C46

Evaluating Performance of the MVS Logger in a CICS Environment

Jim Grauel

Las Vegas, NV **Feb. 10 - Feb. 14, 2003**

Agenda

- Logstream sizing
- Tools
  - DFH0STAT
  - RMF data
  - SMF 88 records
  - SMF 74 records
  - CICS Performance Analyzer
- Test case 1 - tuning using SMF88 data
- Test case 2 - tuning using SMF and RMF data
- Appendix
  - ►SIZING formulas
  - ►DFHLSCU Sample
  - ►IXCMIAPU Samples
  - ►logstream definition
  - ►list logstream
  - ►RMF Samples
  - ►Displays

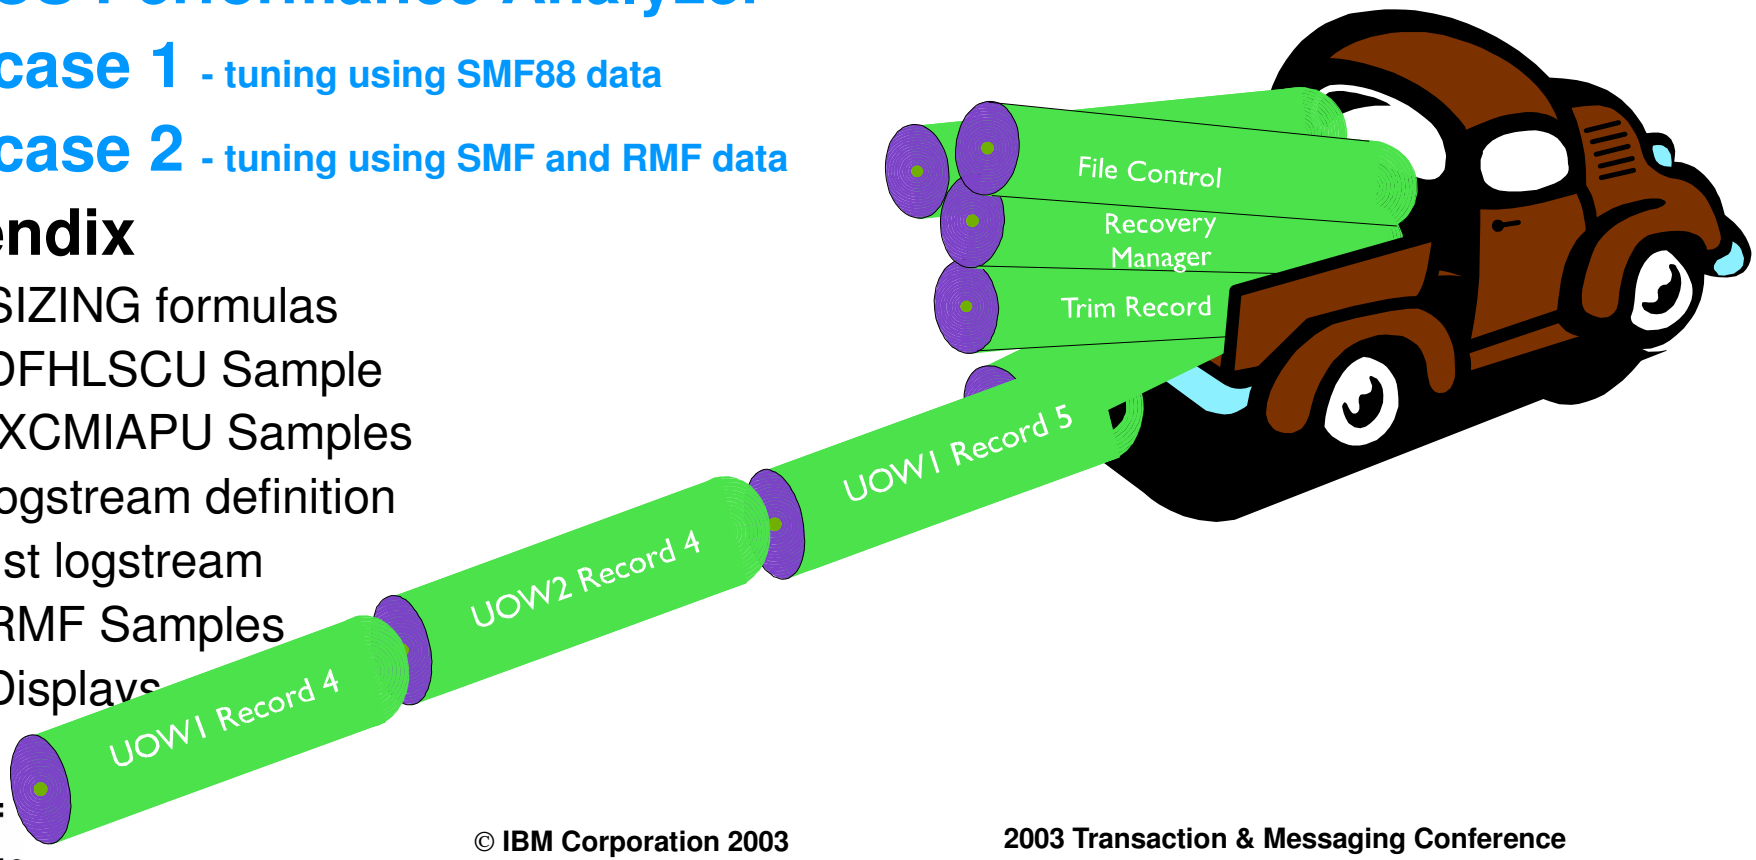

CICS Parameters

- **AKPFREQ**
  - defines the number of records added to the CICS log buffer before an activity keypoint is initiated
  - use CEMT I SYS to view/change
  - Activity KeyPoint
    - records CICS resources on the log as in R410
    - log tail management is initiated for DFHLOG and DFHSHUNT
    - if the oldest records on the log are no longer of interest to CICS they are logically deleted using an IXGDELET call to the MVS logger
    - physical deletion of the 'deleted' data happens during the offload process
- **LGDFINT**
  - specifies the 'log defer' interval
    - the length of time to delay a forced journal write before calling MVS
    - allows coat-tailing of requests
      - i.e. additional records are written in the buffer
    - default value changed to 5 milliseconds in CICS Transaction Server R2.2
  - use CEMT I SYS to view/change

Sizing

- **Structures**
  - **Large enough to hold the sum of data for connected logstreams**
- **Logstreams**
  - **Each CICS system logstream will require enough storage to hold data written in an AKP interval + the duration of the longest UOW**
  - **Control Information (CF logstreams)**
- **Size calculations**
  - **Manual**
  - **DFHLSCU**
    - ➤using pre-CTS journals
    - ➤use worst case day (i.e. heaviest activity)
    - ➤PQ13125 should be applied
    - ➤PQ34671 will correct the HIGHOFFLOAD and LOWOFFLOAD recommendations
    - ➤PQ70064 provides guidance in specifying LS_SIZE <<<<< NEW

Sizing ...

- **Considerations for calculating logstream size**
  - **Number of write requests (LGSWRITES*) in the interval**
  - **Rate of I/O**
    - ➤ LGSWRITES/interval in seconds
  - **Number of bytes written in interval (LGSBYTES*)**
  - **Number of bytes written per I/O**
    - ➤ bytes written/number of write requests (LGSBYTES/LGSWRITES)
  - **HIGHOFFLOAD percentage**
  - **Number of AKPs in the interval (LGSDELETES*)**
  - **Number of offloads during interval**
    - ➤ taken from SMF 88 records
  - **Duration of offload**
    - ➤ calculated based on CTRACE
      - the WOW entries
      - a sample is provided in the appendix

*** reported in the CICS Log Stream Resource Statistics**

MVS Coupling Facility Sizer

- a web based tool for sizing structures
  - IBM Poughkeepsie
    - www.s390.ibm.com/pso
      http://www-1.ibm.com/servers/eserver/zseries/cfsizer/
  - provides an easy to use interface to calculate the structure sizes based on minimum input
  - information from the CICS Transaction Server statistics is used as input

IXCMIAPU - Display of DFHLOG

```
LOGSTREAM NAME(IYOT1.DFHLOG) STRUCTNAME() LS_DATACLAS(LS10MEG)
          LS_MGMTCLAS() LS_STORCLAS() HLQ(GRAUEL) MODEL(NO) LS_SIZE(100) STG_MGMTCLAS() STG_STORCLAS()
          STG_DATACLAS() STG_SIZE(1500) LOWOFFLOAD(19) HIGHOFFLOAD(95) STG_DUPLEX(YES) DUPLEXMODE(UNCOND)
          RMNAME() DESCRIPTION() RETPD(0) AUTODELETE(NO) DASDONLY(YES) DIAG(YES) MAXBUFSIZE(65532)

          LOG STREAM ATTRIBUTES:
          User Data:
          00000000000000000000000000000000000000000000000000000000000000000000

          LOG STREAM CONNECTION INFO:
          SYSTEMS CONNECTED: 1
          SYSTEM    STRUCTURE          CON CONNECTION  CONNECTION
          NAME      VERSION            ID  VERSION     STATE
          --------  -----------------  --  ----------  ----------
          MV55      0000000000000000   00  00000000    N/A

          LOG STREAM DATA SET INFO:
             DATA SET NAMES IN USE: GRAUEL.IYOT1.DFHLOG.<SEQ#>
              Ext.   <SEQ#>    Lowest Blockid    Highest GMT         Highest Local       Status
             -----  --------  ----------------  ------------------  ------------------  --------------
            *00001  A0000004  00000000002846D0  01/16/00 22:33:04  01/16/00 22:33:04
                    A0000005  00000000002F0675  01/16/00 22:33:15  01/16/00 22:33:15  CURRENT
                    A0000006  0000000000000000
              NUMBER OF DATA SETS IN LOG STREAM: 3

          POSSIBLE ORPHANED LOG STREAM DATA SETS:
           NUMBER OF POSSIBLE ORPHANED LOG STREAM DATA SETS: 0

          DSLIST - Data Sets Matching GRAUEL.IYOT1.DFHLOG                 Row 1 of 8
            Command - Enter "/" to select action                Message          Volume
            ------------------------------------------------------------------------------
                     GRAUEL.IYOT1.DFHLOG.A0000004                                *VSAM*
                     GRAUEL.IYOT1.DFHLOG.A0000004.DATA                           PBDA14
                     GRAUEL.IYOT1.DFHLOG.A0000005                                *VSAM*
                     GRAUEL.IYOT1.DFHLOG.A0000005.DATA                           PBDA22
                     GRAUEL.IYOT1.DFHLOG.A0000006                                *VSAM*
                     GRAUEL.IYOT1.DFHLOG.A0000006.DATA                           PBDA21
                     GRAUEL.IYOT1.DFHLOG.PLEXB                                   *VSAM*
                     GRAUEL.IYOT1.DFHLOG.PLEXB.DATA                              PBDA07
```

NOTE: LOG (offload) datasets should *NOT* be deleted unless they are listed as 'orphaned'

DFH0STAT

Run 8- Applid IYOT1 Sysid JIM Jobname IYOT1 Date 01/16/2000 Time 22:34:29 CICS 5.3.0

System Status

MVS Product Name. . . . : MVS/SP6.0.8

Activity Keypoint Frequency. . . . . : 4,000
Logstream Deferred Force Interval. . . . : 30

Logstream Name	Use Count	Status	Sys Log	Structure Name	Max Block Length	DASD Only	Retention Period	Auto Delete	Stream Deletes	Browse Starts	Browse Reads
IYOT1.DFHLOG	1	OK	YES		65,532	YES	0	NO	3	12	0
IYOT1.DFHSHUNT	1	OK	YES	LOG_JG	64,000	NO	0	NO	1	0	0

Logstream Name	Write Requests	Bytes Written	Average Bytes	Buffer Appends	Buffer Full Waits	Force Waits	Current Waiters	Peak Waiters	Retry Errors
IYOT1.DFHLOG	11,218	3,494,503	311	11,734	0	17	0	1	5
IYOT1.DFHSHUNT	0	0	0	0	0	0	0	0	0

Run 10- Applid IYOT1 Sysid JIM Jobname IYOT1 Date 01/16/2000 Time 23:42:38 CICS 5.3.0

Logstream Name	Use Count	Status	Sys Log	Structure Name	Max Block Length	DASD Only	Retention Period	Auto Delete	Stream Deletes	Browse Starts	Browse Reads
IYOT1.DFHLOG	1	OK	YES		65,532	YES	0	NO	13	22	0
IYOT1.DFHSHUNT	1	OK	YES	LOG_JG	64,000	NO	0	NO	1	0	0

Logstream Name	Write Requests	Bytes Written	Average Bytes	Buffer Appends	Buffer Full Waits	Force Waits	Current Waiters	Peak Waiters	Retry Errors
IYOT1.DFHLOG	11,235	4,356,569	387	12,200	11	22	0	1	1
IYOT1.DFHSHUNT	0	0	0	0	0	0	0	0	0

DFH0STAT notes

```
Applid IYOT1     Sysid JIM   Jobname IYOT1       Date 01/30/2000  Time 01:56:47                CICS 5.3.0
```

System Status

```
  MVS Product Name. . . . :  MVS/SP6.0.8
  Activity Keypoint Frequency. . . . . . . :       4,000
  Logstream Deferred Force Interval. . . . :          30
```

Logstream Name	Use count	Status	Sys Log	Structure Name	Max Block Length	DASD Only	Retention Period	Auto Delete	Stream Deletes	Browse Starts	Browse Reads
IYOT1.DFHLOG	1	OK	YES		65,532	YES	0	NO	3	34	0
IYOT1.DFHSHUNT	1	OK	YES	LOG_JG	64,000	NO	0	NO	1	0	0
IYOT1.J02	1	OK	NO		65,532	YES	0	NO	N/A	N/A	N/A

Logstream Name	Write Requests	Bytes Written	Average Bytes	Buffer Appends	Buffer Full Waits	Force Waits	Current Waiters	Peak Waiters	Retry Errors
IYOT1.DFHLOG	9,023	2,546,491	282	9,037	0	2	0	1	1
IYOT1.DFHSHUNT	0	0	0	0	0	0	0	0	0
IYOT1.J02	14	916,800	65,485	9,009	0	0	0	0	0

DFH0STAT is supplied as a sample COBOL program in CICS.SDFHSAMP. It contains self-documenting source code to be compiled and run as a transaction to collect CICS statistics and write them to the JES spool. The output can then be viewed under TSO. The SIT parm SPOOL=YES is required.

A compiled version is available in CICS Transaction Server R2.2.

As shown in the example above, there are a number of interesting statistics produced for logstreams. This same information is available in the CICS shutdown statistics.

Notice in the System Status the OS/390 release is provided (in this case OS/390 R2.8). But of greater importance are the values for Activity Keypoint Frequency (AKPFREQ) and Logstream Deferred Force Interval (LGDFINT).

