DFSMS OAM (Object) Hints and Tips – Performance Tuning and Diagnostics

Session: 22421

Albert Dennes
OAM L3 Service Lead and Development
IBM Corporation
aedennes@us.ibm.com

Derek Erdmann
OAM Level 2 Team Lead
IBM Corporation

Collecting OAM Experiences and Suggestions

- Some customers have had OAM object support in production for many years.
  - What are your experiences?
  - What have you learned?
  - What improvements would you like to see?
- Share your experiences with OAM Development
  - What procedures have you implemented successfully?
  - Even items that seem insignificant can be big bonuses to others
  - Interested in doing a joint presentation?
  - Send an email to aedennes@us.ibm.com and SHARE!

Agenda

- Performance Tuning
  - OAM configuration tuning
    - Environment configuration examples
  - Visualizing an ideal OSMC configuration
  - Data analysis via SMF records
  - Automation interest?
- Diagnostics
  - OSREQ return and reason codes
  - ACS routine guidelines
  - OSREQ STORE and resulting ACS routine flow
  - DB2 BIND diagnostics
  - OSMC cycle running longer than normal
  - OSREQ Query (or RETRIEVE) on DASD taking too long or SYSZTIOT contention
  - Out of synch collection vs catalog entries
- Questions?

Please feel free to stop us and ask any questions!

PERFORMANCE TUNING

OAM Configuration Tuning

- CBROAMxx parmlib member -
  "OAM address space - configuration tuning"
  - Establishes the environment under which the OAM address space runs by taking input configuration keyword and values
  - Crucial part to performance tuning your OAM configuration. Understanding each keyword within CBROAMxx will greatly increase the likelihood of a better operating environment.
  - Can be customized independently of ISMF and SMS
  - Statements include SETDISK, SETOAM, SETOPT, SETOSMC, SETTLIB, OAMXCF, and ONLYIF
  - OAM generally needs to be restarted after updates to CBROAMxx are made

    Note: There are some configuration values that can be updated dynamically via F OAM,UPDATE operator commands

OAM Configuration Tuning
"CBROAMxx - SETOAM Statements"

- MAXTAPERETRIEVETASKS (*tasks*) and MAXTAPESTORETASKS (*tasks*)
  - Global and SG Level - "Controls the number of tape drives for OAM use"
    - The summation of these two keywords is the maximum limit of tape drives that can be read from and written to concurrently by OAM
    - Max of 100 for each
    - OAM **will not** write to and read from more tape drives concurrently than the summation of these values
- SGMAXTAPERETRIEVETASKS (*tasks*) and SGMAXTAPESTORETASKS (*tasks*)
  - Storage Group Level - "Controls a subset of tape drives for each storage group"
    - Ideally, the summation of these storage group values should not exceed the global summation
    - Defaults to 1 if not specified for both keywords.
    - Can be change dynamically via **F OAM,UPDATE,SETOAM**…commands

OAM Configuration Tuning
"CBROAMxx - SETOAM Statements"

"An example of configuration that can lead to starvation..."

```
EDIT       SYS1.PARMLIB(CBROAMCD) - 01.73
Command ===>
****** ******************************** Top of Data ********
000070 SETOAM MAXTAPERETRIEVETASKS(5)
000080        MAXTAPESTORETASKS(5)
000090        MOUNTWAITTIME(1)
000091        DATACLASS(VTSM16G)
000100        STORAGEGROUP(GROUP01
000101                     TAPEUNITNAME(3490)
000102                     SGMAXTAPERETRIEVETASKS(5)
000103                     SGMAXTAPESTORETASKS(5))
000104        STORAGEGROUP(GROUP02
000105                     TAPEUNITNAME(3490)
000106                     SGMAXTAPERETRIEVETASKS(5)
000107                     SGMAXTAPESTORETASKS(5))
000108        STORAGEGROUP(GROUP03
000109                     TAPEUNITNAME(3490)
000110                     SGMAXTAPERETRIEVETASKS(5)
000111                     SGMAXTAPESTORETASKS(5))
000124
```

- Global values are also set to 5 (red)
- Storage group values set to 5 (yellow)
- Total summation of tasks available for all storage groups is 10
- GROUP01 could potentially utilize all of the allocated tasks and starve the other storage groups from tape drives

If (Global ST + Global RET) < sum of all (SG ST + SG RET) then

Possible starvation issues

OAM Configuration Tuning
"CBROAMxx - SETOAM Statements"

"An example of better configuration to eliminate starvation..."

```
EDIT       SYS1.PARMLIB(CBROAMCD) - 01.73
Command ===>
****** ******************************** Top of Data ******
000070 SETOAM MAXTAPERETRIEVETASKS(15)
000080        MAXTAPESTORETASKS(15)
000090        MOUNTWAITTIME(1)
000091        DATACLASS(VTSM16G)
000100        STORAGEGROUP(GROUP01
000101                     TAPEUNITNAME(3490)
000102                     SGMAXTAPERETRIEVETASKS(5)
000103                     SGMAXTAPESTORETASKS(5))
000104        STORAGEGROUP(GROUP02
000105                     TAPEUNITNAME(3490)
000106                     SGMAXTAPERETRIEVETASKS(5)
000107                     SGMAXTAPESTORETASKS(5))
000108        STORAGEGROUP(GROUP03
000109                     TAPEUNITNAME(3490)
000110                     SGMAXTAPERETRIEVETASKS(5)
000111                     SGMAXTAPESTORETASKS(5))
000124
```

- Workloads analyzed and deemed appropriate for tape drive resources available for OAM use to be less than or equal to 30
- Global values are increased to 15 (blue)
- Storage group values maintained at 5 (purple)
- Total summation of tasks available for all storage groups is 30
- This allows each storage group to reach it's specified maximum limit without causing starvation

Note: After analysis it may also be acceptable to decrease the amount of maximum tasks available at the storage group level with a summation of less than or equal to the existing global summation max

If (Global ST + Global RET) = sum of all (SG ST + SG RET) then
Optimal configuration

OAM Configuration Tuning
"CBROAMxx - SETOAM Statements"

"An example of configuration that can lead to resource inefficiency..."

```
EDIT       SYS1.PARMLIB(CBROAMCD) - 01.73
Command ===>
****** ***************************** Top of Data
000070 SETOAM MAXTAPERETRIEVETASKS(100)
000080        MAXTAPESTORETASKS(100)
000090        MOUNTWAITTIME(1)
000091        DATACLASS(VTSM16G)
000100        STORAGEGROUP(GROUP01
000101                     TAPEUNITNAME(3490))
000104        STORAGEGROUP(GROUP02
000105                     TAPEUNITNAME(3490))
000108        STORAGEGROUP(GROUP03
000109                     TAPEUNITNAME(3490))
000124
```

- Global values are set to 100 (red)
- Storage group values defaulted to 1 (implicit).
- Total summation of tasks available for all storage groups is 200
- Every time a read or write request comes in for any of these storage groups OAM will wake up all 200 tasks, but... storage groups are only capable of using at max of 2 tasks

If (Global ST + Global RET) > sum of all (SG ST + SG RET) then
Resource inefficiency

OAM Configuration Tuning
"CBROAMxx - SETOAM Statements"

"An example of better configuration that can optimize resource usage..."

```
EDIT       SYS1.PARMLIB(CBROAMCD) - 01.73
Command ===>
****** ******************************** Top of Dat
000070 SETOAM MAXTAPERETRIEVETASKS(3)
000080        MAXTAPESTORETASKS(3)
000090        MOUNTWAITTIME(1)
000091        DATACLASS(VTSM16G)
000100        STORAGEGROUP(GROUP01
000101                     TAPEUNITNAME(3490))
000104        STORAGEGROUP(GROUP02
000105                     TAPEUNITNAME(3490))
000108        STORAGEGROUP(GROUP03
000109                     TAPEUNITNAME(3490))
000124
```

- Workloads analyzed and deemed appropriate for storage groups to be reduced to a max of 2 tasks (1 for store, 1 for retrieve)
- Global values reduced down to 3 (blue)
- Storage group values remain unchanged at the default of 1 (implicit)
- Total summation of tasks available for all storage groups is now 6
- Every time a read or write request comes in for any of these storage groups OAM will only wake up 6 tasks

