶ History and Trending.



Issue #8: Key Data Sets Running Out of Extents

Dataset Group VSAM Attributes - TEHRLICH - SYSADMIN *ADMIN MODE*

File Edit View Help

Physical

Dataset Attributes

Attributes	Attributes Values	Statistics	Statistic Values	Allocations	Allocation Values
Data Component		Total Records	1	Volume	PRI1
Association	TDSMST.S3310SM3.S3310P14.RKDFDATA	Deleted Records	0	Space Type	CYLIND
Key Length	8	Inserted Records	0	Primary Space	
Relative Key Position		Updated Records	0	Secondary Space	
Average LRECL	1024	Retrieved Records	6	High Allocated RBA	1474560
Maximum LRECL	2048	CI Size	4096	High Used RBA	737280
Buffer Space	12288	CI Splits	0	Percent Free	50
Share Options	(3 , 3)	CA Splits	0	Percent Used	50
RECOVERY		Number CIs per CA	180		
UNIQUE		Free Space per CI	0		
NOERASE		Percent Free CIs	0		
INDEXED		Percent Free Bytes	0		

Dataset Name: TDSMST.S3310SM3.S3310P14.RKDFDATA.DATA

Aggregate Dataset Information

Dataset Type	Maximum MSR	Average MSR	Total I/Os	Maximum I/Os	Volume Count	Creation Date	Expiration Date	Reference Date	Last Backup Date	Owner	SMS Data Class	SMS Management Class	SMS...
Data Component	n/a	n/a	n/a	n/a	1	Jul 01 2005	***NONE***	Aug 08 2005	***NEVER***				

Dataset Name: TDSMST.S3310SM3.S3310P14.RKDFDATA.DATA

Dataset Information by Volume

Sequence Order	Volume	Average MSR	Total I/Os	Extents	Tracks Allocated	Tracks Used	Tracks Free	Percent Free	Percent Used	VSAM High Allocation RBA	VSAM High Used RBA	VSAM High CI RBA	Allocation Flag	Sequence Number
1	PRI169	n/a	n/a	1	30	30	0	50.00	50.00	1474560	737280	0	PRIMARY VOLUME	n/a

Dataset Name: TDSMST.S3310SM3.S3310P14.RKDFDATA.DATA

Ready | Hub Time: Wed, 08/10/2005 11:09 AM | Server Available. | Dataset Group VSAM Attributes - TEHRLICH - SYSADMIN *ADMIN MODE*

Issue #9: Key data sets suffering many CI and CA splits

- Application response is affected by CI (control interval) and CA (control area) splits in VSAM files.
- Splits indicate that a VSAM file has had its logical components separated, potentially increasing response time
- CA splits are generally worse than CI splits from a performance perspective.
- VSAM files must be taken offline and reorganized to rectify the splits, but this requires an application outage.
- From the *Dataset Group Summary* workspace, link to the *VSAM Cluster Report* workspace.



Issue #10: Volume low on space

- Currently the business applications are operating within expected and required service levels.
- A DASD device, while performing well at this time, is close to running out of storage space.
- The business application could abruptly stop working altogether.
- Options include writing a situation to check the percentage of volume space which is currently available (percent not in use).
- Individual data sets can be assessed, and candidates not frequently in use can be migrated using the Take Action feature.
- Select the *SMS Storage Groups Space workspace*. By selecting one of the volumes and drilling down, you can see all the data sets on the volume, as well as their space allocation and storage usage.



IBM Tivoli OMEGAMON XE for Storage on z/OS V3.1.0

Introduction

- IBM Tivoli OMEGAMON XE for Storage on z/OS provides these benefits:
 - ▶ Increases knowledge: Extensive reporting capabilities provide real-time access to reliable, up-to-the-minute information that affects your I/O subsystem.
 - ▶ Enhances system performance: You monitor and manage the applications and storage components that are critical to your business. OMEGAMON XE for Storage alerts you when conditions in your environment meet threshold-based conditions. These alerts notify your storage administrator when the I/O subsystem is not meeting your requirements.
 - ▶ Simplifies I/O subsystem management: You can manage your applications and resources across your I/O subsystem, easily identify trends that could result in I/O problems, and take corrective action before they become issues.



Question and Answer Session



Conclusion

- For More Information

- ▶ Bryan Turriff
Marketing Manager
IBM System z Software
310-727-4603
bturriff@us.ibm.com

- ▶ Rocket Software acquisition of Mainstar Software

http://www.mainstar.com/pdf/000-0123_Mainstar-Rocket_PR.pdf