In the logstream statistics we see each logstream connected to this CICS region. If the logstream is contained in a Coupling Facility (CF) structure the structure name is given; if it is a DASDONLY logstream, the structure name is blank.

In the example, IYOT1.DFHLOG and user journal J02 are DASDONLY logstreams, while IYOT1.DFHSHUNT is a CF logstream connected to structure LOG_JG.

DFH0STAT notes ...

The values under MAX BLOCK Length are worth noting. This value originates in the logstream or structure definition and is returned to CICS when it connects to the logstream. For a CF logstream the blocksize is specified as MAXBUFSIZE on the structure definition. The value specified in MAXBUFSIZE determines the element size for the logstreams in the structure. If the MAXBUFSIZE is specified equal to or less than 65276, the element size is 256; if greater than 65276 the element size is set to 512.

For DASDONLY logstreams, MAXBUFSIZE may be specified on the logstream definition. MAXBUFSIZE defines the largest block that can be written to the logstream. The default value is 65532.

In either case, the MAXBUFSIZE is returned to CICS and is used to determine the CICS logstream buffer size. Note, for user journals, unless the application uses the wait option, the IXGWRITE call is issued when the buffer fills. Refer to the average bytes on J02. This might be a reason to reduce the MAXBUFSIZE on a user journal.

The value given under Stream Deletes is the number of times CICS issued an IXGDELET call to the logger for log tail deletion.

The value under Browse Starts is a count of the number of times a browse start request is issued. You may see some system logstreams with a large value in a low volume system. CICS uses a Browse Start to verify the logger is still operational. In CICS Transaction Server R2.2, the MVS Logger function CHECK_CONNECTION_STATUS is used to verify the logger is operational.

The number of Write Requests is the number of times CICS calls the MVS logger for an IXGWRITE. The number of Buffer Appends may be larger than the number of Write Requests due to calls to the CICS logger domain, which do not include the force option.

Buffer Full and Force waits can be an indication there is a delay in I/O processing. This can also be an indicator the log defer interval is too large. If you consistently see numbers for either of these conditions, the value for LGDFINT may be reduced from the default of 30 to 5. Do not set it to 0. In addition, CF service time (for CF logstreams) or DASD I/O time should be investigated.

Retry Errors is a count of MVS logger errors which have been retried. An example would be 868 errors returned while the staging dataset is being formatted. This can happen with DASDONLY logstreams or if staging datasets are used with a CF logstream. For example:

```
01.48.56 JOB07716  +DFHLG0777 IYOT1
                   A temporary error condition occurred during MVS logger operation IXGWRITE  for log
                   stream IYOT1.DFHLOG. MVS logger codes: X'00000008', X'00000868'.
```

DFH0STAT

```
Applid IYOT1     Sysid IYO1  Jobname IYOT122     Date 08/11/2002  Time 13:50:13        CICS 6.2.0      PAGE    10
Logstream - System Logs
System log - DFHLOG
  Logstream Name . . . . . . . . . . . . . . . : IYOT1.CICS22.DFHLOG        Logstream Status . . . . :          OK
  DASD Only. . . . . . . . . . . . . . . . . . : YES                       Retention Period (days). :           0
  Coupling Facility Structure Name . . . . . :                             Auto Delete. . . . . . . :          NO
  Logstream Writes . . . . . . . . . . . . . :      18,788                 Maximum Block Length . . :      65,532
  Logstream Writes per second. . . . . . . . :       14.74
  Average Bytes per Logstream Write. . . . . :       1,289
  Logstream Deletes (Tail Trims) . . . . . . :         105
  Logstream Query Requests . . . . . . . . . :           3
  Logstream Browse Starts. . . . . . . . . . :           0
  Logstream Browse Reads . . . . . . . . . . :           0
  Logstream Buffer Appends . . . . . . . . . :     104,811
  Logstream Buffer Full Waits. . . . . . . . :           0
  Logstream Force Waits. . . . . . . . . . . :      93,764                 Logstream Current Force Waiters:       3
  Logstream Retry Errors . . . . . . . . . . :           1                 Logstream Peak Force Waiters:         11

Applid IYOT1     Sysid IYO1  Jobname IYOT122     Date 08/12/2002  Time 08:31:13        CICS 6.2.0      PAGE    11
Logstreams - Resource
                          Use          Sys                      Max Block DASD  Retention   Auto   Stream   Browse   Browse
Logstream Name          Count  Status  Log  Structure Name        Length  Only   Period    Delete  Deletes  Starts    Reads
IYOT1.CICS22.DFHLOG         1  OK      YES  LOG_JG_20M           64,000  NO          0    NO        1,896       0        0
IYOT1.CICS22.DFHSHUNT       1  OK      YES  LOG_JG_20M           64,000  NO          0    NO            1       0        0
Logstreams - Requests
                          Write                        Average    Buffer      Buffer     Force   Current    Peak   Retry
Logstream Name          Requests   Bytes Written        Bytes    Appends   Full Waits   Waits   Waiters  Waiters  Errors
IYOT1.CICS22.DFHLOG      445,574     443,483,499          995  1,895,148            0  428,880         1       10       1
IYOT1.CICS22.DFHSHUNT          0               0            0          0            0        0         0        0       0
```

DFH0STAT notes

With CICS Transaction Server R2.2, DFH0STAT has been enhanced. The format has been changed for ease of reading and information has been added to calculate (and display) the number of logstream writes per second.

```
System log - DFHLOG
___________________

  Logstream Name . . . . . . . . . . . . . . : IYOT1.CICS22.DFHLOG      Logstream Status . . . . : OK
  DASD Only. . . . . . . . . . . . . . . . . : YES                      Retention Period (days). : 0
  Coupling Facility Structure Name . . . . . :                          Auto Delete. . . . . . . : NO
  Logstream Writes . . . . . . . . . . . . . :      18,788              Maximum Block Length . . : 65,532
  Logstream Writes per second. . . . . . . . :       14.74
  Average Bytes per Logstream Write. . . . . :       1,289
  Logstream Deletes (Tail Trims) . . . . . . :         105
  Logstream Query Requests . . . . . . . . . :           3
  Logstream Browse Starts. . . . . . . . . . :           0
  Logstream Browse Reads . . . . . . . . . . :           0
  Logstream Buffer Appends . . . . . . . . . :     104,811
  Logstream Buffer Full Waits. . . . . . . . :           0
  Logstream Force Waits. . . . . . . . . . . :      93,764   Logstream Current Force Waiters:   3
  Logstream Retry Errors . . . . . . . . . . :           1   Logstream Peak Force Waiters:      11
```

Monitoring the System Logger

- SMF Type 88 Records
  - **Written periodically**
    - ➤ based on the SMF88 reporting interval
  - **Written at Disconnect time**
  - **Provide Assistance with Tuning**
  - **SMF 88 Subtype 1 records used for Logstream tuning**
    - ➤ interim storage usage
    - ➤ data set switches
  - **SMF 88 Subtype 11 records used for Structure tuning**
    - ➤ Dynamic adjustment of structure entry to element ratio
    - ➤ New as of OS/390 R1.3
  - **See Macro IXGSMF88 for details**
  - **See IXGRPT1 in SYS1.SAMPLIB for generating a report**
    - ➤ supplied in PL/I only
    - ➤ OW28861 improves the formatting
    - ➤ OW36423 provides IXGRPT1J and IXGRPT1L
      - PL/I no longer required

SMF 88 record fields

- **BYT Written by users IXGWRITES (SMF88LWB)**
- **BYT Written to interim storage (SMF88SWB)**
  - amount of storage used in the "interim" storage portion of the logstream.
- **BYT Written to DASD (invoked) (SMF88LDB)**
  - bytes written to the offload dataset -- this data is not rounded to a CI boundary
- **# Writes invoked (SMF88LWI)**
- **BYT Deleted interim ST w/o DASD (SMF88SIB)**
  - due to CICS tail trimming
- **# Deletes w/o DASD write (SMF88SII)**
  - times data deleted from interim storage and the data had *not* been offloaded
- **BYT Deleted interim ST w/DASD (SMF88SAB)**
  - during the offload process, physically deleting the logically deleted data was not enough to reduce the logstream to the LOWOFFLOAD value
- **# Deletes w/write (SMF88SAI)**
  - times data deleted from interim storage where the data had been offloaded
- **# Writes Completed - applies to CF structures**
  - **TYPE1 (SMF88SC1)**
    - the number of writes completed normally
  - **TYPE2 (SMF88SC2)**
    - the number of writes completed while offload in progress
  - **TYPE3 (SMF88SC3)**
    - the number of writes completed with **90% of elements** for the logstream are in use

SMF 88 record fields...

- **Average Buffer size**
  - average size of the data written in the interval
- **Event**
  - Offload (SMF88EO)
    - number of offloads in the interval
  - DASD Shift (SMF88EDS)
    - number of times an additional log dataset is allocated during offload
  - STRC Full (SMF88ESF)
    - number of times a structure full condition was reached
  - Ntry full (SMF88EFS)
    - number of offloads (all logstreams) due to structure reaching 90% entry full
  - STG THLD (SMF88ETT)
    - number of times HIGHOFFLOAD percentage was reached in the staging dataset
  - ST Full (SMF88ETF)
    - number of times staging dataset full
  - Rebuild (SMF88ERI)
    - number of structure rebuilds

SMF 88 field notes

SMF88LWB **"BYT WRITTN BY USERS IXGWRITES"**

This is the actual IXGWRITE "blocklen" summation over the SMF interval, it does not contain any "hot air" assumptions or rounding up to the interim storage medium's storage method boundary.

The computation is the same for CF structure based and DASDONLY logstreams.

SMF88SWB **"BYT WRITTEN TO INTERIM STORAGE"**

This value contains the summation of the amount of space in the interim storage needed to hold the data written (SMF88LWB), it helps to reveal the amount of "hot air" space used to hold the logstream input data.

For CF structure based logstreams, this summation value includes the amount of element space, rounded up to the element boundary, (either 256 or 512 byte) for each IXGWRITE.

For DASDONLY based logstreams, this summation value includes the amount of staging dataset control interval (CI) space (rounded up to the CI boundary (4096 bytes) for each IXGWRITE, "hot air") required to hold the log data written into the logstream.

SMF88SIB **"BYT DELETD INTERIM ST W/O DASD"**

This value contains the summation of the amount of space deleted from the interim storage, including the "hot air", where the data had not been moved to an offload dataset. This value should be close to SMF88SWB.

For CF structure based logstreams, this summation value includes the amount of element space, rounded up to the element boundary for each IXGWRITE. The data is deleted from the CF structure without moving the data to an offload dataset.

For DASDONLY based logstreams, this summation value includes the amount of staging data set control interval (CI) space rounded up to the CI boundary for each IXGWRITE. The data is deleted from the logstream staging data set without moving the data to an offload data set.

SMF88SAB **"BYT DELETD INTERIM ST W/DASD"**

This value contains the summation of the amount of space deleted from the interim storage including the "hot air" , after the log data was moved (offloaded) to DASD. It is related to SMF88SWB.

For CF structure based logstreams, this summation value includes the amount of element space, rounded up to the element boundary for each IXGWRITE. The log data has been moved from the CF structure to an offload dataset.

For DASD Only based logstreams, this summation value includes the amount of staging data set control interval (CI) space, rounded up to the CI boundary for each IXGWRITE. The log data has been moved to an offload data set.

It's worth noting that in this case, logger does not read the staging dataset log data copy, it reads the local buffer (dataspace) copy and writes it to the offload data set. However, the SMF88SAB value represents the amount of space used in the staging data set since this resource is of limited size.

SMF 88 notes

Determine the largest average buffer size for logstreams in the structure. This value is used to determine the entry to element ratio for a CF logstream. The element size is determined by the MAXBUFSIZE value specified when the structure is defined (using IXCMIAPU). When the value is less than 65276, the element size is 256. If the value specified is greater than 65276, the element size is 512.

If the element size is 256 and the write is for 734 bytes, 3 elements are required.

The entry to element ratio is dynamically adjusted based on a current snapshot of the usage of all logstreams connected to the structure. The snapshot will be taken about every 30 minutes and the ratio adjusted if necessary. If there are 3 logstreams in the structure, with average buffer sizes of 1800, 1450, and 734, logger will use the current real time average (how many records of each size) to define the entry to element ratio. If most of the records are in the 1800 range the average would be 1:8. Add 4 bytes for logger data, (1800+4)/256 = 7.046 resulting in a ratio of 1:8.

This means we expect to use 8 elements for each entry. When shorter records are written to the logstream with an average size of 734, (for example 200 bytes) we still use 1 entry, but only 1 of the seven assumed elements. The net effect is more entries are used than predicted leading to an entry full condition before HIGHOFFLOAD can be reached. Offload will be triggered, but there is no extra room to write additional records until some space is recovered. In this situation the NTRY FULL count will be equal to or greater than the number of offloads.

The important point to remember is a logstream performs best when it is contained in a structure where all logstreams have like characteristics (i.e. average buffer size and amount of data written).

Also, remember writes greater than 4K are written asynchronously which can increase response time.
Refer to Washington System Center Flash 10159 for additional information.

Another point of interest is a comparison of SMF88SWB (bytes written to the "primary" or "interim" storage) with SMF88LWB (bytes requested to be written on the IXGWRITE calls). This will give you an idea of how much gas is in the logstream. For CF logstreams, SMF88SWB represents the full amount of element storage, up to the 256 or 512 byte element boundary required. For DASD only logstreams, SMF88SWB represents the full amount of Control Interval storage used in the staging data set to hold the data written. Note, if the same data were written to a DASD only logstream and a CF logstream, SMF88SWB will normally be larger for the DASD only.

SMF 88 notes

Using the SMF 88 report produced via IXGRPT1:

NOTE: OW36423 should be applied, - it increases several counter sizes and removes the dependency on PL/1 runtime libraries

1. For DFHLOG and DFHSHUNT the number of "BYT DELETD INTERIM ST W/O DASD" should be very close to the "BYT WRITTN TO INTERIM STORAGE". A value in the "BYT DELETD INTERIM ST W/DASD" indicates data is being offloaded and then deleted, costing additional processing and I/O. The BYT WRITTN to DASD (INVOKED) should be very low.