If (Global ST + Global RET) = sum of all (SG ST + SG RET) then
Optimized resource use

OAM Configuration Tuning
"CBROAMxx - SETOAM Statements"

"An example of configuration that can lead to allocation recovery..."

```
EDIT        SYS1.PARMLIB(CBROAMCD) - 01.77
Command ===>
****** ***************************** Top of Data ****
000070 SETOAM MAXTAPERETRIEVETASKS(9)
000080        MAXTAPESTORETASKS(9)
000090        MOUNTWAITTIME(1)
000091        DATACLASS(VTSM16G)
000100        STORAGEGROUP(GROUP01
000101                     TAPEUNITNAME(3490)
000102                     SGMAXTAPERETRIEVETASKS(3)
000103                     SGMAXTAPESTORETASKS(3))
000106        STORAGEGROUP(GROUP02
000107                     TAPEUNITNAME(3490)
000108                     SGMAXTAPERETRIEVETASKS(3)
000109                     SGMAXTAPESTORETASKS(3))
000110        STORAGEGROUP(GROUP03
000111                     TAPEUNITNAME(3490)
000112                     SGMAXTAPERETRIEVETASKS(3)
000113                     SGMAXTAPESTORETASKS(3))
000114
```

- Global values are set to 9 (red)
- Storage group values set to 3 (yellow)
- Total summation of tasks available for all storage groups is 18
- Tape drives available for OAM is less than the amount of resources we are requesting... in this case we have 6 drives available for OAM use

If (Global ST + Global RET) > # tape drives available to OAM then
Allocation recovery

OAM Configuration Tuning
"CBROAMxx - SETOAM Statements"

"An example of better configuration that prevents allocation recovery..."

```
EDIT       SYS1.PARMLIB(CBROAMCD) - 01.77
Command ===>
****** ***************************** Top of Data ****
000070 SETOAM MAXTAPERETRIEVETASKS(3)
000080        MAXTAPESTORETASKS(3)
000090        MOUNTWAITTIME(1)
000091        DATACLASS(VTSM16G)
000100        STORAGEGROUP(GROUP01
000101                     TAPEUNITNAME(3490)
000102                     SGMAXTAPERETRIEVETASKS(1)
000103                     SGMAXTAPESTORETASKS(1))
000106        STORAGEGROUP(GROUP02
000107                     TAPEUNITNAME(3490)
000108                     SGMAXTAPERETRIEVETASKS(1)
000109                     SGMAXTAPESTORETASKS(1))
000110        STORAGEGROUP(GROUP03
000111                     TAPEUNITNAME(3490)
000112                     SGMAXTAPERETRIEVETASKS(1)
000113                     SGMAXTAPESTORETASKS(1))
000114
```

- Workloads analyzed and deemed appropriate for storage groups to be reduced to a max of 2 task (1 for store, 1 for retrieve).
- Global values reduced to 3 (blue).
- Storage group values reduced to 1 (purple).
- Total summation of tasks available for all storage groups is 6
- Tape drives available for OAM now matches the amount specified for this configuration

Note: After analysis it may also be acceptable to increase the amount of driver resources available to OAM if the workloads do verify the need for that many resources

If (Global ST + Global RET) <= # tape drives available to OAM then

No allocation recovery risk

OAM Configuration Tuning
"CBROAMxx - SETOAM Statements"

- TAPEDRIVESTARTUP (*threshold in megabytes)*
  Storage Group Level - "Determines when OAM is to start additional tape drive"
  - Threshold value in MB that specifies the amount of write data that needs to be pending before OAM will start an additional tape drive for this storage group
  - If not explicitly specified, defaults to 9999 MB
  - To use more than one tape drive for a storage group to write object data to tape, the threshold value must be low enough to trigger the startup of the additional tape drive
  - Maximum number of tape drives used is still limited to what was specified on global MAXTAPESTORETASKS as well as what was specified for SGMAXTAPESTORETASKS at the storage group level
  - Can be change dynamically via **F OAM,UPDATE,SETOAM**…commands

If (# of MB write data pending / # of tape drives currently writing) > TAPEDRIVESTARTUP value then
  OAM will attempt to use another tape drive
Else
  OAM will continue to use existing tape drives even if there is a higher specification on MAXTAPESTORETASKS

OAM Configuration Tuning
"CBROAMxx - SETOAM Statements"

"An example of configuration that can lead to in-efficiencies with tape drive usage for heavy workloads.."

```
EDIT       SYS1.PARMLIB(CBROAMCD) - 01.77
Command ===>
****** ***************************** Top of Data ****
000070 SETOAM MAXTAPERETRIEVETASKS(15)
000080        MAXTAPESTORETASKS(15)
000090        MOUNTWAITTIME(1)
000091        DATACLASS(VTSM16G)
000100        STORAGEGROUP(GROUP01
000101                     TAPEUNITNAME(3490)
000102                     SGMAXTAPERETRIEVETASKS(5)
000103                     SGMAXTAPESTORETASKS(5))
000106        STORAGEGROUP(GROUP02
000107                     TAPEUNITNAME(3490)
000108                     SGMAXTAPERETRIEVETASKS(5)
000109                     SGMAXTAPESTORETASKS(5))
000110        STORAGEGROUP(GROUP03
000111                     TAPEUNITNAME(3490)
000112                     SGMAXTAPERETRIEVETASKS(5)
000113                     SGMAXTAPESTORETASKS(5))
```

- SGMAXTAPESTORETASKS max at 5 tasks for GROUP01.
- TAPEDRIVESTARTUP not specified, therefore default of 9999 MB is taken
- During heavy workloads, GROUP01 is only utilizing 1 tape drive even though a max of 5 is specified for this storage group.

If SGMAXTAPESTORE > # tape drives currently used While(HeavyWriteWorkloads) then Tape Drive Usage Inefficiency

OAM Configuration Tuning
"CBROAMxx - SETOAM Statements"

"An example of a better configuration that will optimize tape drive usage for heavy workloads.."

```
EDIT       SYS1.PARMLIB(CBROAMCD) - 01.77
Command ===>
****** ******************************** Top of Data ****
000070 SETOAM MAXTAPERETRIEVETASKS(15)
000080        MAXTAPESTORETASKS(15)
000090        MOUNTWAITTIME(1)
000091        DATACLASS(VTSM16G)
000100        STORAGEGROUP(GROUP01
000101                     TAPEUNITNAME(3490)
000102                     TAPEDRIVESTARTUP(1)
000103                     SGMAXTAPERETRIEVETASKS(5)
000104                     SGMAXTAPESTORETASKS(5))
000106        STORAGEGROUP(GROUP02
000107                     TAPEUNITNAME(3490)
000108                     TAPEDRIVESTARTUP(1)
000109                     SGMAXTAPERETRIEVETASKS(5)
000110                     SGMAXTAPESTORETASKS(5))
000111        STORAGEGROUP(GROUP03
000112                     TAPEUNITNAME(3490)
000113                     TAPEDRIVESTARTUP(1)
000114                     SGMAXTAPERETRIEVETASKS(5)
000115                     SGMAXTAPESTORETASKS(5))
000116
```

- SGMAXTAPESTORETASKS max at 5 tasks for GROUP01
- TAPEDRIVESTARTUP keyword added with value of 1MB (purple)
- During heavy workloads, GROUP01 is now utilizing 5 tape drives since the threshold of 1 MB was reached for pending write data

Note: There may be cases where utilizing all tape drives specified on SGMAXTAPESTORETASKS is undesirable... in those cases fine tuning of the TAPEDRIVESTARTUP value is needed to achieve proper balance. OP CMD is handy here.