Factors:

- **Long running CICS tasks**
  - **> this is not average response time, but how long the tasks which use recoverable resources are in the system and causing log records to be written during each activity keypoint interval.**
  - **> if message DFHLG0743 is not being issued for DFHLOG with each activity keypoint, a long running task is preventing tail trimming. However, it is not unusual to see infrequent DFHLG0743 messages for DFHSHUNT. Units of work may exist on DFHSHUNT for extended periods of time. Examples are conversational tasks which have updated a recoverable resource, and mirror tasks awaiting a forget flow from the connected region. In CICS Transaction Server the forget flow (part of 2 phase commit) is carried with the next task attach to flow across the MRO link. This improves performance on most links, but if the usage is low, the log records for the mirror may reside on DFHSHUNT for an extended period of time. A DFHLG0743 message being issued for DFHSHUNT indicates units of work, which had been inactive, have completed. APARs PQ22563 and PQ14796 reduce data recorded on DFHSHUNT. PQ56315 (R1.3) and PQ56341 (R2.2) should be applied to correct the 'forget flow' problem.**
- **AKPFREQ is set too high**
  - **> use the DFHRM0205 messages to determine how often an activity keypoint is being taken.**
- **Interim storage**
  - **> CF logstream:**
    - **structure size is too small, or staging dataset cannot hold as much data as the logstream in the structure.**
    - **incorrect entry : element ratio - this can happen when unlike logstreams are in the same structure; the ratio is based on the worst (largest) bufsize. Dynamic changes will happen on no less than 30 minute intervals.**
  - **> DASDONLY: the allocation size of the staging dataset may be insufficient.**
- **HIGHOFFLOAD should be set no higher than 85%**
- **LOWOFFLOAD should be set in the range of 40 - 60 %**
  - **> For user journals, all data should be offloaded each time an offload is initiated**
    - **HIGHOFFLOAD should be set no higher than 85%**
    - **LOWOFFLOAD should be set to 0**

SMF 88 notes ...

2. Under # WRITES COMPLETED (note this is for CF logstreams only)
   TYPE1 -- normal - this number should be high
   TYPE2 -- normal, but the logstream is over the HIGHOFFLOAD value
   TYPE3 -- writes issued when 90% of the *elements* for the logstream are in use
   - look for HIGHOFFLOAD set to 90% or higher
   - tail trimming not happening (see #1 above)
   - CICS is filling the space above HIGHOFFLOAD point faster than the logstream is being offloaded

3. Under EVENTS
   - NTRY FULL indicates the number of times all logstreams in the structure were offloaded due to reaching 90% of the structure *entries* in use.
     - this could be the result of the entry to element ratio being too large, or a poorly behaving application which is causing many small records to be written to a logstream which normally contains large records.
   - OFFLOADs are good if they are being triggered by the HIGHOFFLOAD value. However, offloads are bad if they are triggered by an NTRY FULL condition. In addition, for a CF logstream, offloads should not be triggered by reaching the HIGHOFFLOAD value on the staging dataset rather than the CF logstream. (see STG THLD below).
   - DASD Shifts indicates the number of times an additional offload dataset is allocated
     - for DFHLOG and DFHSHUNT this number should be very small, otherwise too much data is being offloaded. (see item 1)
     - verify the allocation size of the offload dataset
       -- for user journals each offload dataset should be capable of holding multiple offloads of the logstream
       -- if the size has not been specified in the logstream definition (LS_SIZE) or in the SMS data class the size will be determined by either the installation ACS (Automatic Class Selection) routines or the value specified in the ALLOCxx member of SYS1.PARMLIB -- which defaults to 2 tracks.
   - STRC Full indicates the number of times a structure full condition was reached -- this should always be 0
   - STG THLD number of times the HIGHOFFLOAD percentage was reached in the staging dataset
     -- this is good for DASD only logstreams but should not happen for CF logstreams. If numbers are seen here for a CF logstream, the staging dataset needs to be increased so that it will hold at least as much data as the logstream in the structure. A value here for a CF logstream also indicates the CF is either volatile or failure dependent.
   - Rebuild indicates the number of structure rebuilds in the interval -- if this happens on a regular basis, it needs investigation.

CF LOGSTREAMS

-LOGSTREAM NAME-	STRUCTURE NAME--	BYT WRITTN BY USERS IXGWRITES	BYT WRITTN TO INTERIM STORAGE	BYT WRITTN TO DASD	#WRITES INVOKED	---# WRITES COMPLETED------ TYPE1	TYPE2	TYPE3	AVERAGE BUFFER SIZE
		BYT DELETD INTERIM ST W/O DASD	# DELETES W/O DASD WRITE	BYT DELETD INTERIM ST W/DASD	# DELETS W/ WRITE	----------EVENT---------- OFF- LOAD	DASD SHFT	STRC FULL	NTRY FULL	STG THLD	STG FULL	RE- BLD
(1) AKPFREQ 4000 NO Syncpoints LGDFINT 30 5M structure												
01/08/00 1:15:00 AM (SMF INTERVAL TIMESTAMP 'B369F6110AD00000'X)												
IYOT1.DFHLOG	LOG_JG	3124035	5508864	2659354	10411	9891		520		0		300
		0	0	2346714	7816	10	2	0	0	0	0	0
(2) AKPFREQ 4000 with Syncpoints LGDFINT 30 5M structure												
01/08/00 4:30:00 PM (SMF INTERVAL TIMESTAMP 'B36AC295C3200000'X)												
IYOT1.DFHLOG	LOG_JG	3493285	5909504	1755243	11223	10445		686		92		311
		1865798	5745	1567363	4697	10	1	0	0	0	0	0
(3) AKPFREQ 1000 with Syncpoints LGDFINT 5 5M structure												
01/08/00 5:30:00 PM (SMF INTERVAL TIMESTAMP 'B36ACFFEFD600000'X)												
IYOT1.DFHLOG	LOG_JG	4362966	6781696	0	11244	11190		54		0		388
		4305206	11043	0	0	6	0	0	0	0	0	0
(4) AKPFREQ 1000 with Syncpoints LGDFINT 5 -- 3 Regions connected to the structure 5M structure												
01/08/00 7:00:00 PM (SMF INTERVAL TIMESTAMP 'B36AE41CD4C00000'X)												
IYOT1.DFHLOG	LOG_JG	4295002	6713600	2728029	11240	10700		362		178		382
		1963968	6629	2552309	4393	34	2	0	0	0	0	0
IYOT3.DFHLOG	LOG_JG	4340339	6758144	2167961	11238	10788		428		22		386
		2040630	6883	2011881	3902	30	2	0	0	0	0	0
(5) AKPFREQ 1000 with Syncpoints LGDFINT 5 -- 3 Regions connected to the structure 20M structure												
01/10/00 11:30:00 PM (SMF INTERVAL TIMESTAMP 'B36DA43146E00000'X)												
IYOT1.DFHLOG	LOG_JG	4348827	6768128	0	11248	11216		32		0		386
		3348643	9327	0	0	3	0	0	0	0	0	0
IYOT3.DFHLOG	LOG_JG	116	256	0	1	1		0		0		116
		0	0	0	0	0	0	0	0	0	0	0

CF Logstream Notes

Evaluating the performance of logstream requires examination of the SMF 88 records produced by the logger. It also helps to understand the CICS parameters which affect logstream activity. For DFHLOG and DFHSHUNT the goal is to reduce the BYT DELETD INTERIM ST W/DASD to as close to zero as possible. However, there may be regions where circumstances cause an increase in the amount of data required for a given period during the day. The most important factor is the length of the longest unit of work, as this defines the amount of data which must be available for backout.

On the prior page is a series of runs changing one or two parms which can have a significant effect on logstream operation. All runs were made using 2 transactions which cause 5000+ log writes each. HIGHOFFLOAD is set to 85% with LOWOFFLOAD set to 50%. Please note, DFHSHUNT is not shown, due to presentation space limitation.

The first run was made with AKPFREQ set to 4000 (the default) LGDFINT set to 30 (the default before CICS Transaction Server R2.2) with a 5M structure. In this case the application did not issue syncpoints. A single structure is used for both DFHLOG and DFHSHUNT, which is not recommended, due to differences in logstream characteristics.

Since logstreams for IYOT1 (DFHLOG and DFHSHUNT) were the only logstreams connected, each was allocated about 2.5M.

The SMF88 data shows 10 offloads, 2 DASD shifts, and 2346714 bytes deleted with DASD I/O. Notice there were 7816 logstream deletes with DASD write. Also notice there were 520 TYPE2 writes, meaning 520 writes took place while an offload was in progress. A point of interest is the logger writes data to the offload dataset in 4K (4096 bytes) CIs. Dividing 2346714 by 4096 tells us it cost 573 I/O to offload the data. Note after application of OW31383 for DFSMS, offload dataset CISIZE can be specified up to 24K.

From the report, field BYT WRITTN to DASD (data written to the offload dataset), seems to be incorrect. This field reflects the actual user data which is written to the offload dataset. This is in contrast to the data written to the CF logstream which is rounded to the element size.

In run #2, the application was changed to issue syncpoints, everything else remained the same. The improvement in the number of bytes offloaded (1567363) is the result of deleting a larger number of records during the activity keypoint process. Notice there were still 10 offloads, but only 1 DASD shift. These changes are both in the correct direction. However, on the negative side there are 92 TYPE3 writes, indicating 90% of the elements for the logstream are in use.

In run #3, AKPFREQ was set to 1000 and LGDFINT was set to 5. The number of bytes offloaded was reduced to 0, and there were no DASD shifts. Note, 6 offloads occurred but, because the LOWOFFLOAD threshold was reached each time, no data was written to the offload datasets.

In run #4, the only change was to start 2 additional regions which connected to the same structure (LOG_JG) for DFHLOG and DFHSHUNT. This reduces the storage available for each logstream to about .83M (5M/6). For IYOT1.DFHLOG, we see a large amount of data being offloaded (2552309 bytes), 2 DASD shifts and 178 TYPE3 writes.

For run #5, the structure size was increased to 20M, or 3.3M per logstream. Once again we see the logstream is performing much better. The number of offloads is reduced, due to the increase in structure size. A point to remember is the LOGSNUM value defines how many logstreams may be connected to this structure. With LOGSNUM set to 10, each logstream would be reduced to 2M each.

DASDONLY LOGSTREAMS

-LOGSTREAM NAME- STRUCTURE NAME--	BYT WRITTN BY USERS IXGWRITES	BYT WRITTN TO INTERIM STORAGE	BYT WRITTN TO DASD	#WRITES INVOKED	---# WRITES COMPLETED------ TYPE1	TYPE2	TYPE3	AVERAGE BUFFER SIZE
	BYT DELETD INTERIM ST W/O DASD	# DELETES W/O DASD WRITE	BYT DELETD INTERIM ST W/DASD	# DELETS W/ WRITE	----------EVENT---------- OFF- LOAD	DASD SHFT	STRC FULL	NTRY FULL	STG THLD	STG FULL	RE- BLD
(6) AKPFREQ 4000 LGDFINT 30 STG_SIZE 2518 HIGHOFFLOAD 95 LOWOFFLOAD 19								
01/16/00 5:00:00 PM (SMF INTERVAL TIMESTAMP 'B374D83610400000'X)								
IYOT1.DFHLOG *DASDONLY*	3495557	46129152	2078345	11225	0	0	0	311
	15839232	3867	23478272	5694	5	1	0	0	280	1	0
(7) AKPFREQ 1000 LGDFINT 30 STG_SIZE 2518 HIGHOFFLOAD 95 LOWOFFLOAD 19								
01/16/00 5:30:00 PM (SMF INTERVAL TIMESTAMP 'B374DEEAAD600000'X)								
IYOT1.DFHLOG *DASDONLY*	4353623	47009792	334416	11242	0	0	0	387
	38436864	9288	692224	99	5	2	0	0	5	0	0
01/16/00 5:36:21 PM (SMF INTERVAL TIMESTAMP 'B374E056BF160400'X)								
IYOT1.DFHLOG *DASDONLY*	161602	167936	132141	4	0	0	0	40400
	7876608	1854	135168	3	1	0	0	0	0	0	0
(8) AKPFREQ 4000 LGDFINT 30 STG_SIZE 1500 HIGHOFFLOAD 95 LOWOFFLOAD 19								
01/16/00 10:45:00 PM (SMF INTERVAL TIMESTAMP 'B37525531F300000'X)								
IYOT1.DFHLOG *DASDONLY*	3494503	46100480	2628475	11222	0	0	0	311
	11407360	2785	30142464	7322	8	6	0	0	384	4	0
(9) AKPFREQ 1000 LGDFINT 30 STG_SIZE 2518 HIGHOFFLOAD 80 LOWOFFLOAD 50								
01/16/00 11:15:00 PM (SMF INTERVAL TIMESTAMP 'B3752C07BC500000'X)								
IYOT1.DFHLOG *DASDONLY*	4353874	46997504	0	11239	0	0	0	387
	46366720	11121	0	0	6	0	0	0	6	0	0
(10) AKPFREQ 1000 LGDFINT 5 STG_SIZE 3500 HIGHOFFLOAD 80 LOWOFFLOAD 50								
01/16/00 11:45:00 PM (SMF INTERVAL TIMESTAMP 'B37532BC59700000'X)								
IYOT1.DFHLOG *DASDONLY*	4356569	46985216	0	11235	0	0	0	387
	46292992	11103	0	0	4	0	0	0	4	0	0

DASDONLY Logstream notes

Evaluating the performance of a DASDONLY logstream also requires examination of the SMF 88 records produced by the logger. As with CF logstreams, it helps to understand the CICS parameters which affect logstream activity. For DFHLOG and DFHSHUNT the goal is to reduce the BYT DELETD INTERIM ST W/DASD to as close to zero as possible. However, there may be regions which have circumstances which increase the amount of data required for a given period during the day. The most important factor is the length of the longest unit of work, as this defines the amount of data which must be available for backout.

On the prior page is a series of runs changing one or two parms which can have a significant effect on logstream operation. All runs were made using 2 transactions which cause 5000+ log writes each.

The first DASDONLY run (#6) was made with AKPFREQ set to 4000 (the default) LGDFINT set to 30 (the default prior to CICS Transaction Server R2.2) a staging dataset size of 2518 (STG_SIZE on the logstream definition), HIGHOFFLOAD set to 95% and LOWOFFLOAD set to 19%. HIGHOFFLOAD, LOWOFFLOAD, and STG_SIZE were set based on the output from DFHLSCU (prior to PQ34671).

The SMF88 data shows 23478272 bytes were offloaded to the offload dataset, via 5 offloads, with 1 DASD shift. Also notice there were 5694 logstream deletes with DASD (offload dataset) write. There were 280 writes which either reached or were over the HIGHOFFLOAD (STG THLD) value. A staging dataset full condition was reached once. A point of interest is the logger writes data to the offload dataset in 4K (4096 bytes) CIs. Dividing 23478272 by 4096 tells us it cost 5732 I/O to offload the data. Note after application of OW31383 for DFSMS, offload dataset CISIZE can be specified up to 24K.

From the report, field BYT WRITTN to DASD (data written to the offload dataset), seems to be incorrect. This field reflects the actual user data which is written to the offload dataset. This is in contrast to the data written to the staging dataset, which is rounded to a CI (4K) boundary.

In run #7, AKPFREQ was dropped to 1000. The effect of the change was a reduction in offloaded data, and the number in STG THLD now maps to the number of offloads. The number of DASD shifts is still a concern but is directly related to the amount of data being offloaded and the value specified in LS_SIZE.

The second entry for run #7 is for the offload when CICS disconnects from the logstream. The interesting point is the average buffer size and the fact there is 1 offload. Remember, all data is offloaded from a logstream when the last connector disconnects, the staging dataset is then freed and de-allocated.

In run number 8, just to show a very poor example, I dropped the STG_SIZE to 1500, leaving everything else the same. The number of bytes offloaded increased dramatically, as expected. The number of offloads, DASD shifts, STG THLD count, and the number of STG FULL conditions all increased.

For run #9, STG_SIZE was set back to 2518 and HIGHOFFLOAD set to 80% with LOWOFFLOAD set to 50%. The results were very positive, the number of bytes offloaded is zero, and the number of offloads is directly tied to the number of times STG THLD was hit.

In run #10, STG_SIZE was set to 3500. In this case everything looks good except being over allocated may result in a large amount of extra data to be kept, causing the size of the logger dataspace to be larger than needed. This can lead to increased paging in a system with limited central storage and increased CICS startup time due to the logger formatting the staging dataset.

Coupling Facility Activity Report

```
COUPLING FACILITY NAME = SSCF04
TOTAL SAMPLES(AVG) =    899   (MAX) =    899   (MIN) =    899
```

COUPLING FACILITY USAGE SUMMARY

STRUCTURE SUMMARY

TYPE	STRUCTURE NAME	STATUS	CHG	ALLOC SIZE	% OF CF STORAGE	# REQ	% OF ALL REQ	AVG REQ/ SEC	LST/DIR ENTRIES TOT/CUR	DATA ELEMENTS TOT/CUR	LOCK ENTRIES TOT/CUR	DIR REC/ DIR REC XI'S
LIST	DSN510PB_SCA	ACTIVE		10M	1.0%	3586	3.5%	1.99	16K	32K	N/A	N/A
									152	212	N/A	N/A
	LOG_JG	ACTIVE		5M	0.5%	26011	25.6%	14.45	4317	13K	N/A	N/A
									1960	4806	N/A	N/A
	LOG_RRS_TEST	ACTIVE		5M	0.5%	3835	3.8%	2.13	4547	14K	N/A	N/A
									875	1044	N/A	N/A

PROCESSOR SUMMARY

```
COUPLING FACILITY          9674        MODEL C05      CFLEVEL   7
AVERAGE CF UTILIZATION  (% BUSY)            1.1      LOGICAL PROCESSORS:   DEFINED   2    EFFECTIVE   1.5
```

C O U P L I N G F A C I L I T Y A C T I V I T Y

```
    OS/390                    SYSPLEX PLEXB              START 01/08/2000-18.30.00          INTERVAL 030.00.000
    REL. 02.08.00             RPT VERSION 2.7.0          END   01/08/2000-19.00.00          CYCLE 01.000 SECONDS
```

COUPLING FACILITY STRUCTURE ACTIVITY

STRUCTURE NAME = LOG_JG TYPE = LIST

SYSTEM NAME	# REQ TOTAL AVG/SEC	REQUESTS	# REQ	% OF ALL	-SERV TIME(MIC)- AVG	STD_DEV	DELAYED REQUESTS REASON	# REQ	% OF REQ	AVG TIME(MIC) /DEL	STD_DEV	/ALL
MV55	26011	SYNC	24K	91.7%	284.0	1314.7						
	14.45	ASYNC	2157	8.3%	2028.7	2429.1	NO SCH	8	0.4%	1471	1582	5.4
		CHNGD	9	0.0%	INCLUDED IN ASYNC							
							DUMP	0	0.0%	0.0	0.0	

CF Activity Report ..

For Coupling Facility logstreams, the coupling facility activity report can provide important information about the structures. Determine the number of logstreams associated with the structure in question, from either the SMF 88 data, or the logstream definitions. Using the information for structure LOG_JG , the structure size is 5M, which is .5% of the total CF storage. This structure did 25.6% of the requests to the CF in the 30 minute interval shown. The average request rate was 14.45 per second.

Logstreams are placed in LIST type structures. Under the LST/DIR column there are 2 lines per structure. The first line gives the size (number) of entries and the second line is the number in use. The entries are in a common pool for the structure (the entries for all logstreams come from the same pool). Dividing the number of data elements (13K) by the number of Entries (4317) gives an entry to element ratio of 1:4 . The entry to element ratio is dictated by the worst behaving logstream. If the number of entries in use reaches 90% of the total number of entries for the structure, the logger will force an offload of ALL logstreams in the structure.

Notice the current usage (second row of data) indicates a 1:3 ratio (4806/1960), rounded to the next whole number. This indicates that although we have one or more logstreams in the structure, which have an average buffer size of around 768 bytes (3*256), most of the records currently in the structure are much smaller.

In the Data Elements column, the first line gives the number of data elements in the structure, the second line gives the number in use. An important point to remember is the number of data elements is equally divided among the connected logstreams. So, if the number of data elements is 13K and there are 3 logstreams connected to the structure, each logstream has 4.3K data elements.

The Coupling Facility Activity report provides information on the request activity for the structure. From the MV55 system, there were 26011 requests with an average of 14.45 per second. 24K of the requests were synchronous with average service time of 284 microseconds. Anything less than 300 microseconds is acceptable but a G5 processor is capable of service times in the 60 - 100 microsecond range.

The average, or mean represents the middle of the distribution of a set of individual measurements. The standard deviation measures the spread or variation of the individual measurements. 66% of all observations lie within plus or minus 1 standard deviation. 95% of all observations lie within plus or minus 2 standard deviations. Some of the tests showed an average SYNC time of 284 microseconds and a standard deviation of 1315 microseconds for the population of SYNC requests. This indicates 95% of all SYNC requests for this test would lie between 0 and 2914 microseconds. This also indicates some portion of the CF configuration is non-responsive, and causing large variability in individual measurements.

CF Activity Report ...

The most frequently seen reasons for a non-responsive CF are either the use of shared CF CPs, or use of DYNDISP=YES for the CF LPAR. Looking at the RMF CF Usage Summary report, whenever the number of logical processors defined is greater than the effective logical processor, the configuration may be seeing performance issues due to dynamic dispatching or shared CP. The example shows the SSCF04 LPAR has 2 logical processors defined but the "effective" is only 1.5.

For production CICS regions DYNDISP=NO is recommended. DYNDISP is specified on the CF LPAR configuration frame .

Another point of caution - if the CF is actually an LPAR in the same machine as the MVS image, and the LPARs share CPs, ALL SYNC requests will be converted to ASYNC requests. Neither CICS nor the MVS logger have control (or knowledge of the change) and the reports will show the requests as SYNC but the service times will be elongated.

Under the delayed requests we see 8 of 2157 (.4%) ASYNC requests were delayed due to no subchannel available. Below 10% is okay.

Coupling Facility Activity Report - 2

```
COUPLING FACILITY NAME = SSCF04
TOTAL SAMPLES(AVG) =   898   (MAX) =   898   (MIN) =   897
```

COUPLING FACILITY USAGE SUMMARY

STRUCTURE SUMMARY

TYPE	NAME	STATUS	CHG	ALLOC SIZE	% OF CF STORAGE	# REQ	% OF ALL REQ	AVG REQ/ SEC	LST/DIR ENTRIES TOT/CUR	DATA ELEMENTS TOT/CUR	LOCK ENTRIES TOT/CUR	DIR REC/ DIR REC XI'S
LIST	ISTMNPS	ACTIVE		12M	1.2%	0	0.0%	0.00	15K	30K	N/A	N/A
									1	0	N/A	N/A
	IXCDEF	ACTIVE		8M	0.8%	23562	30.0%	13.09	1862	1845	N/A	N/A
									1	20	N/A	N/A
	LOG_JG	ACTIVE		20M	2.0%	23826	30.3%	13.24	15K	61K	N/A	N/A
									5217	18K	N/A	N/A

PROCESSOR SUMMARY

```
COUPLING FACILITY          9674        MODEL C05      CFLEVEL   8
AVERAGE CF UTILIZATION  (% BUSY)            1.2      LOGICAL PROCESSORS:   DEFINED   1    EFFECTIVE   1.0
```

C O U P L I N G F A C I L I T Y A C T I V I T Y

```
   OS/390                    SYSPLEX PLEXB              START 02/02/2000-19.00.00          INTERVAL 030.00.000
   REL. 02.08.00             RPT VERSION 2.7.0          END   02/02/2000-19.30.00          CYCLE 01.000 SECONDS
```

```
COUPLING FACILITY NAME = SSCF04
```

COUPLING FACILITY STRUCTURE ACTIVITY

STRUCTURE NAME = LOG_JG TYPE = LIST

SYSTEM NAME	# REQ TOTAL AVG/SEC		REQUESTS # REQ	% OF ALL	-SERV TIME(MIC)- AVG	STD_DEV	DELAYED REQUESTS REASON	# REQ	% OF REQ	AVG TIME(MIC) /DEL	STD_DEV	/ALL
MV55	23826	SYNC	23K	95.3%	125.1	268.6						
	13.24	ASYNC	1113	4.7%	2237.6	2098.7	NO SCH	0	0.0%	0.0	0.0	0.0
		CHNGD	0	0.0%	INCLUDED IN ASYNC							
							DUMP	0	0.0%	0.0	0.0	

CF Activity Report -2 ..

In this sample the structure size for LOG_JG is 20M, which is 2% of the total CF storage. This structure did 30.3% of the requests to the CF in the 30 minute interval shown. The average request rate was 13.24 per second. During this test, there were 3 CICS regions (IYOT1, IYOT3, and IYOT4) with DFHLOG and DFHSHUNT (6 logstreams) connected to structure LOG_JG. IYOT4 was executing on MV56.

From the MV55 system, there were 23826 requests with an average of 13.24 per second. 23K of the requests were synchronous with an average service time of 125.1 microseconds. You will note this is a significant improvement from the 284 microseconds on the prior report.

In the prior runs (reference the first CF Activity Report) the 9674 had the following configuration:
The 9674 is a model C05 (6 CPs available) total storage 8192k. There are 5 CF LPARs defined :-
CF01 2048K 4 non dedicated non capped CPs weighted 100
CF02 1024K 2 non dedicated non capped CPs weighted 100
CF03 1024K 2 non dedicated non capped CPs weighted 100
CF04 1024K 2 non dedicated non capped CPs weighted 100 <==========
CF05 1024K 2 non dedicated non capped CPs weighted 100
DYNDISP is set to NO in all LPARs.

The change which caused the improvement was giving a dedicated CP to CF04.

The 9674 is a model C05 (6 CPs available) total storage 8192k. There are 5 CF LPARs defined :-
CF01 2048K 4 non dedicated non capped CPs weighted 100
CF02 1024K 2 non dedicated non capped CPs weighted 100
CF03 1024K 2 non dedicated non capped CPs weighted 100
CF04 1024K 1 dedicated non capped CP weighted 100 <==========
CF05 1024K 2 non dedicated non capped CPs weighted 100
DYNDISP is set to NO in all LPARs.

The STD_DEV being 268 indicates much less variability in the samples, i.e. a more consistent service time compared to 1314.7 in the prior report. With the reduction in service times comes a reduction in the CICS task CPU times.

Notice in the Processor summary, there is 1 logical processor defined and the "effective" is now 1.

Additional Tools

- **Workload activity reports**
  - produced from SMF 70 to 79 records using the RMF post processor
    - **ERBRMFPP**
  - provide a performance view of the region(s)
- **CICS Performance Analyzer**
  - used to process the SMF 110 records written by CICS
  - in release 2 process the SMF 88 records

Workload Activity Reports

W O R K L O A D A C T I V I T Y

PAGE 1

z/OS V1R2 SYSPLEX WSCZPLEX START 01/01/2002-23.15.00 INTERVAL 000.04.59 MODE = GOAL
RPT VERSION V1R2 RMF END 01/01/2002-23.19.59

POLICY ACTIVATION DATE/TIME 11/21/2001 12.21.14

REPORT CLASS(ES)

REPORT BY: POLICY=WLMPOL REPORT CLASS=RIYOT1
DESCRIPTION =CICS Report Class for Jim G.

TRANSACTIONS		TRANS.-TIME HHH.MM.SS.TTT		--DASD I/O--		--SERVICE--		SERVICE RATES-		PAGE-IN RATES		----STORAGE----	
AVG	1.00	ACTUAL	0	SSCHRT	11.4	IOC	343	ABSRPTN	631	SINGLE	0.0	AVG	8517.43
MPL	1.00	EXECUTION	0	RESP	1.9	CPU	9348	TRX SERV	631	BLOCK	0.0	TOTAL	8517.43
ENDED	0	QUEUED	0	CONN	1.6	MSO	175902	TCB	0.8	SHARED	0.0	CENTRAL	8517.43
END/S	0.00	R/S AFFINITY	0	DISC	0.1	SRB	3566	SRB	0.3	HSP	0.0	EXPAND	0.00
#SWAPS	0	INELIGIBLE	0	Q+PEND	0.2	TOT	189159	RCT	0.0	HSP MISS	0.0		
EXCTD	0	CONVERSION	0	IOSQ	0.0	/SEC	631	IIT	0.0	EXP SNGL	0.0	SHARED	1.00
AVG ENC	0.00	STD DEV	0					HST	0.0	EXP BLK	0.0		
REM ENC	0.00							APPL %	0.4	EXP SHR	0.0		
MS ENC	0.00												

REPORT BY: POLICY=WLMPOL REPORT CLASS=RLOGER
DESCRIPTION =Report for System Logger

TRANSACTIONS		TRANS.-TIME HHH.MM.SS.TTT		--DASD I/O--		--SERVICE--		SERVICE RATES-		PAGE-IN RATES		----STORAGE----	
AVG	1.00	ACTUAL	0	SSCHRT	12.3	IOC	1	ABSRPTN	8	SINGLE	0.0	AVG	5115.59
MPL	1.00	EXECUTION	0	RESP	1.1	CPU	85	TRX SERV	8	BLOCK	0.0	TOTAL	5115.59
ENDED	0	QUEUED	0	CONN	0.8	MSO	1181	TCB	0.0	SHARED	0.0	CENTRAL	5115.59
END/S	0.00	R/S AFFINITY	0	DISC	0.1	SRB	1085	SRB	0.1	HSP	0.0	EXPAND	0.00
#SWAPS	0	INELIGIBLE	0	Q+PEND	0.2	TOT	2352	RCT	0.0	HSP MISS	0.0		
EXCTD	0	CONVERSION	0	IOSQ	0.0	/SEC	8	IIT	0.0	EXP SNGL	0.0	SHARED	0.00
AVG ENC	0.00	STD DEV	0					HST	0.0	EXP BLK	0.0		
REM ENC	0.00							APPL %	0.0	EXP SHR	0.0		
MS ENC	0.00												

Workload Activity Reports notes

RMF (Resource Measurement Facility) provides a wealth of information which is invaluable in the resolution of performance problems. This information can be used to measure the impact of logstream usage and definition changes.

The prior page contains a WLM Workload Activity Report which presents data collected for report classes RIYOT1 (a CICS region) and RLOGER (the MVS logger address space). Report classes are defined using the WLM ISPF panels (=WLM).