If SGMAXTAPESTORE = # tape drives currently used While(HeavyWriteWorkloads) then

Tape Drive Usage Optimized

OAM Configuration Tuning
"CBROAMxx - SETOAM Statements"

- DEMOUNTWAITTIME (seconds)

  Global and SG level – "How long in a non-busy environment to keep a volume mounted if not being read/written"

  - Specifies the time, in seconds, that OAM waits before demounting and deallocating a tape drive that OAM is not currently using
  - In a busy environment this keyword is ignored
  - Seconds specification of 1-9999
  - Default is 120 seconds
  - Special conditions –
    - If a new tape drive allocation comes in and OAM has already used the maximum # of tape drives specified, then this keyword is ignored and we demount the volume on the drive in order to mount another.
    - If OAM is cancelled, this keyword will be ignored and demounts will occur
    - If a request to vary the drive offline is sent while this keyword is in effect for that drive, the drive cannot vary offline until the specified time elapses

OAM Configuration Tuning
"CBROAMxx - SETOAM Statements"

"An example of configuration that can lead to tape volume thrashing for intermittent work.."

```
EDIT        SYS1.PARMLIB(CBROAMCD) - 01.79
Command ===>
****** ***************************** Top of Data ******
000070 SETOAM MAXTAPERETRIEVETASKS(15)
000080        MAXTAPESTORETASKS(15)
000090        MOUNTWAITTIME(1)
000091        DATACLASS(VTSM16G)
000100        STORAGEGROUP(GROUP01
000101                     TAPEUNITNAME(3490)
000103                     SGMAXTAPERETRIEVETASKS(5)
000104                     SGMAXTAPESTORETASKS(5))
000106        STORAGEGROUP(GROUP02
000107                     TAPEUNITNAME(3490)
000109                     SGMAXTAPERETRIEVETASKS(5)
000110                     SGMAXTAPESTORETASKS(5))
000111        STORAGEGROUP(GROUP03
000112                     TAPEUNITNAME(3490)
000114                     SGMAXTAPERETRIEVETASKS(5)
000115                     SGMAXTAPESTORETASKS(5))
```

- DEMOUNTWAITTIME not explicitly specified, therefore default of 120 seconds is used
- Read/write request for the GROUP01 come in every 180 seconds
- Constant thrashing of mount/demount observed within SYSLOG hindering tape performance

OAM Configuration Tuning
"CBROAMxx - SETOAM Statements"

"An example of better configuration that can prevent tape volume thrashing for intermittent work.."

```
EDIT        SYS1.PARMLIB(CBROAMCD) - 01.80
Command ===>
****** ***************************** Top of Data ****
000070 SETOAM  MAXTAPERETRIEVETASKS(15)
000080         MAXTAPESTORETASKS(15)
000090         MOUNTWAITTIME(1)
000091         DATACLASS(VTSM16G)
000100         STORAGEGROUP(GROUP01
000101                      DEMOUNTWAITTIME(240)
000102                      TAPEUNITNAME(3490)
000103                      SGMAXTAPERETRIEVETASKS(5)
000104                      SGMAXTAPESTORETASKS(5))
000106         STORAGEGROUP(GROUP02
000107                      TAPEUNITNAME(3490)
000109                      SGMAXTAPERETRIEVETASKS(5)
000110                      SGMAXTAPESTORETASKS(5))
000111         STORAGEGROUP(GROUP03
000112                      TAPEUNITNAME(3490)
000114                      SGMAXTAPERETRIEVETASKS(5)
000115                      SGMAXTAPESTORETASKS(5))
000116
```

- DEMOUNTWAITTIME specified with a value of 240 seconds (blue)
- Read/write request for GROUP01 come in every 180 seconds
- Now volumes remain mounted past the intermittent work point of 180 seconds due to the demountwaittime specification of 240 seconds
- Thrashing prevented

"Data and workload analysis of your systems is critical in determining optimum performance configuration settings."

OAM Configuration Tuning
"CBROAMxx - SETOAM Statements"

- TAPEDISPATCHERDELAY (*seconds*)

  Global Level – "How long to wait in a busy environment before demounting a tape volume"

  - Specifies the time in seconds that OAM will wait before demounting a tape volume, even if other work is available for this drive
  - Seconds specification of 1-60
  - Default is 0
  - This delay allows time for a new read request to come into OAM that requires the currently mounted tape volume and can greatly reduce the number of mounts and demounts of volumes for certain applications

OAM Configuration Tuning
"CBROAMxx - SETOAM Statements"

"An example of configuration that can lead to tape volume thrashing for busy work.."

```
EDIT        SYS1.PARMLIB(CBROAMCD) - 01.80
Command ===>
****** ***************************** Top of Data ****
000070 SETOAM MAXTAPERETRIEVETASKS(15)
000080        MAXTAPESTORETASKS(15)
000090        MOUNTWAITTIME(1)
000091        DATACLASS(VTSM16G)
000100        STORAGEGROUP(GROUP01
000102                     TAPEUNITNAME(3490)
000103                     SGMAXTAPERETRIEVETASKS(5)
000104                     SGMAXTAPESTORETASKS(5))
000106        STORAGEGROUP(GROUP02
000107                     TAPEUNITNAME(3490)
000109                     SGMAXTAPERETRIEVETASKS(5)
000110                     SGMAXTAPESTORETASKS(5))
000111        STORAGEGROUP(GROUP03
000112                     TAPEUNITNAME(3490)
000114                     SGMAXTAPERETRIEVETASKS(5)
000115                     SGMAXTAPESTORETASKS(5))
000116
```

- TAPEDISPATCHERDELAY not explicitly specified, therefore default of 0 seconds is used.
- Currently processing reads on a volume every 5 seconds
- Read request for a different volume come in every second
- No more available drives/tasks available. All being used
- Constant thrashing of mount/demount observed within SYSLOG hindering tape performance

OAM Configuration Tuning
"CBROAMxx - SETOAM Statements"

"An example of better configuration that can prevent tape volume thrashing for busy work.."

```
EDIT       SYS1.PARMLIB(CBROAMCD) - 01.80
Command ===>
****** ***************************** Top of Data ******
000070 SETOAM MAXTAPERETRIEVETASKS(15)
000080        MAXTAPESTORETASKS(15)
000090        MOUNTWAITTIME(1)
000091        DATACLASS(VTSM16G)
000092        TAPEDISPATCHERDELAY(15)
000100        STORAGEGROUP(GROUP01
000102                     TAPEUNITNAME(3490)
000103                     SGMAXTAPERETRIEVETASKS(5)
000104                     SGMAXTAPESTORETASKS(5))
000106        STORAGEGROUP(GROUP02
000107                     TAPEUNITNAME(3490)
000109                     SGMAXTAPERETRIEVETASKS(5)
000110                     SGMAXTAPESTORETASKS(5))
000111        STORAGEGROUP(GROUP03
000112                     TAPEUNITNAME(3490)
000114                     SGMAXTAPERETRIEVETASKS(5)
000115                     SGMAXTAPESTORETASKS(5))
```

- TAPEDISPATCHERDELAY specified with a value of 15 seconds
- Currently processing reads on a volume every 5 seconds
- Read request for a different volume come in every second
- No more available drives/tasks available. All being used.
- Since we increased the TAPEDISPATCHERDELAY time more than 5 seconds, we are able to wait for the next read for the same volume even though a pending read for another volume is waiting
- We are able to keep this volume mounted and prevent thrashing

We may also need to investigate further and see if we need more drive tasks specified with the MAXTAPERETRIEVETASKS to allow for more parallel type workloads to be active.

OAM Configuration Tuning
"CBROAMxx - SETOSMC Statements"

- CYCLEWINDOW (*start_mode*)

  Global Level - "Specifies the start window for OSMC for a given storage group"

  - Can choose either the STARTONLY mode or the STARTSTOP mode.
  - If not explicitly specified, defaulted to STARTONLY
  - CYCLE START TIME and CYCLE END TIME attributes are used in conjunction with CYCLEWINDOW to start (or stop) the OSMC process for a storage group automatically
    - These values are defined in each Object or Object Backup storage group through ISMF

  Note: If you manually start the OSMC cycle by issuing a F OAM,START,OSMC or F OAM,START,STORGRP, the CYCLEWINDOW mode and the start and stop times for the storage group cycle are ignored

OAM Configuration Tuning
"CBROAMxx - SETOSMC Statements"