The report interval is listed in the start and end times at the top of the page. A word of caution, the minimum interval is defined by the INTVAL() parm in the SMFPRMxx member of SYS1.PARMLIB. In the samples collected, the interval was set to 5 minutes : INTVAL(05) /* SMF GLOBAL RECORDING INTERVAL */
It's also important to ensure the SMF 70 to 79 records are being collected, along with the CICS 110 records. Which records are to be collected is also defined in the SMFPRMxx member.

```
SUBSYS(STC,EXITS(IEFACTRT),INTERVAL(SMF,SYNC),
                        TYPE(0,30,70:79,88,89,90,99,110,245))
SUBSYS(OMVS,NOEXITS,INTERVAL(SMF,SYNC),
                        TYPE(0,30,70:79,90,88,89,99,110,245))
```

When the reports are formatted, it's possible to report a larger interval than was specified in the SMFPRMxx member, by using the DINTV parm for ERBRMPFF. However, don't forget the length of the minimum interval is the value specified for INTVAL. One word of caution -- SMF88 data which is formatted using IXGRPT1 does not have the ability to summarize at a larger interval than the interval used for data collection (the INTVAL value specified in the current SMFPRMxx).

The following fields should be noted in the reports:

- MPL -- Multiprogramming level, number of address spaces active in this service/ report class during the interval
- TCB -- provides the CPU seconds accumulated in TCB mode during the collection interval
- SRB -- provides the CPU seconds accumulated in SRB mode during the collection interval
- APPL% -- percentage of a engine (CP) used during the collection interval
- Under STORAGE:
  - AVG -- average number of central and expanded storage frames allocated to ASIDs in the report class
- Under PAGE-IN RATES:
  - SINGLE The average rate at which pages are read into central storage
- Under DASD I/O:
  - SSCHRT -- Number of start subchannels (SSCH) per second in the reported interval
  - RESP -- Average DASD response time (in milliseconds)

Reference z/OS V1R2.0 RMF Report Analysis SC33-7991-01
or OS/390 Resource Measurement Facility Report Analysis SC28-1950-04

CICS Performance Analyzer

- **Post processor for CICS SMF 110 records**
  - SMF 88 (MVS Logger) and SMF 101 (DB2) added in R2.
- **For comparison of data the following summary was run for each sample:**

```
V1R1M0                                          CICS Performance Analyzer
                                                  Performance Summary

SUMM0001 Printed at 20:13:58  2/06/2002     Data from 19:45:00  2/06/2002 to 19:50:00  2/06/2002
Interval summary for Logger measurements
```

APPLID	#Tasks	Avg Response Time	Max Response Time	Avg User CPU Time	Avg Dispatch Time	Avg Suspend Time	Avg DispWait Time	Avg JC Wait Time	Avg FCAMRq	Avg FC Wait Time	Avg ENQDelay Time
IYOT1	4520	.6635	1.2136	.0089	.0125	.6510	.0163	.0002	84	.6503	.0000
IYOT2	4406	.6803	1.2372	.0087	.0124	.6679	.0170	.0194	84	.6481	.0000
IYOT3	4409	.6806	1.2507	.0087	.0123	.6683	.0175	.0196	84	.6482	.0000

- **To examine an individual task:**

```
V1R1M0                                          CICS Performance Analyzer
                                                  Performance List

LIST0001 Printed at 21:08:22  1/24/2002    Data from 16:00:04  1/24/2002                APPLID IYOT1
```

Tran	Term	TaskNo	Stop Time	Response Time	Dispatch Time	User CPU Time	Suspend Time	DispWait Time	FCGET	FCPUT	FC Total	FC Wait Time	FCAMRq
JOR2		728	16:00:04.168	7.9594	1.0615	.0370	6.8978	.0210	42	42	84	.2352	944

CICS PA notes

In addition to using the RMF data, a means of summarizing the CICS SMF 110 records was needed in order to understand the impact of each change on task CPU and response time. There are a number of tools which might have been used, but I wanted to see what CICS Performance Analyzer could do. The results speak for themselves in that I was able to define the fields of interest (standard CICS CMF fields) and summarize the sample period to match the RMF data collected.

The first sample on the prior page shows the summary which will be used to monitor the impact on the CICS transaction in the samples to follow in the presentation. The transactions are summarized by APPLID to show the number of transactions which were executed during the interval. The average response time (AVG Response Time) and average CPU time (AVG User CPU Time) indicate throughput and CPU cost per transaction.

Dividing the average CPU time by the average dispatch time (AVG Dispatch Time) gives the CPU to dispatch ratio. I like to see this ratio above 80%. In this case it was 71%, this indicates the processor was being 'stolen' from CICS to service higher priority work in the system. This can be caused by a number of external factors such as a high number of page faults, an excessive number of I/O interrupts being processed by the CP, a higher weighted LPAR needing an engine, etc.

The average dispatch wait time (AVG DispWait Time) is an indicator of how well CICS is able to dispatch tasks which are waiting on the dispatchable queue. This is the interval between the time when an ECB is posted and CICS redispatches the task. Large values here are often an indication of a CPU shortage. The average JC (Journal Control) wait (AVG JC Wait) was included as another indicator of how well the journal/logging function was performing.

In the sample listed, IYOT1 executed 4520 tasks in the 5 minute interval from 19:45 to 19:50 on February 6, the average response time was .6635 with an average CPU cost of .0089 seconds. Notice most of the time was spent suspended waiting for file control activity (AVG Suspend Time is .6510 seconds and AVG FC Wait time is .6503 seconds).

At one point during the data collection it was necessary to understand at a task level where the time was being spent. The second example on the prior page shows the information used to define why there was a large variation in the results of running the sample tasks. The problem was due to exclusive conflicts on the VSAM file. The significant key fields are the number of FCGET, FCPUT and FCAMRQ. FCAMRQ is the number of actual VSAM calls made on behalf of this task. It was obvious as I watched the tasks on a CEMT screen that some tasks were basically 'stranded' for exclusive control conflict resolution. In many cases there were several tasks executed prior to the task awaiting exclusive control conflict resolution being re-dispatched -- or so it seemed. In reality, the task was getting re-dispatched as a result of the every_so_often dispatcher scan when the ECB was found to be posted (exclusive control conflict uses an internal hand posted ECB). When the task is re-dispatched, it retried the request only to find another task has obtained the CI first, thus the reason for the number reported in FCAMRQ being larger than the sum of the gets (FCGET) and puts (FCPUT).

This problem is being addressed in CICS Transaction Server R2.2, by VSAM Control Interval Deadlock Avoidance change.

The net result was, I changed the workload to eliminate the exclusive control conflict so any delays in the transaction could be attributed to the operation of the logstream.

CICS PA R2 Logger Report

SYSTEM LOGGER ACTIVITY REPORT (IXGRPT1)

-LOGSTREAM NAME----STRUCTURE NAME--	BYT WRITTN BY USERS IXGWRITES BYT DELETD INTERIM ST W/O DASD	BYT WRITTN TO INTERIM STORAGE # DELETES W/O DASD WRITE	BYT WRITTN TO DASD INTERIM ST W/DASD	#WRITES INVOKED # DELETS W/ WRITE	---# WRITES COMPLETED------ TYPE1 TYPE2 TYPE3 -----EVENT----- OFF- LOAD	DASD SHFT	STRC FULL	NTRY FULL	STG THLD	STG FULL	AVERAGE BUFFER SIZE RE- BLD
02/05/02 4:20:00 AM (SMF INTERVAL TIMESTAMP 'B7246703ADC00000'X)											
IYOT1.DFHLOG LOG_JG_5M	23899926	41643776	0	86606	86289	317	0	275			
	22628806	82001	0	0	8	0	0	0	0	0	0

V1R2M0 — CICS Performance Analyzer — System Logger - List

LOGR0001 Printed at 15:51:45 2/07/2002 — Data from 23:20:00:00 2/04/2002 to 23:24:00:00 2/04/2002 — Page 1

Logstream name	Structure name	MVSID	Flag	Interval expired at	Level
IYOT1.DFHLOG	LOG_JG_5M	SYSD		04:20:00.00 2/05/2002	SP7.0.2

IXGWRITES

Count	Total Bytes	Average Bytes	Bytes Writn to Interim Storage
86606	23900K	276	41644K

DELETIONS

Count With DASD Write	Count Without DASD Write	Bytes After Offload w. DASD	Bytes Int Stor w/o DASD Write
0	82001	0	22629K

EVENTS

Offloads	Staging Threshld	Demand DASD Shifts	Staging Full	Entry Full	Struct Full	Demand Init'd Offloads	Minimum Block Length	Maximum Block Length
8	0	0	0	0	0	0	116	4573

EVENTS

Type1	Type2	Type3	Struct Rebuilds Init'd	Struct Rebuilds Complt'd
86289	317	0	0	0

DASD Writes

Count	Total Bytes	Average	Waits
0	0	0	0

CICS PA Logger report notes

In release 2 of the CICS Performance Analyzer a new function has been added to report the data contained in the SMF 88 records for CICS logstreams. The prior page provides a comparison example of the SMF88 data as reported using IXGRPT1 (supplied in SYS1.SAMPLIB) and the new report using CICS Performance Analyzer.

There are a couple of fields (Minimum Block Length, Maximum Block Length, and Bytes After Offload w.DASD) which are not listed when using the IXGRPT1 sample.

The samples shown are 1 SMF interval 4:19-4:20 A.M. on February 5, 2002.

On the following page is an example of the logstream summary function which can be used to summarize SMF88 data for an interval greater than the SMF collection interval. IXGRPT1 does not provide this function.

CICS PA R2 Logger Report...

V1R2M0 **CICS Performance Analyzer**

System Logger - Logstream Summary

LOGR0002 Printed at 15:51:45 2/07/2002 Data from 23:20:00:00 2/04/2002 to 23:24:00:00 2/04/2002 Page 4

Logstream name	MVSID	Structure name	First interval start	Last interval stop	Total Interval
IYOT1.IYO1.DFHJ03	SYSD	*DASDONLY*	04:19:00.00 2/05/2002	04:24:00.00 2/05/2002	0000:05:00

	IXGWRITES				DELETIONS			
	Count	Total Bytes	Average Bytes	Bytes Writn to Interim Storage	Count With DASD Write	Count Without DASD Write	Bytes After Offload w. DASD	Bytes Int Stor w/o DASD Write
Total	48005	63980K	1333	197956K	48506	0	200118K	0
Rate(/Sec)	160	213266		659852	161	0	667061	0
Minimum	9261	12300K		38134K	9081	0	37515K	0
Maximum	9946	13255K		41001K	10624	0	43778K	0

	EVENTS							
	Offloads	Staging Threshld	Demand DASD Shifts	Block Length	Staging Full	Entry Full	Struct Full	Demand Init'd Offloads
Total	32	8492	273		32	0	0	0
Rate(/Sec)	0	28	0		0	0	0	0
Minimum	6	0	51	210	6	0	0	0
Maximum	7	0	60	9390	7	0	0	0

	EVENTS					DASD Writes			
	Type1	Type2	Type3	Struct Rebuilds Init'd	Struct Rebuilds Complt'd	Count	Total Bytes	Average	Waits
Total	0	0	0	0	0	305	75794K	0	0
Rate(/Sec)	0	0	0	0	0	1	252647		0
Minimum	0	0	0	0	0	57	14205K		0
Maximum	0	0	0	0	0	66	16566K		0

Environment for the second set of tests

- **Hardware configuration**
  - 2064 (Freeway) processor model 116
    - **16 processors**
    - **LPAR (SYSD) has 4 shared CPs**
  - Coupling Facility
    - **2064 model R100 with 2 dedicated CPs**
  - 2105 (SHARK) 3390-3 DASD
    - **PAV (Parallel Access Volume) is dynamic**
- **Software**
  - z/OS V1.4
  - CICS Transaction Server R2.2
    - **up to three identical (as close as possible) regions**
    - **each CICS region and the MVS logger was placed in a separate report class**

Environment Notes ...

The following tests were run on a 2064 model 116 with 2105(SHARK)configured as 3390-3 DASD devices and a standalone CF (2064 R100). The LPAR has 2 CPs.

Data collected includes the SMF 88 records, SMF 70 to 79 records, and the SMF 110(CICS)records. The data was combined in the following charts by sample. Prior to running each test, the logger address space was canceled (FORCE IXGLOGR,ARM) and restarted (S IXGLOGR). The SMF datasets were switched prior to each sample run. The system was left to idle for 10 minutes in order to establish the 'idle' level of the storage used by the logger address space. The tests are run for 30 minutes; a 5 minute interval is then selected for the display from 15 minutes into the run.

Evaluation of the SMF 88 data allows us to fine tune the logstream activity. But to carry the tuning to the next level it is necessary to look at the system resources required to sustain that performance. The workload activity reports (taken from the SMF 70 - 79 records), provide the CPU used, average DASD I/O and response time, TCB and SRB times, and the number of storage frames consumed.

To ease analysis, a separate WLM reporting class has been defined for each CICS region and the MVS Logger. The data is formatted using the RMF Post Processor (ERBRMFPP). The following sample JCL formats the data associated with report classes (RCLASS) RLOGER, RIYOT1, RIYOT2, RIYOT3, and RIYOT4. RLOGER is the MVS Logger address space, while RIYOT1 - 4 are the CICS regions used in the following tests.

```
//S1       EXEC PGM=ERBRMFPP,REGION=0M
//MFPMSGDS DD SYSOUT=*
//SYSPRINT DD SYSOUT=*
//MFPINPUT DD DISP=SHR,
//            DSN=&SMFIN1
//SYSIN DD *
 SYSOUT(O)
 SYSRPTS(WLMGL(RCLASS(RLOGER)))
 SYSRPTS(WLMGL(RCLASS(RIYOT1)))
 SYSRPTS(WLMGL(RCLASS(RIYOT2)))
 SYSRPTS(WLMGL(RCLASS(RIYOT3)))
 SYSRPTS(WLMGL(RCLASS(RIYOT4)))
 DINTV(0005)
 SUMMARY(INT)
```

Combining the information found in the workload activity reports with the information from the SMF 88 records helps us to understand the costs associated with a given logstream configuration. CICS Performance Analyzer was used to summarize transaction data from the CICS 110 records.

Introduction notes ..

The workload consists of up to 3 CICS regions (CICS Transaction Server R130) executing under z/OS R1.2.

For the test runs, transaction (JORM) was started via terminal input. The application reads and writes 7 records from/to a VSAM file (FILEA), issues a syncpoint, then repeats the process 6 times. It then issues an EXEC CICS Start for ten (10) transactions (JORA-J). Each of these transactions reads and writes 7 records from/to a unique VSAM (FILEA-J), issues a syncpoint and repeats the process six times. It will issue an EXEC CICS START for itself. This will repeat up to the number of repetitions passed from JORM (x'7500'). The 10 tasks run in parallel for 30 minutes and the region is canceled. In cases where 3 regions are used, the process is the same in each region.

In some cases the only logstreams for each CICS region are DFHLOG and DFHSHUNT. The logstreams are varied between a CF logstream and DASDONLY. No data is ever written to DFHSHUNT during the tests. In other cases, a forward recovery log is added to each region.

The data in the presentation is laid out with the SMF88 data shown for each run followed by a combination of data from RMF and the CICS 110 records summarized using the CICS Performance analyzer.

Type	CICS TCB	DASD I/O	APPL%	# Tran/ Second	Average CPU/Task	Average Resp	Average JC Wait	Average File WT	Logr TCB	Logr SRB	DASD I/O	APPL%	Storage Frames	Idle Frames	Net Frames
DASD	0.4	8.3	0.2	.197	0.005	50.1582	50.0519	.0988	0.1	0	2.3	0.1	2571.42	1025.06	1546.36

The data on the left is for the CICS address space; the data on the right starting with LOGR TCB is for the MVS logger address space (IXGLOGR)

- Type is the logstream type
  - DASD for DASDONLY
  - CF64 for Coupling Facility w/MAXBUFSIZE(64000)
  - CF32 for Coupling Facility w/MAXBUFSIZE(32000)
- CICS TCB is the CICS TCB time in seconds for the interval taken from the RMF data
- DASD I/O is the average number of I/Os per second initiated during the interval
- Storage Frames is the average number of 4K page frames in use by the address space. For the logger, the number of idle (i.e. no activity) pages is subtracted from the average Storage Frames to determine the increase due to CICS logging activity in the interval.

Initially it was difficult to define the logger steady state until OPERLOG was varied off. With OPERLOG running, the amount of storage required for logger would grow as batch jobs ran. To vary OPERLOG off, so only syslog is used, issue the z/OS command: `v operlog,hardcpy,off` .

A  symbol has been placed on several samples to indicate the best response and throughput for that configuration. Those samples can be used to perform a configuration comparison for cost and throughput.

DASDONLY LOGSTREAMS

```
                                  BYT WRITTN   BYT WRITTN   BYT WRITTN                                          AVERAGE
                                   BY USERS    TO INTERIM    TO DASD     #WRITES  ---# WRITES COMPLETED------   BUFFER
-LOGSTREAM NAME- STRUCTURE NAME-- IXGWRITES    STORAGE                   INVOKED  TYPE1     TYPE2     TYPE3     SIZE
                                 BYT DELETD   # DELETES   BYT DELETD  # DELETS  ---------------EVENT--------------
                                 INTERIM ST   W/O DASD    INTERIM ST    W/      OFF- DASD STRC NTRY STG  STG    RE-
                                 W/O DASD     WRITE       W/DASD       WRITE   LOAD SHFT FULL FULL THLD FULL   BLD
-------------------------------------------------------------------------------------------------------------------------
(14) AKPFREQ 1000 LGDFINT 30 STG_SIZE 2518 HIGHOFFLOAD 80 LOWOFFLOAD 50
12/11/02  2:20:00 AM  (SMF INTERVAL TIMESTAMP 'B8A8CC80A7400000'X)
 IYOT1.CICS22.DFHLOG *DASDONLY*   3032896     9834496          0      2401        0         0         0      1263
                                  7585792        1852          0         0        1    0    0    0    1    0    0

Type CICS  DASD APPL%  # Tran/ Average   Average   Average Average  Logr Logr  DASD  APPL% Storage  Idle      Net
     TCB   I/O         Second  CPU/Task  Resp      JC Wait File WT  TCB  SRB   I/O         Frames   Frames    Frames
DASD 1.6   37.8  0.6    1.57     .0044    11.11     11.042  0.0608    0.1  0.1   8.4   0.1   2949.35  859.09    2090.26
                                                                                    Logstream Writes/sec: 10.42
*************************************************************************************************************************
(15) AKPFREQ 1000 LGDFINT 5 STG_SIZE 2518 HIGHOFFLOAD 80 LOWOFFLOAD 50
12/11/02  3:05:00 AM  (SMF INTERVAL TIMESTAMP 'B8A8D68F92F00000'X)
 IYOT1.CICS22.DFHLOG  *DASDONLY*  11240854    32935936          0      8041        0         0         0      1397
                                  30801920        7520    1368064       334        4    0    0    0    4    0    0

Type CICS  DASD APPL%  # Tran/ Average   Average   Average Average  Logr Logr  DASD  APPL% Storage  Idle      Net
     TCB   I/O         Second  CPU/Task  Resp      JC Wait File WT  TCB  SRB   I/O         Frames   Frames    Frames
DASD 5.4   140.4 2.2      4     .0039     3.0288    2.9631  0.059    0.1  0.2   27.3  0.1   2850.76   859.08   1991.68
                                                                                    Logstream Writes/sec:  29.85
*************************************************************************************************************************
(16) AKPFREQ 1000 LGDFINT 00 STG_SIZE 2518 HIGHOFFLOAD 80 LOWOFFLOAD 50
12/11/02  3:50:00 AM  (SMF INTERVAL TIMESTAMP 'B8A8E09E7EA00000'X) )
 IYOT1.CICS22.DFHLOG *DASDONLY*   184275302   699846656          0    170861        0         0         0      1078
                                  701902848      171363          0         0       87    0    0    0  156    0    0
Type CICS  DASD APPL%  # Tran/ Average   Average   Average Average  Logr Logr  DASD  APPL% Storage  Idle      Net
     TCB   I/O         Second  CPU/Task  Resp      JC Wait File WT  TCB  SRB   I/O         Frames   Frames    Frames
DASD 85.3  2277  35.2   56.83    .0038     0.1845    0.1224  0.0556   0.2  3.9   572.8 1.8   3092.09  859.08   2233.01
                                                                                    Logstream Writes/sec:  569.75
```

1

DASDONLY LOGSTREAMS

As we continue the evaluation of logger/logstream performance, in addition to the SMF 88 records, RMF data (SMF 70 to 78 records) and performance data from the CICS 110 SMF records will be used. The CICS data is formatted using the CICS Performance Analyzer utility.

The format of the presentation pages has been altered to merge data from all three sources into a single report.

The SMF interval has been set to 5 minutes. The data used is for a 5 minute interval, after the workload has been running for at least 10 minutes.

The transactions are designed to be very quick, simply reading and writing 42 records of a VSAM file. A syncpoint is issued after every 6 read/write operations. In the initial tests, the files (there is a separate file for each transaction) are defined with RECOVERY(BACKOUTONLY). In later tests, a forward recovery log was added.

Sample 14. AKPFREQ 1000 LGDFINT 30 STG_SIZE 2518 HIGHOFFLOAD 80 LOWOFFLOAD 50

Type	CICS TCB	DASD I/O	APPL%	# Tran/ Second	Average CPU/Task	Average Resp	Average JC Wait	Average File WT	Logr TCB	Logr SRB	DASD I/O	APPL%	Storage Frames	Idle Frames	Net Frames
DASD	1.6	37.8	0.6	1.57	.0044	11.11	11.042	0.0608	0.1	0.1	8.4	0.1	2949.35	859.09	2090.26

Logstream Writes/sec: 10.42

In this test, DFHLOG was defined as a DASDONLY logstream with STG_SIZE set to 2518, which equates to roughly 10.3M (2518CIs*4096 CISIZE). AKPFREQ is set to 1000 and LGDFINT set to 30 with HIGHOFFLOAD 80 and LOWOFFLOAD 50. There are no user journals in this test. Notice in the SMF 88 data there was 1 offload in the interval. The problem is not the size of the logstream, but the fact very little work is getting through the CICS region.

Notice the CICS TCB time is very low, which supports the long response times. Also note that CICS is only issuing an average of 37.8 DASD I/O per second. The DASD I/O is a direct reflection of the number of file control requests in the region. Finally notice the CICS region is only using .6% of an engine during this interval.

In the logger address space the CPU, SRB and APPL% values are very low indicating a minimum of work is being passed to the logger. The average number of DASD I/O is 8.4 per second and the average net storage frames is 2090.26. The number of storage frames (a 4K page) is an important value to watch, especially in an LPAR which is storage constrained, i.e. short on CSTOR (Central Storage). In this case the logger is averaging 8.56M (2090.26 * 4096) of storage.

The average response time per transaction is 11.11 seconds with an average CPU per task at .0044 seconds. Note there is an average of 11.042 seconds spent waiting for journal I/O to complete.

DASDONLY LOGSTREAMS notes...

Sample 15. AKPFREQ 1000 LGDFINT 5 STG_SIZE 2518 HIGHOFFLOAD 80 LOWOFFLOAD 50

Type	CICS TCB	DASD I/O	APPL%	# Tran/ Second	Average CPU/Task	Average Resp	Average JC Wait	Average File WT	Logr TCB	Logr SRB	DASD I/O	APPL%	Storage Frames	Idle Frames	Net Frames
DASD	5.4	140.4	2.2	4	.0039	3.0288	2.9631	0.059	0.1	0.2	27.3	0.1	2850.76	859.08	1991.68

Logstream Writes/sec: 29.85

The change in this sample is to drop LGDFINT to 5. This single change had a tremendous positive impact.

Notice the average response has gone from 11.11 to 3.0288 seconds. The TCB time reflects the fact the region is processing more work, along with the fact it is now averaging 140.4 DASD I/O per second, the CPU per transaction is now .0039. Note the JCWait time has dropped to 2.9631 seconds. This clearly indicates transaction throughput was being throttled by log defer interval (LGDFINT) being set to 30.

The logger DASD I/O increased from 8.4 to 27.3 per second. The average storage used by the logger has decreased slightly to 8.16M (1991.68 * 4096).

Sample 16. AKPFREQ 1000 LGDFINT 00 STG_SIZE 2518 HIGHOFFLOAD 80 LOWOFFLOAD 50

Type	CICS TCB	DASD I/O	APPL%	# Tran/ Second	Average CPU/Task	Average Resp	Average JC Wait	Average File WT	Logr TCB	Logr SRB	DASD I/O	APPL%	Storage Frames	Idle Frames	Net Frames
DASD	85.3	2277	35.2	56.83	.0038	0.1845	0.1224	0.0556	0.2	3.9	572.8	1.8	3092.09	859.08	2233.01

In this case I wanted to see what would happen if the log defer interval was removed, so LGDFINT was set to 00. The results are interesting. The average response time dropped to 0.1845, with a huge increase in the DASD I/O for both CICS and the logger. The increase of CICS TCB time from 5.4 to 85.3 shows a large increase in the amount of work being processed in the CICS region. Note there is still an average of 0.1224 seconds per task spent waiting for log writes to complete. The DASD requests in the logger (from the RMF data) are averaging 1.3 milliseconds. Each task is doing 42 read/updates and replaces, accounting for .1092 seconds of wait time. Since this is a DASDONLY logstream there is wait and redispatch time for each request. The average logger storage frames is (9.15M).

Note: these tests were run on a very fast processor (z900 turbo), use of LGDFINT=0 on a slower processor could create enough overhead to offset any gains in throughput.

DASDONLY LOGSTREAMS

```
                                  BYT WRITTN  BYT WRITTN  BYT WRITTN                                        AVERAGE
                                   BY USERS    TO INTERIM    TO DASD    #WRITES  ---# WRITES COMPLETED------  BUFFER
-LOGSTREAM NAME- STRUCTURE NAME-- IXGWRITES    STORAGE                  INVOKED   TYPE1     TYPE2     TYPE3    SIZE

                                 BYT DELETD  # DELETES   BYT DELETD  # DELETS   ---------------EVENT--------------
                                 INTERIM ST   W/O DASD   INTERIM ST     W/      OFF- DASD STRC NTRY STG   STG    RE-
                                 W/O DASD     WRITE      W/DASD       WRITE    LOAD SHFT FULL FULL THLD FULL   BLD
```

(17) 3 Regions AKPFREQ 1000 LGDFINT 5 STG_SIZE 2518 HIGHOFFLOAD 80 LOWOFFLOAD 50

12/11/02 5:00:00 AM (SMF INTERVAL TIMESTAMP 'B8A8F043ED400000'X)

```
IYOT1.CICS22.DFHLOG  *DASDONLY*  72431138   254234624         0      62069        0        0        0      1166
                                 243875840      59540   7192576      1756       32    0   0    0   47    0    0
```

Type	CICS TCB	DASD I/O	APPL%	# Tran/ Second	Average CPU/Task	Average Resp	Average JC Wait	Average File WT	Logr TCB	Logr SRB	DASD I/O	APPL%	Storage Frames	Idle Frames	Net Frames
DASD	35.9	989.1	14.9	21.05	.0040	0.4672	0.3685	0.0914	0.1	7.3	1053	3.3	18783.9	858.08	17925.8

Logstream Writes/sec: 191.62

****

(18) 3 Regions AKPFREQ 1000 LGDFINT 00 STG_SIZE 2518 HIGHOFFLOAD 80 LOWOFFLOAD 50

12/11/02 5:50:00 AM (SMF INTERVAL TIMESTAMP 'B8A8FB70F3200000'X)

```
IYOT1.CICS22.DFHLOG  *DASDONLY*  101452897   388829184         0      94929        0        0        0      1068
                                 380203008      92823   6266880      1530       48    0   0    0   81    0    0
```

Type	CICS TCB	DASD I/O	APPL%	# Tran/ Second	Average CPU/Task	Average Resp	Average JC Wait	Average File WT	Logr TCB	Logr SRB	DASD I/O	APPL%	Storage Frames	Idle Frames	Net Frames
DASD	49.9	1253	20.9	31.33	0.0041	0.3352	0.2367	0.0912	0.2	7.8	1131	3.7	18147.2	861.08	17286.12

Logstream Writes/sec: 317.43

****

DASDONLY LOGSTREAMS- 3 Regions

Sample 17. 3 Regions AKPFREQ 1000 LGDFINT 5 STG_SIZE 2518 HIGHOFFLOAD 80 LOWOFFLOAD 50

12/11/02 5:00:00 AM (SMF INTERVAL TIMESTAMP 'B8A8F043ED400000'X)

```
    IYOT1.CICS22.DFHLOG  *DASDONLY*  72431138   254234624          0     62069        0         0         0      1166
                                     243875840      59540    7192576      1756       32     0    0    0   47     0    0
```

Type	CICS TCB	DASD I/O	APPL%	# Tran/ Second	Average CPU/Task	Average Resp	Average JC Wait	Average File WT	Logr TCB	Logr SRB	DASD I/O	APPL%	Storage Frames	Idle Frames	Net Frames
DASD	35.9	989.1	14.9	21.05	.0040	0.4672	0.3685	0.0914	0.1	7.3	1053	3.3	18783.9	858.08	17925.8

Logstream Writes/sec: 191.62

In this sample STG_SIZE was set to 2518 and 3 identical regions were started, each running the same workload. Comparing with sample 15, we see an increase in the CPU used in the CICS region, but there is also a significant increase in the transaction rate (sample 15 was 4 while sample 17 is 21.05). Also notice the reduction in JC Wait time in sample 17 is 0.3685 compared to 2.9631 in sample 15. This would indicate the logger is responding quicker when there are additional regions using logstreams. Even considering the logger is handling requests from 3 regions the amount of DASD I/O is up significantly. The SMF data reports the DASD I/O time in sample 17 is 1.7 ms compared to 1.2 ms in sample 15.