"An example of configuration that can lead to runaway OSMC processing..."

```
EDIT       SYS1.PARMLIB(CBROAMAD) - 01.12                Columns 0000
Command ===>                                               Scroll ==
****** ******************************** Top of Data ***********************
000070 SETOAM MAXTAPERETRIEVETASKS(1)
000080        MAXTAPESTORETASKS(1)
000090        MOUNTWAITTIME(1)
000091        DATACLASS(VTSM16G)
000100        STORAGEGROUP(GROUP01 TAPEUNITNAME(3490))
000101        STORAGEGROUP(GROUP02 TAPEUNITNAME(3490))
000102        STORAGEGROUP(GROUP03 TAPEUNITNAME(3490))
000103        STORAGEGROUP(GROUP04 TAPEUNITNAME(3490))
000104        STORAGEGROUP(GROUP05 TAPEUNITNAME(3490))
000110
000120 SETDISK STORAGEGROUP(GROUP04 L2TYPE(ZFS) L2DIR(/z/oam/group04))
000130
000140 SETOSMC FIRSTBACKUPGROUP(IMAFIRST)
000150         SECONDBACKUPGROUP(IMSECOND)
000160         MAXRECALLTASKS(5)
****** ******************************** Bottom of Data ********************
```

```
                          OBJECT STORAGE GROUP DISPLAY
Command ===>

CDS Name  . . . . . : OAMWRK.EZU180.SCDS
Storage Group Name  : GROUP01

Description  :


Qualifier  . . . . . : GROUP01
Cycle Start Time . . . . . : 23
      End Time . . . . . . : 4
```

- CYCLEWINDOW not explicitly specified, default of STARTONLY taken
- Storage group definition has CYCLESTART at 23 (11pm) and CYCLEEND at 4 (4am)
- OSMC cycle starts at 11pm as indicated, but.. continues to run and starts new work well past 4am
- Cycle still running during production workloads drawing CPU/resources leading to slowdowns

OAM Configuration Tuning
"CBROAMxx - SETOSMC Statements"

"An example of better configuration that can finely tune OSMC timing..."

```
EDIT       SYS1.PARMLIB(CBROAMAD) - 01.12                  Columns 00001
Command ===> _                                               Scroll ===
****** ******************************** Top of Data ************************
000070 SETOAM MAXTAPERETRIEVETASKS(1)
000080        MAXTAPESTORETASKS(1)
000090        MOUNTWAITTIME(1)
000091        DATACLASS(VTSM16G)
000100        STORAGEGROUP(GROUP01 TAPEUNITNAME(3490))
000101        STORAGEGROUP(GROUP02 TAPEUNITNAME(3490))
000102        STORAGEGROUP(GROUP03 TAPEUNITNAME(3490))
000103        STORAGEGROUP(GROUP04 TAPEUNITNAME(3490))
000104        STORAGEGROUP(GROUP05 TAPEUNITNAME(3490))
000110
000120 SETDISK STORAGEGROUP(GROUP04 L2TYPE(ZFS) L2DIR(/z/oam/group04))
000130
000140 SETOSMC FIRSTBACKUPGROUP(IMAFIRST)
000150         SECONDBACKUPGROUP(IMSECOND)
000160         MAXRECALLTASKS(5)
000170         CYCLEWINDOW(STARTSTOP)
****** ******************************** Bottom of Data *********************
```

```
                          OBJECT STORAGE GROUP DISPLAY
Command ===> _

CDS Name  . . . . . : OAMWRK.EZU180.SCDS
Storage Group Name  : GROUP01

Description  :


Qualifier  . . . . . : GROUP01
Cycle Start Time . . . . . : 23
      End Time . . . . . . : 4
```

- CYCLEWINDOW specified with STARTSTOP value (blue).
- Storage group definition has CYCLESTART at 23 (11pm) and CYCLEEND at 4 (4am)
- OSMC cycle starts at 11pm as indicated and stops processing new work at 4 am.
- Cycle finishes processing before heavy production workloads come in ensuring appropriate system resources are available.

OAM Configuration Tuning

- IEFSSNxx parmlib member -
  - “OAM subsystem - configuration tuning"
    - Establishes the environment under which the OAM subsystem runs by taking input configuration keyword and values
    - Also a crucial part to performance tuning your OAM configuration
    - TIME=, MSG=, OTIS=, UPD=, MOS=, LOB=, QB=, DP=, and D= values (D added in V2R3)
    - Any change here generally requires a re-IPL
      - Note: with the exception of changing between Multiple and Classic OAM configurations via F OTIS,DELSUB and SETSSI ADD in V2R3

OAM Configuration Tuning
"IEFSSNxx – OAM subsystem options"

- UPD=x

  "Whether DB2 updates for ODPENDDT and ODLREFDT should be performed"

  - UPD=Y specifies that the pending action date and the last reference date be updated on all OSREQ RETRIEVEs
    - This is the default
  - UPD=C specifies that the pending action date and the last reference date be updated on all OSREQ RETRIEVES and OSREQ CHANGEs
  - UPD=N specifies no updates to be done on OSREQ RETREIVE

  "UPD=N can reduce unnecessary retrieval and update of objects' directory entries during the OSMC cycle if your installation's management classes do not use the TIME SINCE LAST USE or EXPIRE AFTER DAYS USAGE parameters."

  Note: Regardless of the setting for UPD=, the ODLREFDT field is not updated for RETRIEVEs that result in RECALL or for objects currently in RECALL mode

OAM Configuration Tuning
"IEFSSNxx – OAM subsystem options"

- QB=x

  "Whether OSREQ QUERY results in a call into OAM to retrieve backup information"

  - QB=Y causes every OSREQ QUERY to result in a call into OAM for each backup copy. The query returns a complete backup retrieval order key (volser_sectorLocation) for each backup copy. If a backup copy does not exist, then OAM is not called and the order keys contain zeros.
    - This is the default
  - QB=N do not call OAM on OSREQ QUERY

  "QB=N can reduce unnecessary calls into the OAM address space for each backup copy on an OSREQ QUERY request."

  Note: Many applications issue an OSREQ QUERY prior to an OSREQ RETRIEVE and do not want or need the queried information on the backups. In addition, gathering the backup information can extend the overall retrieval response for the application.

OAM Configuration Tuning
"IEFSSNxx – OAM subsystem options"

- LOB=x

  "Whether OAM uses DB2 LOB support for objects greater than 32KB"

  - LOB=P indicates that a partial list of storage groups are capable of storing greater than 32KB in LOB structures

    "This causes OAM to query DB2 on every store to check if the LOB table exists for a given storage group."

  - LOB=A specifies that for all storage groups objects greater than 32KB be stored in a LOB storage structure when stored to DB2

    "This results in optimal performance when you want to store large objects to DB2, because OAM does not query DB2 to see if the LOB base table view exists."

Note: A LOB storage structure must be used for objects greater than 256M. Also, MOS= specification must be compatible with object sizes being stored.

OAM Configuration Tuning

- OAM proc
  "OAM address space – configuration tuning for started proc"
  - Started task for the OAM address space
  - Defined in sample job CBRAPROC
  - Input parameters OSMC=,MAXS=,UNLOAD=,EJECT=,APLAN=, REST= and D= (D= added in V2R3)
  - Contains characteristics about OAM address space behavior dealing with restart timing, OSMC task configuration, optical related processing, and DB2 connection specification

OAM Configuration Tuning
"OAM Proc"

- MAXS

  "Controls the number of storage groups that OSMC processes concurrently"
  - Default value of 2
  - Works in conjunction with TAPEDRIVESTARTUP, MAXTAPESTORETASKS, and MAXTAPERETRIEVETASKS
  - Do not specify a MAXS value larger than the number of drives that are available for OSMC processing.
  - Consider other types of OSMC resources before adjusting MAXS value, i.e. MOVEVOL, RECOVERY, or RECYCLE commands
  - Use OAMplex for concurrent OSMC processing
    - Even if you are not using an OAMplex for multiple system access to data, or cross system applications
  - In an OAMplex, the OSMC processing system name for each Object or Object Backup storage group controls where OSMC processing is done for that storage group
    - Using this parameter and separating hardware between storage groups can balance workload across systems for OSMC processing
    - Localize hardware and highest usage to reduce XCF overhead

OAM Configuration Tuning
"OAM Proc"

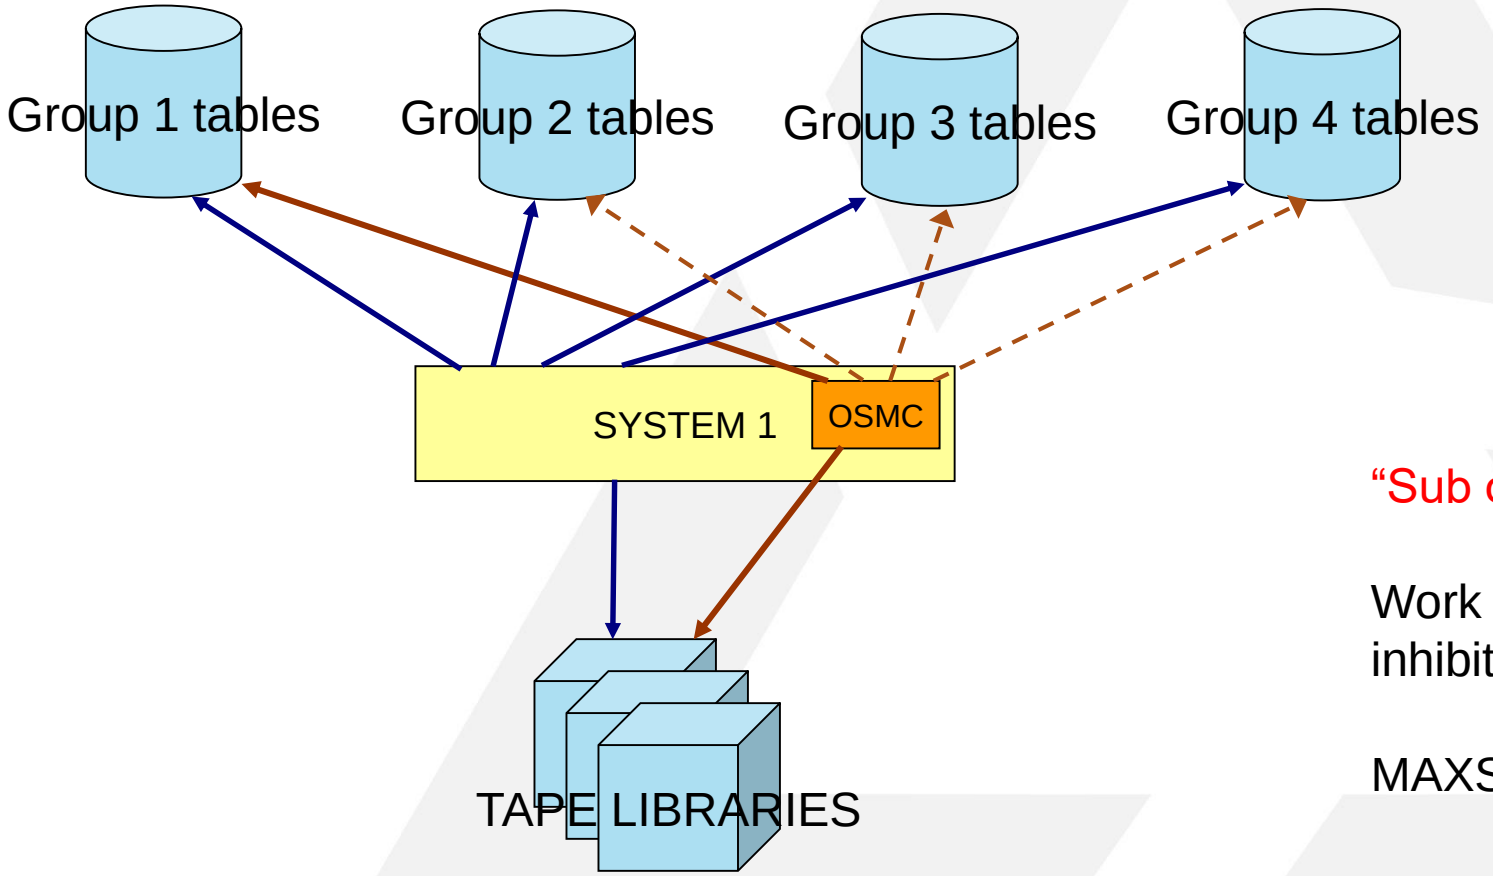

"Sub optimal case"

Work sequentially queued on one OSMC task inhibiting performance

MAXS set to 1

OAM Configuration Tuning
"OAM Proc"

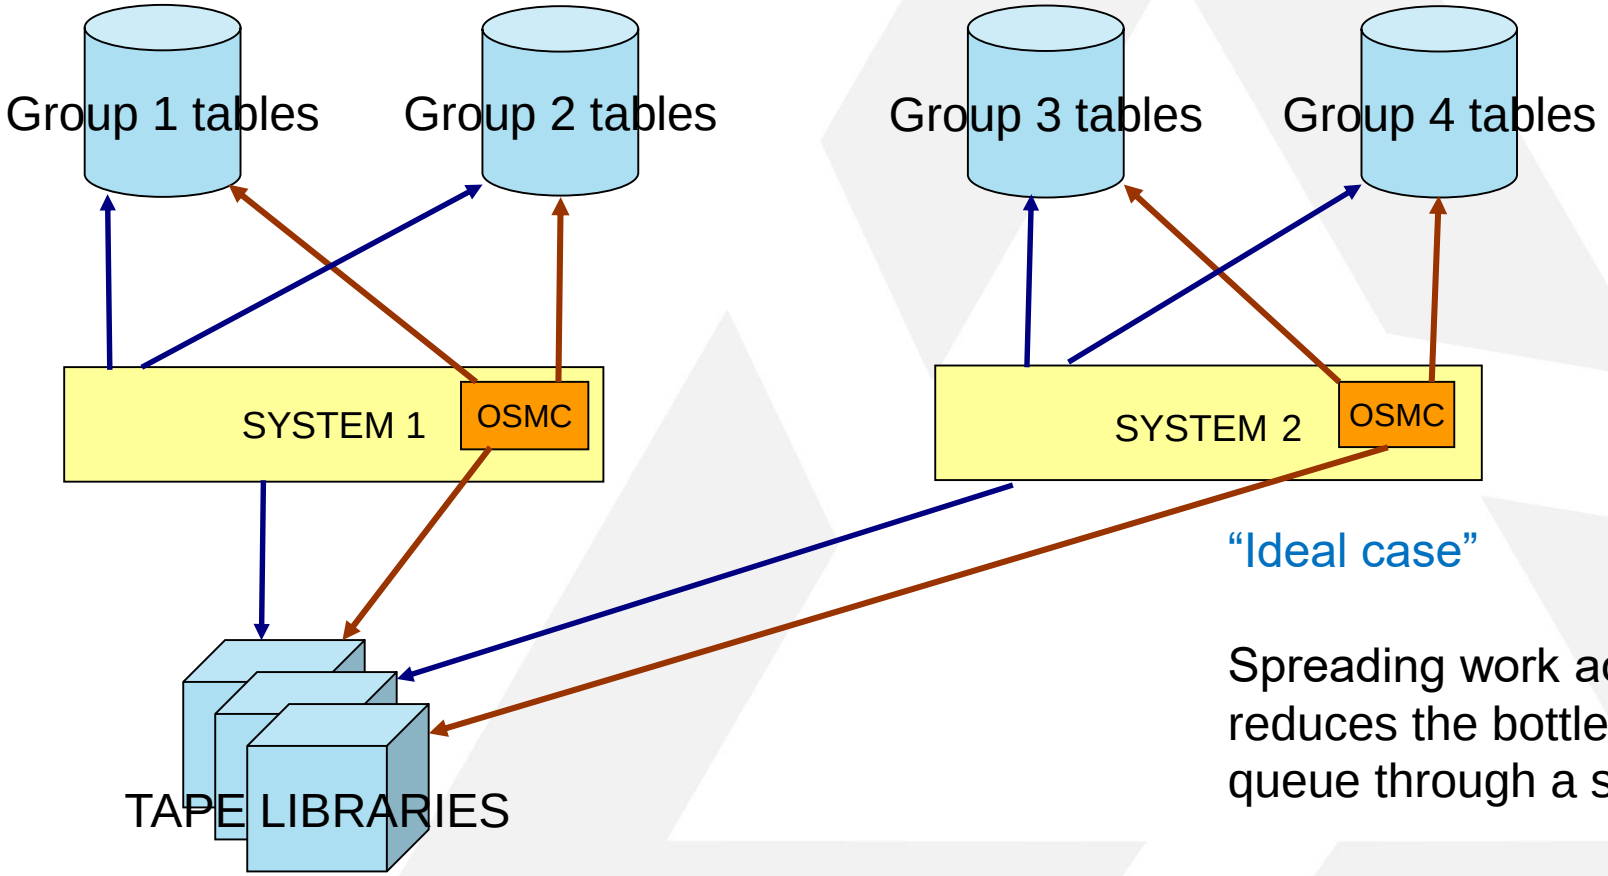

"Ideal case"

Spreading work across multiple OAM's / OSMC's reduces the bottleneck of a single input work queue through a single address space.