With the increase in throughput more data is being offloaded, indicating the size of the logstream should be increased.

The amount of storage used by the logger has increased from 1991.68 to 17925.8, indicating the use of three regions is causing more data to be retained in the logger data space.

Sample 18. 3 Regions AKPFREQ 1000 LGDFINT 00 STG_SIZE 2518 HIGHOFFLOAD 80 LOWOFFLOAD 50

12/11/02 5:50:00 AM (SMF INTERVAL TIMESTAMP 'B8A8FB70F3200000'X)

```
 IYOT1.CICS22.DFHLOG *DASDONLY*  101452897  388829184          0     94929        0         0         0      1068
                                 380203008      92823    6266880      1530       48    0    0    0   81     0    0
```

Type	CICS TCB	DASD I/O	APPL%	# Tran/ Second	Average CPU/Task	Average Resp	Average JC Wait	Average File WT	Logr TCB	Logr SRB	DASD I/O	APPL%	Storage Frames	Idle Frames	Net Frames
DASD	49.9	1253	20.9	31.33	0.0041	0.3352	0.2367	0.0912	0.2	7.8	1131	3.7	18147.2	861.08	17286.12

Logstream Writes/sec: 317.43

In this sample I again wanted to see the effect of removing the log defer interval, this time in a 3 region environment. From the CICS side, the amount of workload processed (CICS TCB Time) and the number of DASD I/O requests indicated a significant increase in throughput. The reduction in JC Wait time shows the positive impact of removing the log defer interval. It's interesting to note, on the logger side there is actually a decrease in the storage frames with an increase in the DASD I/O. Logger storage is now at 70.8M.

Once again the fact data is being offloaded is an indication the size of the staging dataset should be increased.

Note: these tests were run on a very fast processor (z900 turbo), use of LGDFINT=0 on a slower processor could create enough overhead to offset any gains in throughput.

Locate the CF logstream

- **Display the Logstream**

```
D LOGGER,LOGSTREAM,LSN=IYOT1.CICS22.DFHLOG
IXG601I   22.42.12  LOGGER DISPLAY 252
INVENTORY INFORMATION BY LOGSTREAM
LOGSTREAM                  STRUCTURE         #CONN  STATUS
---------                  ---------         ------ ------
IYOT1.CICS22.DFHLOG        LOG_JG2_5M        000000 AVAILABLE
```

- **Display the structure**

```
D XCF,STRUCTURE,STRNAME=LOG_JG2_5M
IXC360I  20.48.46  DISPLAY XCF 239
STRNAME: LOG_JG2_5M
 STATUS: NOT ALLOCATED
 POLICY INFORMATION:
  POLICY SIZE     : 30000 K
  POLICY INITSIZE: 5000 K
  POLICY MINSIZE  : 0 K
  FULLTHRESHOLD   : 80
  ALLOWAUTOALT    : NO
  REBUILD PERCENT: 1
  DUPLEX          : ALLOWED
  PREFERENCE LIST: CF2     CF1
  ENFORCEORDER    : NO
  EXCLUSION LIST IS EMPTY
```

Locate the logstream notes..

Using the display logger command (D LOGGER,LOGSTREAM,LSN=IYOT1.CICS22.DFHLOG), it can be seen the logstream is allocated to structure LOG_JG2_5M.

```
D LOGGER,LOGSTREAM,LSN=IYOT1.CICS22.DFHLOG
IXG601I   22.42.12  LOGGER DISPLAY 252
INVENTORY INFORMATION BY LOGSTREAM
LOGSTREAM                  STRUCTURE         #CONN  STATUS
---------                  ---------         ------ ------
IYOT1.CICS22.DFHLOG        LOG_JG2_5M        000000 AVAILABLE
```

A display XCF command is issued to display the structure characteristics and location. The display indicates LOG_JG2_5M is allocated in coupling facility CF2.

```
D XCF,STRUCTURE,STRNAME=LOG_JG2_5M
IXC360I  20.48.46  DISPLAY XCF 239
STRNAME: LOG_JG2_5M
STATUS: NOT ALLOCATED
POLICY INFORMATION:
POLICY SIZE    : 30000 K
POLICY INITSIZE: 5000 K
POLICY MINSIZE : 0 K
FULLTHRESHOLD  : 80
ALLOWAUTOALT   : NO
REBUILD PERCENT: 1
DUPLEX         : ALLOWED
PREFERENCE LIST: CF2    CF1
ENFORCEORDER   : NO
EXCLUSION LIST IS EMPTY
```

Display the Coupling Facility

```
D XCF,CF
IXC361I  18.44.53  DISPLAY XCF 222
  CFNAME      COUPLING FACILITY
  CF1         002064.IBM.02.000000010B2E
              PARTITION: 5   CPCID: 00
  CF103       002064.IBM.02.000000010F91
              PARTITION: 3   CPCID: 00
  CF2         002064.IBM.02.000000010B2E
              PARTITION: 6   CPCID: 00
```

```
D XCF,CF,CFNAME=CF2
IXC362I  18.17.06  DISPLAY XCF 216
CFNAME: CF2
 COUPLING FACILITY      :   002064.IBM.02.000000010B2E
                            PARTITION: 6   CPCID: 00
 POLICY DUMP SPACE SIZE:    10000 K
 ACTUAL DUMP SPACE SIZE:    10240 K
 STORAGE INCREMENT SIZE:    256 K

 NO SYSTEMS ARE CONNECTED TO THIS COUPLING FACILITY
 NO STRUCTURES ARE IN USE BY THIS SYSPLEX IN THIS
    COUPLING FACILITY
```

```
D XCF,CF,CFNAME=CF1
IXC362I  18.18.35  DISPLAY XCF 218
CFNAME: CF1
 COUPLING FACILITY      :   002064.IBM.02.000000010B2E
                            PARTITION: 5   CPCID: 00
 POLICY DUMP SPACE SIZE:    10000 K
 ACTUAL DUMP SPACE SIZE:    10240 K
 STORAGE INCREMENT SIZE:    256 K

 NO SYSTEMS ARE CONNECTED TO THIS COUPLING FACILITY
 NO STRUCTURES ARE IN USE BY THIS SYSPLEX IN THIS
    COUPLING FACILITY
```

```
D XCF,CF,CFNAME=CF103
IXC362I  18.19.19  DISPLAY XCF 220
CFNAME: CF103
  COUPLING FACILITY      : 002064.IBM.02.000000010F91
                           PARTITION: 3   CPCID: 00
  POLICY DUMP SPACE SIZE: 10000 K
  ACTUAL DUMP SPACE SIZE: 10240 K
  STORAGE INCREMENT SIZE: 256 K

  CONNECTED SYSTEMS:
  SYSA      SYSB      SYSD

  STRUCTURES:
  DSNZPLEX_LOCK1  DSNZPLEX_SCA       ISGLOCK
  ISTGSTR01       IXCSTR1            IXCSTR2
  IXCSTR3         IXCSTR4            J2CKPT1
  RRS_ARCHIVE     RRS_RESTART        RRS_RMD
  SYSIGGCAS_ECS   SYSZWLM_WORKUNIT   SYSZWLM_0B2E2064
```

Locate the logstream notes..

The displays on the prior pages provide information to identify the location of each coupling facility.

The D XCF,CF output shows 2 of the coupling facilities (CF1 and CF2) are executing in the 2064 with serial number 10B2E. The third CF is executing on the 2064 with serial number 10F91.

```
D XCF,CF
IXC361I  18.44.53  DISPLAY XCF 222
  CFNAME      COUPLING FACILITY
  CF1         002064.IBM.02.000000010B2E
              PARTITION: 5   CPCID: 00
  CF103       002064.IBM.02.000000010F91
              PARTITION: 3   CPCID: 00
  CF2         002064.IBM.02.000000010B2E
              PARTITION: 6   CPCID: 00
```

Also note the displays do not provide any coupling facility type information for CF1 and CF2. This indicates they are integrated couple facilities (ICF), which are LPARs on the 2064. Since they are in the same physical CEC as the MVS image, they are considered 'failure dependent'. In a 'failure dependent' situation, and STG_DUPLEX(YES) has been specified in the logstream definition, the logger will duplex the data to the staging dataset rather than the dataspace.

Reviewing the CICS joblog showed a staging dataset was allocated to the IYOT1.CICS22.DFLOG logstream. This was evidenced by a DFHLG0777 message during startup. Unfortunately, the staging datasets are released during the disconnect process.

The net result is the response to the log writes can be no quicker than the sum of the CF write time plus the DASD I/O time. This can explain why performance is less than expected when using the CF logstream.

CF LOGSTREAMS -Alter AKPFREQ

```
                                BYT WRITTN   BYT WRITTN   BYT WRITTN                                        AVERAGE
                                 BY USERS    TO INTERIM    TO DASD     #WRITES   ---# WRITES COMPLETED------  BUFFER
-LOGSTREAM NAME- STRUCTURE NAME-- IXGWRITES   STORAGE                  INVOKED   TYPE1     TYPE2     TYPE3    SIZE

                               BYT DELETD  # DELETES   BYT DELETD  # DELETS   --------------EVENT--------------
                               INTERIM ST   W/O DASD   INTERIM ST     W/      OFF- DASD STRC NTRY STG  STG   RE-
                               W/O DASD     WRITE      W/DASD       WRITE    LOAD SHFT FULL FULL THLD FULL  BLD
```

(19A) AKPFREQ 1000 LGDFINT 00 20M structure HIGHOFFLOAD 80 LOWOFFLOAD 50

11/20/02 0:15:00 AM (SMF INTERVAL TIMESTAMP 'B88E49656A900000'X)

```
IYOT1.CICS22.DFHLOG LOG_JG2_20M  203010589   291244032          0    459807   458832       975          0       441
                                 204272897      462665          0         0       38    0    0    0    0    0    0
```

Type	CICS TCB	DASD I/O	APPL%	# Tran/ Second	Average CPU/Task	Average Resp	Average JC Wait	Average File WT	Logr TCB	Logr SRB	DASD I/O	APPL%	Storage Frames	Idle Frames	Net Frames
CF64	114.2	2323	45.6	55.29	0.0057	0.1808	0.0585	0.1046	0.2	11.8	31.6	4.0	11743.3	939.61	10803.69

Logstream Writes/sec: 1423.85

****

(19B) AKPFREQ 2500 LGDFINT 00 20M structure HIGHOFFLOAD 80 LOWOFFLOAD 50

11/20/02 0:30:00 AM (SMF INTERVAL TIMESTAMP 'B88E4CBFB9200000'X)

```
IYOT1.CICS22.DFHLOG LOG_JG2_20M  202501766   290984192          0    460652   459624      1028          0       439
                                 203081243      461994          0         0       40    0    0    0    0    0    0
```

Type	CICS TCB	DASD I/O	APPL%	# Tran/ Second	Average CPU/Task	Average Resp	Average JC Wait	Average File WT	Logr TCB	Logr SRB	DASD I/O	APPL%	Storage Frames	Idle Frames	Net Frames
CF64	113.7	2324	45.4	55.34	0.0057	0.1803	0.0585	0.1047	0.2	11.9	31.6	4.0	12903.3	939.61	11963.69

Logstream Writes/sec: 1470.53

****

(19C) AKPFREQ 4000 LGDFINT 00 20M structure HIGHOFFLOAD 80 LOWOFFLOAD 50

11/20/02 0:40:00 AM (SMF INTERVAL TIMESTAMP 'B88E4EFBED800000'X)

```
IYOT1.CICS22.DFHLOG  LOG_JG2_20M  34224020    49171968     425972     77860    77860         0          0       439
                                  35502061       80766     391092       872        7    0    0  121     0    0    0
```

Type	CICS TCB	DASD I/O	APPL%	# Tran/ Second	Average CPU/Task	Average Resp	Average JC Wait	Average File WT	Logr TCB	Logr SRB	DASD I/O	APPL%	Storage Frames	Idle Frames	Net Frames
CF64	114.0	2334	45.5	55.51	0.0057	0.1800	0.0576	0.1057	0.2	11.9	31.6	4.1	13665.5	939.61	12725.89

Logstream Writes/sec: 1486.29

CF LOGSTREAMS AKPFREQ notes ..

The samples on the prior page (19A, 19B, and 19C) were run using a 20M structure in a standalone, nonvolatile coupling facility, with a single CP LPAR. The samples were taken at three different points in a single CICS execution. Sample 19A had AKPFREQ set to 1000 with LGDFINT of 0.

```
(19A) AKPFREQ 1000 LGDFINT 00 20M structure HIGHOFFLOAD 80 LOWOFFLOAD 50
11/20/02  0:15:00 AM  (SMF INTERVAL TIMESTAMP 'B88E49656A900000'X)
IYOT1.CICS22.DFHLOG LOG_JG2_20M  203010589   291244032          0    459807   458832       975          0       441
                                 204272897      462665          0         0       38    0    0    0    0    0    0
Type CICS  DASD  APPL% # Tran/ Average   Average  Average Average  Logr Logr  DASD  APPL% Storage  Idle      Net
     TCB   I/O         Second  CPU/Task  Resp     JC Wait File WT  TCB  SRB   I/O         Frames   Frames    Frames
CF64 114.2 2323  45.6  55.29   0.0057    0.1808   0.0585  0.1046   0.2  11.8  31.6  4.0   11743.3  939.61  10803.69
                                                                                Logstream Writes/sec: 1423.85


*************************************************************************************************************
(19C) AKPFREQ 4000 LGDFINT 00 20M structure HIGHOFFLOAD 80 LOWOFFLOAD 50
11/20/02  0:40:00 AM  (SMF INTERVAL TIMESTAMP 'B88E4EFBED800000'X)
IYOT1.CICS22.DFHLOG  LOG_JG2_20M  34224020    49171968      425972     77860    77860         0          0       439
                                  35502061       80766      391092       872        7    0    0  121    0    0    0

Type CICS  DASD  APPL% # Tran/ Average   Average  Average Average  Logr Logr  DASD  APPL% Storage  Idle      Net
     TCB   I/O         Second  CPU/Task  Resp     JC Wait File WT  TCB  SRB   I/O         Frames   Frames    Frames
CF64 114.0 2334  45.5  55.51   0.0057    0.1800   0.0576  0.1057   0.2  11.9  31.6  4.1   13665.5  939.61  12725.89
                                                                                Logstream Writes/sec: 1486.29
```

Notice there is very little difference in throughput for the three samples. However data is being offloaded in sample 19C and notice the increase in storage frames (an increase of 1922.2 frames) from sample 19A to 19C.