MAXS set to 2
Group 1 and Group 2 processing system set to SYS1
Group 3 and Group 4 processing system set to SYS 2

OAM Configuration Tuning

- CBRHADUX

  "OSMC AUTO-DELETE Installation Exit"

  - Used to verify whether or not an object should be deleted via OSMC processing
  - Default behavior is to not allow deletion – RC 12 and to not invoke CBRHADUX again on subsequent delete requests issued by this storage group
  - Using this exit can cause OSMC performance problems if you have not properly established your expiration criteria in your SMS management classes
    - For example, if using the default behavior "never delete", then management class should have expiration criteria set to NEVER expire. Otherwise, OSMC will continually pick these objects for processing every time it is started for that storage group.

OAM Configuration Tuning

OAM SMF records are a great source of data for analysis

"SMF 85 subtypes 78/79/87 to monitor drive usage (device number/volume)"

- For instance, a drive demount SMF record will show how many objects were read or written while mounted, how long mounted, etc. Data can be used to determine how efficient drives are being used
- OSMC records show total numbers, total MB, read or written to different media types, expired, etc.
- Summary records, like demount, OSMC storage group processing, MOVEVOL processing, can provide performance or planning data

OAM Configuration Tuning – Automation Interest?

"Any interest in an automated OAM configuration tuning advisor tool?

Potentially capable of performing data analysis on OAM SMF record subtypes (workload activity) and as a result - provide estimated "optimal" configuration settings for your specific environment?"

Note: If RFEs are received for this type of tool, delivery can not be promised. Plans are always subject to change.

DIAGNOSTICS

OSREQ Return and Reason Codes

- **OSREQ sets return and reason codes during processing of OSREQ functions**
  - Surfaced by OAM or the application that invokes the OSREQ macro
  - Documented in DFSMSDfp Diagnosis
- **Each byte of the reason code has unique meaning:**
  - **Byte 0:** Contains a unique OSREQ reason code
  - **Byte 1:** Contains an internal OSREQ function code **x** (used by IBM service)
  - **Bytes 2 and 3:** Represent the bytes x and y (in the table) where
    - **y** is an error indicator
    - **z** is reserved for IBM use, unless otherwise indicated
  - **Example:**

8 (X'08')	X'30'	x	y	z	Object already exists.
8 (X'08')	X'30'	x	X'01'	z	Directory entry already exists.
	X'30'	x	X'02'	z	Object segment already exists.

OSREQ Return and Reason Codes (cont.)

- **Some reason codes require additional work by the end user in debugging**
- Example:

12 (X'0C')	X'74'	x	y	z	DB2 detected failure (other). yz represents the DB2 SQL error. Convert the yz to decimal, and look up the resulting SQL code in IBM Information Management Software for z/OS Solutions Information Center at http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp. **yz** DB2 SQL code

- **How to convert yz to the corresponding SQL error code**
  - Use two's complement
    - Subtract the reason code from x'FFFF', then add 1
    - Convert from hex to decimal, then add the minus sign
  - Example:
    - Last two bytes of reason code are x'FCDB'
    - X'FFFF' – x'FCDB' = x'324' + 1 = x'325' = 805
    - SQL error code is -805

ACS routine guidelines

OAM Read-Only Variables	
&MEMHLQ	&MEMN
&MEMLLQ	&MEMNQUAL

The ACS routines will be invoked by OAM in 4 environments (&ACSENVIR)

- **STORE** – an OSREQ STORE was issued for an object
- **CHANGE** – an OSREQ CHANGE was issued for an object
- **CTRANS** – OSMC has detected that an object has met criteria for a class transition
- **ALLOC** – an OSREQ STORE was issued for an object and the object is being directed to tape

ACS routines invoked (&ACSENVIR)	STORE	CHANGE	CTRANS	*ALLOC
DATA CLASS	NO	NO	NO	YES
STORAGE CLASS	**YES	***YES	YES	YES
MANAGEMENT CLASS	**YES	***YES	YES	NO
STORAGE GROUP	**YES	NO	NO	YES

- *ACS environment of ALLOC is invoked during allocation
- **When storing the first object into a new collection, the ACS routines are entered. For existing collections, the ACS routines are NOT entered
- ***Only entered if that construct was specified on the OSREQ CHANGE request

OSREQ STORE and resulting ACS routine flow

- Flow has changed slightly for z/OS 2.3 and beyond
- Collections no longer stored/referenced in the Catalog
- Collections only stored/referenced in the DB2 Collection Name table

- Potentially 3 calls to the ACS routines
  - First call occurs if this is the first OSREQ STORE of a new collection (**&ACSENVIR = 'STORE'**)
  - Second call occurs if a specific storage class or management class is specified on the OSREQ STORE (**&ACSENVIR = 'STORE'**)
  - Third call occurs if the object is being directed to tape (**&ACSENVIR = 'ALLOC')**

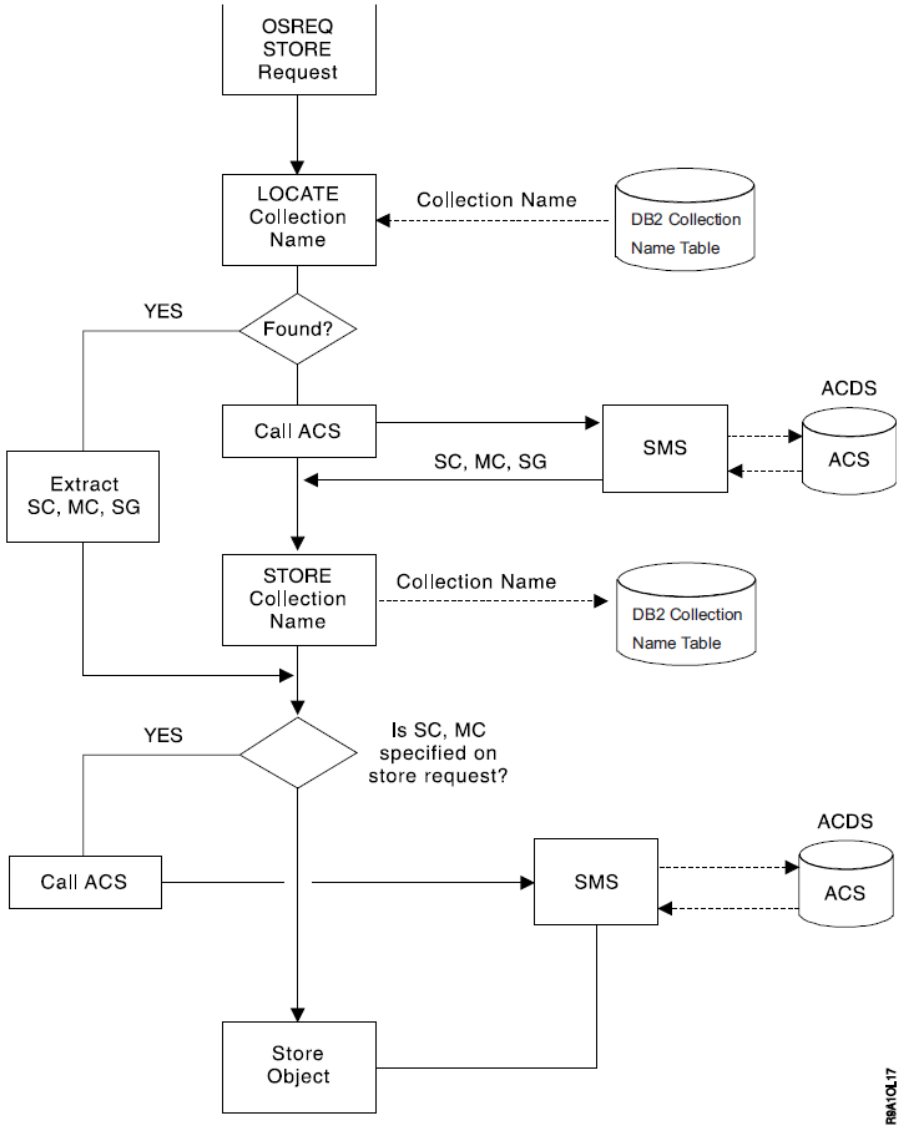

DB2 BIND Diagnostics : SQL Code –805

- Surfaced when the **last 2 bytes** of OSREQ reason code x'74xyz' are x'FCDB'

- Indicates a BIND issue with either:
  1) One of the OAM binds (CBRPBIND, CBRHBIND, CBRABIND, CBRIBIND)
  2) Application binds

- Steps to resolve:
  1) If OSREQ reason code surfaced in application, use TSO OSREQ command to issue same function against that ***collection name object name***
     a) If this does NOT surface the same reason code and ends successfully, the issue is with the application binds
     b) If this does surface the same reason code, the issue is with the OAM binds

DB2 BIND Diagnositics: SQL Code –805 (cont.)

2) Steps to resolve (cont.)
   a) Re-run the OAM binds applicable in your installation
   b) If any of the steps end with an **RC8 or higher**, review and correct any errors and re-run the binds
   c) If the steps end with an **RC0 or RC4**, re-run the TSO OSREQ command
   d) If the command still surfaces the same reason code
      1) Compare bind jobs to those supplied in SYS1.PARMLIB(CBRxBIND) members
      2) Modify your custom bind jobs if CBRxBIND jobs indicate changes not applied to your custom bind jobs
      3) Verify correct DBRM library is pointed to in **DBRMLIB DD** statement in bind jobs
      4) If any recent OAM maintenance applied and non-standard methods were used to apply maintenance (such as copying load modules manually), check to make sure all modules/DBRMs were copied successfully
   e) If the issue is still present, open a PMR with OAM Q&A to assist

OSMC Cycle Running Longer than Normal

- OSMC finishing in 2x to 3x or more time than normal
- Review CBR9370I messages surfaced
- Example:

```
06.00.50 STC85422  CBR9370I OSMC Detail for GROUP00: 190    190
     190                          READ      READ      READ      READ
     190                         DISK1     DISK2       OPT      TAPE
     190              WORK Q:        0         0         0         0
     190              WAIT Q:                   0         0         0
     190              DONE:      14452         0         0         0
     190                         WRITE     WRITE     WRITE     WRITE     WRITE
     190                         DISK1     DISK2       OPT     TAPE1     TAPE2
     190              WORK Q:        0         0         0         0         0
     190              WAIT Q:                   0         0         0         0
     190              DONE:          0         0         0     14452         0
     190                         WRITE     WRITE       DIR
     190                       BACKUP1   BACKUP2     UPDTS
     190              WORK Q:        0         0         0
     190              WAIT Q:        0         0         0
     190              DONE:          0         0    652056
     190              End of Display Detail
 06.00.51 STC85422  CBR9048I Storage Group GROUP00 has successfully
                    completed processing.
```

Indicates many objects with ODPENDDT <= current date only had directory entry updated

OSMC Cycle Running Longer than Normal (cont.)

- Why may this indicate a problem?
  - OSMC selects for processing every object in the Object Storage Group with a pending action date (ODPENDDT) <= the date of the OSMC execution
  - For every object selected, OAM will determine if object should be expired, backed up, or transitioned
  - If an object still has an **ODPENDDT <= the current date** after OSMC has completed processing, it will be selected during subsequent OSMC cycle

- The root cause
  - The last management class an object is designed to have has **transition criteria** set
    - Transition criteria should only be set if the object should be assigned a new storage class or management class
  - ACS routines incorrectly set up
    - Results in new management class not being assigned during OSMC processing as was intended

OSMC Cycle Running Longer than Normal (cont.)

- How to verify if the issue is present
  1. Issue the following SPUFI on the object storage group
     SELECT * FROM OAMADMIN.CBR_MGT_CLASS_TBL C,
     hlq.OSM_OBJ_DIR D
     WHERE ODPENDDT <= date the last OSMC run finished (format 'yyyy-mm-dd')
     AND D.ODMCNUM = C.ODMCNUM;
  2. Review the output to determine which management class (**ODCLMCNM**) the objects are associated with
  3. Review the settings for this management class in the **ACTIVE** SMS configuration
  4. If transition criteria are specified, this will cause the object to be selected during the subsequent OSMC execution of this object storage group

Prevent CPU churn!

OSMC Cycle Running Longer than Normal (cont.)

- Other possibilities
  - CBRHADUX not allowing expiration even though the management class setting (or explicit expiration date) determine objects should expire
  - SETOAM DEMOUNTWAITTIME(seconds) may need to be increased if there is constant demount/mount activity of same volume
  - MAXTAPERETREIVETASKS tasks not being set efficiently can cause starvation among storage groups
  - DB2 may be using tablespace scan versus index scan. If large changes in the number of entries in an Object Directory Table, execute RUNSTATs and the OAM binds

Recommendation: Use D SMS,OSMC,TASK(taskname) command in 5 minute intervals to surface CBR9370I summary message

OSREQ Query (or RETRIEVE) on DASD taking too long or SYSZTIOT contention

- **QB=x** – Specifies whether an OSREQ QUERY request results in a call into the OAM address space to retrieve the backup retrieval order keys. Specified on IEFSSNxx parmlib member.
  - **QB=Y** – OSREQ QUERY requests result in a call into the OAM address space for each backup copy. **This is the default.**
  - **QB=N** – OSREQ QUERY requests do not result in a call into the OAM address space for each backup copy
- If **QB=Y** and a backup copy is present, an **ENQ on SYSZTIOT** is required, which impacts OPEN/CLOSE/ALLOCATION/DEALLOCATION processing for tape
- Some applications issue a QUERY to acquire object length before doing the actual RETREIVE. If the default, QB=Y, is active, and a backup copy is present, this may cause unneeded, additional work.

Check with the application invoking OSREQ to determine if the additional information is needed

Out of Synch Collection vs Catalog Entries

Collection and Catalog changes

- With our V2R3 support and with co-existence APAR OA51129 for V2R1 and V2R2, we have removed the dependency on the catalog for OAM collection entry processing
- Collections will now only be maintained in DB2
- This eliminates future out of synch conditions and OAM having to maintain the information in two places, however:

****Warning** -** This will surface any existing out of synch conditions that you currently have, which would have been caused by SPUFI deleting entries in the DB2 collection table

****Warning** -** Do not delete catalog collection entries that belong to OAM

****Warning** -** If modifications need to be made to OAM's DB2 tables, ensure that DB2 image copies are made of those tables

V2R1 and V2R2 environments without OA51129 applied

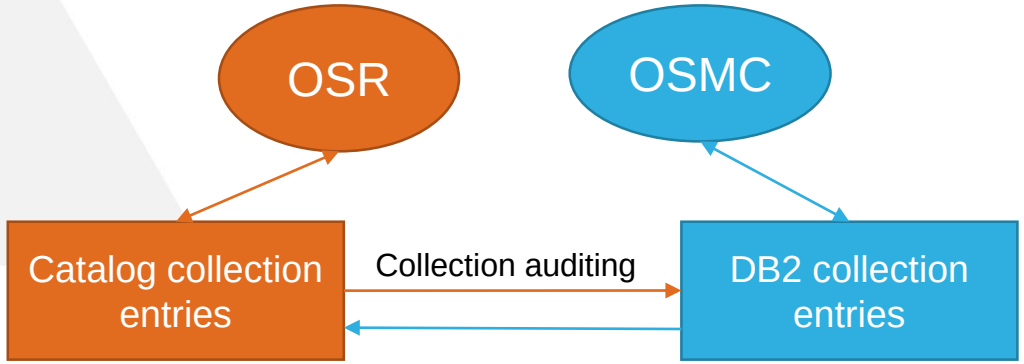

V2R3 or V2R1/V2R2 with OA51129 applied

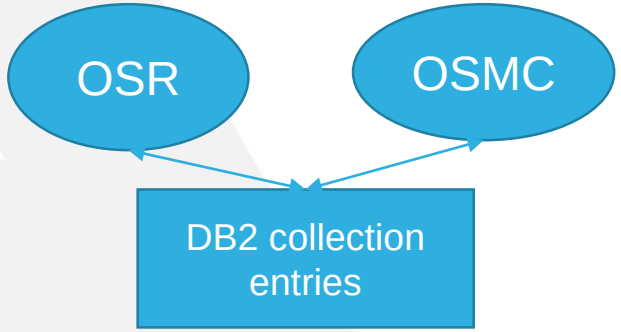

Out of Synch Collection vs Catalog Entries – Information APAR II14842

Information APAR II14842

- To help with identifying out of synch conditions as well as provide procedures and on what to do if this condition exists, we have created informational APAR II14842
- This APAR contains explicit instructions on how to identify this condition as well as step by step instructions on what to do if you find yourself in any of these of out of synch cases
- We have also included a new catalog vs collection entry automation tool called CATDB2CP (downloadable via the APAR) that is capable of determining these out of synch conditions and returns back to you vital information needed to get things sorted out
  - Note: For users that choose not to use the automation tool, we do provide manual step by step instructions for identifying out of synch conditions as well

Out of Synch Collection vs Catalog Entries - Information APAR II14842

SHARE
EDUCATE • NETWORK • INFLUENCE

CATDB2CP

- We have identified 2 parent situations that most out of synch conditions stem from
  - Collection entries are found in catalog but are missing from DB2
  - Collection entries are found in catalog and found in DB2, but… the entry found in DB2 has a different collection ID.
- This automation tool is capable of identifying both of these cases by scanning the OAM catalog as well as querying the DB2 collection table and performing comparisons
- If the tool finds any discrepancies, it will output to you information gathered for that entry from either the catalog, DB2, or both in some cases
  - Collection ID, collection name, storage class, management class, and storage group information found for that collection entry so that remedial action can be taken to correct the issue
- The tool will come with explicit step by step instructions on how to correct the out of synch condition pertinent to your environment
  - 8 fix scenarios identified

Out of Synch Collection vs Catalog Entries - Information APAR II14842

1st Case – Collection entry found in catalog but not in DB2

(Figure 1)

```
********************************************************************** Top of Data ***********************************
The following collection entries are found in catalog but are missing from the OAM DB2 collection table.

Note: if there is any collection entry in the OAM DB2 collection table that has the same collection ID as a missing entry,
      it will be output below inside parentheses.

++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++
COLLECTION ID  COLLECTION NAME                                STORAGE CLASS  MANAGEMENT CLASS  STORAGE GROUP
-------------  ----------------------------------------------------------------------------------------------------------
      6        GROUP06                                        DASDSC         DASDTAPE          GROUP06
                  (Collection entry GROUP08
                   in DB2 with the same collection ID 6 - SC: DASDSC - MC: DASDTAPE - SG: GROUP08)

      7        GROUP07                                        DASDSC         DASDTAPE          GROUP07
                  (Collection entry GROUP09
                   in DB2 with the same collection ID 7 - SC: DASDSC - MC: DASDTAPE - SG: GROUP09)

      8        GROUP10                                        DASDSC         DASDTAPE          GROUP10

********************************************************************** Bottom of Data ********************************
```

Out of Synch Collection vs Catalog Entries - Information APAR II14842

2nd Case – Collection entry found in catalog and found in DB2 with a different collection ID.

(Figure 2)

```
********************************************************************************** Top of Data **********************************************
--------------------------------------------------------------------------------------------------------------------------------------------------
The following collection entries are found in catalog and are also found in the DB2 collection table with a DIFFERENT collection ID.

Note: if there is any collection entry in the OAM DB2 collection table that has the same collection ID as a catalog entry,
      it will be output below inside parentheses.

++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++
               COLLECTION ID  COLLECTION NAME                                  STORAGE CLASS  MANAGEMENT CLASS  STORAGE GROUP
--------------------------------------------------------------------------------------------------------------------------------------------------

Catalog entry:         1        GROUP01                                         DB2DASD        MCD01             GROUP01
DB2 entry:             2        GROUP01                                         DB2DASD        MCD01             GROUP01
--------------------------------------------------------------------------------------------------------------------------------------------------
Catalog entry:         2        GROUP01.A                                       DB2DASD        MCD01             GROUP01
DB2 entry:             3        GROUP01.A                                       DB2DASD        MCD01             GROUP01
            (Collection entry GROUP01
             in DB2 with the  same collection ID 2 - SC: DB2DASD - MC: MCD01 - SG: GROUP01)
--------------------------------------------------------------------------------------------------------------------------------------------------
Catalog entry:         3        GROUP01.B                                       DB2DASD        MCD01             GROUP01
DB2 entry:             4        GROUP01.B                                       DB2DASD        MCD01             GROUP01
            (Collection entry GROUP01.A
             in DB2 with the  same collection ID 3 - SC: DB2DASD - MC: MCD01 - SG: GROUP01)
--------------------------------------------------------------------------------------------------------------------------------------------------
```

Questions?

Please remember to take a moment to fill out your session evaluation...

DFSMS OAM (Object) Hints and Tips – Performance Tuning and Diagnostics

Session: 22421

Albert Dennes
OAM L3 Service Lead and Development
IBM Corporation
aedennes@us.ibm.com

Derek Erdmann
OAM Level 2 Team Lead
IBM Corporation

Your ***input*** greatly helps us in delivering better and more useful material for you!

The Race to 15,000!

Don't miss Thursday's Farewell Reception!

SHARE
EDUCATE • NETWORK • INFLUENCE

PERSON

PROGRAM

PROJECT

Be a part of the...

HOOPLA!

Complete your session evaluations at www.share.org/evaluation and see your progress on the leaderboard!

REFERENCES

OAM Publications

- ✓ z/OS DFSMS Object Access Method Application Programmer's Reference, SC23-6865
- ✓ z/OS DFSMS Object Access Method Planning, Installation, and Storage Administration Guide for Object Support, SC23-6866
- ✓ z/OS DFSMS Object Access Method Planning, Installation, and Storage Administration Guide for Tape Libraries, SC23-6867
- ✓ z/OS DFSMSdfp Diagnosis, SC23-6863
- ✓ z/OS System Messages Vol 4 (CBD-DMO), SA38-0671
- ✓ z/OS DFSMS Using the New Functions, SC23-6857
- ✓ z/OS DFSMSdfp Storage Administration Reference, SC23-6860
- ✓ z/OS Summary of Message and Interface Changes, SA23-2300

Trademarks and Disclaimers

AIX*	DFSMSdfp	DS6000	IBM*	MQSeries*	Redbooks*	System Storage	Tivoli*	z/OS*
BladeCenter*	DFSMSdss	DS8000*	IBM eServer	MVS	REXX	System x*	WebSphere*	zSeries*
BookManager*	DFSMShsm	Easy Tier	IBM logo*	OS/390*	RMF	System z	z10 BC	z Systems
DataPower*	DFSMSrmm	FICON*	IMS	Parallel Sysplex*	SYSREXX	System z9*	z10 EC	
DB2*	DFSORT	FlashCopy*	InfinBand*	PR/SM	RMF	System z10	z/Architecture*	
DFSMS	Domino*	HiperSockets	Language Environment*	RACF*	SYSREXX	System z10 Business Class	zEnterprise*	

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

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.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

Java and all Java based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and

Linux is a registered 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.

OpenStack is a trademark of OpenStack LLC. The OpenStack trademark policy is available on the OpenStack website.

TEALEAF is a registered trademark of Tealeaf, an IBM Company.

Windows Server and the Windows logo are trademarks of the Microsoft group of countries.

Worklight is a trademark or registered trademark of Worklight, an IBM Company.

UNIX is a registered trademark of The Open Group in the United States and other countries.

* Other product and service names might be trademarks of IBM or other companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

This information provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g, zIIPs, zAAPs, and IFLs) ("SEs"). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the “Authorized Use Table for IBM Machines” provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html (“AUT”). No other workload processing is authorized for execution on an SE. IBM offers SE at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

Trademarks and Disclaimers

NOTES:

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices are suggested US list prices and are subject to change without notice. Starting price may not include a hard drive, operating system or other features. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Any proposed use of claims in this presentation outside of the United States must be reviewed by local IBM country counsel prior to such use.

The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM makes no representation or warranty regarding third-party products or services including those designated as ServerProven, ClusterProven or BladeCenter Interoperability Program products. Support for these third-party (non-IBM) products is provided by non-IBM Manufacturers.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.