CF LOGSTREAMS

```
                                BYT WRITTN  BYT WRITTN  BYT WRITTN                                          AVERAGE
                                 BY USERS   TO INTERIM    TO DASD    #WRITES  ---# WRITES COMPLETED------  BUFFER
-LOGSTREAM NAME- STRUCTURE NAME-- IXGWRITES   STORAGE                INVOKED   TYPE1     TYPE2     TYPE3    SIZE

                               BYT DELETD  # DELETES   BYT DELETD  # DELETS   ---------------EVENT---------------
                               INTERIM ST   W/O DASD   INTERIM ST     W/     OFF- DASD STRC NTRY STG  STG    RE-
                               W/O DASD     WRITE      W/DASD       WRITE    LOAD SHFT FULL FULL THLD FULL   BLD
```

```
(20) AKPFREQ 1000 LGDFINT 5 5M structure HIGHOFFLOAD 80 LOWOFFLOAD 50
12/12/02  0:45:00 AM  (SMF INTERVAL TIMESTAMP 'B8A9F9222BB00000'X)
 IYOT1.CICS22.DFHLOG  LOG_JG2_5M  4560403     5881856          0      7417     7371        46          0       614
                                  4418914        7190          0         0        6    0    0    0    0    0    0

Type CICS  DASD APPL%  # Tran/ Average   Average  Average Average  Logr Logr  DASD  APPL% Storage  Idle      Net
      TCB  I/O         Second  CPU/Task  Resp     JC Wait File WT  TCB  SRB   I/O         Frames   Frames    Frames
CF32  2.7  54.2  1.1    1.22    .0055    7.7556   7.6727   0.0739  0.1  0.2   0.3   0.1   2833.95  868.81   1965.14
                                                                                Logstream Writes/sec: 35.25

*************************************************************************************************************
(21) AKPFREQ 1000 LGDFINT 00 5M structure HIGHOFFLOAD 80 LOWOFFLOAD 50
12/12/02  0:55:00 AM  (SMF INTERVAL TIMESTAMP 'B8A9FB5E60100000'X)
 IYOT1.CICS22.DFHLOG  LOG_JG2_5M  320447741   531816704          0   1039948  1025325     14622          0       308
                                  320157518     1039384          0         0      592    0    0    0    0    0    0

Type CICS  DASD  APPL% # Tran/ Average   Average  Average Average  Logr Logr  DASD  APPL% Storage  Idle      Net
      TCB  I/O         Second  CPU/Task  Resp     JC Wait File WT  TCB  SRB   I/O         Frames   Frames    Frames
CF32  201  3457  78.1   82.25    .0071    0.1216   0.0079  0.1027  0.6  14.1  0.3   4.9   3728.27  861.81   2866.46
                                                                                Logstream Writes/sec: 3535.9
*************************************************************************************************************
```

CF LOGSTREAMS notes .

Sample 20. AKPFREQ 1000 LGDFINT 5 5M structure HIGHOFFLOAD 80 LOWOFFLOAD 50

```
12/12/02  0:45:00 AM  (SMF INTERVAL TIMESTAMP 'B8A9F9222BB00000'X)
 IYOT1.CICS22.DFHLOG  LOG_JG2_5M  4560403     5881856          0      7417     7371        46          0       614
                                  4418914        7190          0         0        6    0    0    0    0    0    0
```

Type	CICS TCB	DASD I/O	APPL%	# Tran/ Second	Average CPU/Task	Average Resp	Average JC Wait	Average File WT	Logr TCB	Logr SRB	DASD I/O	APPL%	Storage Frames	Idle Frames	Net Frames
CF32	2.7	54.2	1.1	1.22	.0055	7.7556	7.6727	0.0739	0.1	0.2	0.3	0.1	2833.95	868.81	1965.14

Logstream Writes/sec: 35.25

In this sample, there is a single region with logstream DFHLOG defined to a 5M structure (LOG_JG2_5M). DFHSHUNT was placed in a 20M structure. MAXBUFSIZE is set at 32000. DFHLOG is the only logstream in the LOG_JG2_5M structure. It should be noted there was no activity to DFHSHUNT during the sample. AKPFREQ is set to 1000, LGDFINT 5 with HIGHOFFLOAD 80 and LOWOFFLOAD 50.

The response time is comparable to sample 15, but you will notice the CPU time is higher due to a CF logstream being used.

The logger storage use is 8M (1965.14*4096) which is slightly lower than the 8.16M shown in sample 15.

****

Sample 21. AKPFREQ 1000 LGDFINT 00 5M structure HIGHOFFLOAD 80 LOWOFFLOAD 50

```
12/12/02  0:55:00 AM  (SMF INTERVAL TIMESTAMP 'B8A9FB5E60100000'X)
 IYOT1.CICS22.DFHLOG  LOG_JG2_5M  320447741   531816704          0   1039948  1025325     14622          0       308
                                  320157518     1039384          0         0      592    0    0    0    0    0    0
```

Type	CICS TCB	DASD I/O	APPL%	# Tran/ Second	Average CPU/Task	Average Resp	Average JC Wait	Average File WT	Logr TCB	Logr SRB	DASD I/O	APPL%	Storage Frames	Idle Frames	Net Frames
CF32	201	3457	78.1	82.25	.0071	0.1216	0.0079	0.1027	0.6	14.1	0.3	4.9	3728.27	861.81	2866.46

Logstream Writes/sec: 3535.9

In sample 21, there is again a single region with logstream DFHLOG defined to a 5M structure (LOG_JG2_5M). DFHSHUNT was placed in a 20M structure. MAXBUFSIZE is set at 32000. AKPFREQ is set to 1000, LGDFINT remained at 0 with HIGHOFFLOAD 80 and LOWOFFLOAD 50.

The numbers show a significant change between sample 19 and 20. Setting the log defer interval to zero causes the data to be written sooner. Notice the CICS TCB time is now up to 201 seconds in the 5 minute interval, the DASD I/O is running over 3400 requests per second, which corresponds to the transaction rate of 82.25 per second.

However, the down side is the increase in logger storage frames.

CF LOGSTREAMS - 3 Regions

```
                                BYT WRITTN   BYT WRITTN   BYT WRITTN                                                AVERAGE
                                 BY USERS    TO INTERIM    TO DASD     #WRITES  ---# WRITES COMPLETED------   BUFFER
-LOGSTREAM NAME- STRUCTURE NAME-- IXGWRITES    STORAGE                  INVOKED   TYPE1      TYPE2      TYPE3     SIZE

                               BYT DELETD  # DELETES   BYT DELETD  # DELETS   ---------------EVENT--------------
                               INTERIM ST   W/O DASD   INTERIM ST     W/      OFF- DASD STRC NTRY STG  STG    RE-
                               W/O DASD     WRITE      W/DASD       WRITE    LOAD SHFT FULL FULL THLD FULL   BLD
```

(22) 3 Regions AKPFREQ 1000 LGDFINT 00 20M structure HIGHOFFLOAD 80 LOWOFFLOAD 50

12/12/02 5:40:00 AM (SMF INTERVAL TIMESTAMP 'B8AA3B1234C00000'X)

```
 IYOT1.CICS22.DFHLOG LOG_JG2_20M  174330570 286120448        0    549782   546544      3237          0       317
                                  174750469    551118        0         0      112    0    0    0    0    0    0
```

Type	CICS TCB	DASD I/O	APPL%	# Tran/ Second	Average CPU/Task	Average Resp	Average JC Wait	Average File WT	Logr TCB	Logr SRB	DASD I/O	APPL%	Storage Frames	Idle Frames	Net Frames
CF60	118.5	1891	48.5	45.35	0.0070	0.2221	0.0186	0.1921	0.4	41.9	137.1	14.3	24660.7	859.08	23801.62

Logstream Writes/sec: 1785.19

****

CF Logstream notes

Sample 22. 3 Regions AKPFREQ 1000 LGDFINT 00 20M structure HIGHOFFLOAD 80 LOWOFFLOAD 50

```
12/12/02  5:40:00 AM  (SMF INTERVAL TIMESTAMP 'B8AA3B1234C00000'X)
 IYOT1.CICS22.DFHLOG LOG_JG2_20M  174330570 286120448        0    549782   546544      3237          0       317
                                  174750469    551118        0         0      112    0    0    0    0    0    0

Type CICS  DASD  APPL% # Tran/ Average   Average  Average Average  Logr Logr  DASD  APPL% Storage  Idle     Net
      TCB  I/O         Second  CPU/Task  Resp     JC Wait File WT  TCB  SRB   I/O         Frames   Frames   Frames
CF60 118.5 1891  48.5  45.35   0.0070   0.2221   0.0186  0.1921   0.4  41.9 137.1  14.3 24660.7  859.08 23801.62
                                                                                Logstream Writes/sec: 1785.19
```

In sample 22 there are three regions, all with DFHLOG allocated to a 20M CF structure (each logstream will have approximately 6.6M). The DFHSHUNT logstreams were allocated to other structures. AKPFREQ was set to 1000. No data is being offloaded.

The results proved interesting. The file wait time took a significant increase which was reflected in the response time. Further investigation found there was a conflict at the device level for several of the VSAM files. This contention also reduced the throughput accounting for the CICS TCB time being lower than in the single region samples.

In terms of storage, the logger storage frames have increased to 23801.62 (97.5M). This is a direct result of the number of logstreams and the increase in response time, causing data to be retained in the dataspace for a longer time.

When these samples were run under CICS Transaction Server R1.3, there was a variation in the response times associated with inconsistencies in the JCWait time. This was caused by the problem described in PQ57850. The net result was a bursting effect, resulting in the log defer interval often being raised to the ICV value (partition exit) which was set to 3000 (3 seconds). The problem is corrected in CICS Transaction Service R2.2 base code.

However, there was one perplexing question – why is the logger issuing DASD I/O at 137 I/O per second. This is inconsistent with the use of CF logstreams when no data being offloaded. Post run investigation highlighted the fact the system programmer (me) had not removed the forward recovery log (defined as DASDONLY) from the IYOT3 region.

CF LOGSTREAMS - 32000 vs 64000

```
                                  BYT WRITTN  BYT WRITTN  BYT WRITTN                                          AVERAGE
                                   BY USERS   TO INTERIM    TO DASD    #WRITES  ---# WRITES COMPLETED------   BUFFER
-LOGSTREAM NAME- STRUCTURE NAME-- IXGWRITES   STORAGE                  INVOKED  TYPE1     TYPE2     TYPE3    SIZE

                                BYT DELETD  # DELETES  BYT DELETD  # DELETS  ---------------EVENT--------------
                                INTERIM ST  W/O DASD  INTERIM ST     W/     OFF- DASD STRC NTRY STG  STG   RE-
                                W/O DASD    WRITE     W/DASD       WRITE    LOAD SHFT FULL FULL THLD FULL  BLD
```

(23) 1 Region AKPFREQ 1000 LGDFINT 0 5M structure HIGHOFFLOAD 80 LOWOFFLOAD 50 MAXBUFSIZE 64000

```
12/12/02  4:40:00 AM  (SMF INTERVAL TIMESTAMP 'B8AA2DA8FA800000'X)
 IYOT1.CICS22.DFHLOG  LOG_JG2_20M 334322147   498134016          0    837644   834539      3105          0       399
                                  330752824      828725          0         0       64    0    0    0    0    0     0
```

Type	CICS TCB	DASD I/O	APPL%	# Tran/ Second	Average CPU/Task	Average Resp	Average JC Wait	Average File WT	Logr TCB	Logr SRB	DASD I/O	APPL%	Storage Frames	Idle Frames	Net Frames
CF64	195.8	3774	78.4	89.8	0.0060	0.1114	0.0203	0.0807	0.2	22.3	0.3	7.5	10734.4	858.09	9876.31

****

(24) 1 Region AKPFREQ 1000 LGDFINT 0 5M structure HIGHOFFLOAD 80 LOWOFFLOAD 50 MAXBUFSIZE 32000

```
12/12/02  3:40:00 AM  (SMF INTERVAL TIMESTAMP 'B8AA203FC0400000'X)
 IYOT1.CICS22.DFHLOG  LOG_JG2_20M 365200243   532300288          0    867099   864023      3077          0       421
                                  364781949      865815          0         0       62    0    0    0    0    0     0
```

Type	CICS TCB	DASD I/O	APPL%	# Tran/ Second	Average CPU/Task	Average Resp	Average JC Wait	Average File WT	Logr TCB	Logr SRB	DASD I/O	APPL%	Storage Frames	Idle Frames	Net Frames
CF32	207.5	4155	82.6	98.89	0.0058	0.1012	0.0204	0.070	0.2	19.4	0.3	6.5	11241.5	867.08	10373.92

CF LOGSTREAMS - 32000 vs 64000

Sample 23. 1 Region AKPFREQ 1000 LGDFINT 0 5M structure HIGHOFFLOAD 80 LOWOFFLOAD 50 MAXBUFSIZE 64000

```
12/12/02  4:40:00 AM  (SMF INTERVAL TIMESTAMP 'B8AA2DA8FA800000'X)
 IYOT1.CICS22.DFHLOG  LOG_JG2_20M 334322147   498134016          0    837644   834539      3105         0       399
                                  330752824      828725          0         0       64     0    0    0    0    0    0

Type CICS  DASD  APPL% # Tran/ Average   Average  Average Average  Logr Logr  DASD  APPL% Storage  Idle     Net
     TCB   I/O         Second  CPU/Task  Resp     JC Wait File WT  TCB  SRB   I/O         Frames   Frames   Frames
CF64 195.8 3774  78.4   89.8   0.0060    0.1114   0.0203  0.0807   0.2  22.3  0.3   7.5   10734.4  858.09   9876.31
**********************************************************************************************************
```

Sample 24. 3 Region AKPFREQ 1000 LGDFINT 0 5M structure HIGHOFFLOAD 80 LOWOFFLOAD 50 MAXBUFSIZE 32000
1 user journal (forward recovery log) allocated to the 20M CF structure

```
12/12/02  3:40:00 AM  (SMF INTERVAL TIMESTAMP 'B8AA203FC0400000'X)
 IYOT1.CICS22.DFHLOG  LOG_JG2_20M 365200243   532300288          0    867099   864023      3077         0       421
                                  364781949      865815          0         0       62     0    0    0    0    0    0

 Type CICS  DASD  APPL% # Tran/ Average   Average  Average Average  Logr Logr  DASD  APPL% Storage  Idle     Net
      TCB   I/O         Second  CPU/Task  Resp     JC Wait File WT  TCB  SRB   I/O         Frames   Frames   Frames
CF32 207.5 4155  82.6   98.89  0.0058    0.1012   0.0204  0.070    0.2  19.4  0.3   6.5   11241.5  867.08   10373.92
```

Samples 23 and 24 show a comparison between specifying MAXBUFSIZE(64000) and MAXBUFSIZE(32000) for the LOG_JG2_20M structure used for DFHLOG. It's interesting to note the throughput and response time improved with 32000 but the number of storage frames increased by 497.61 (2.03M).

The numbers indicate around a 10% drop in response time at a cost of 2M of storage in the logger dataspace.

This option is worth considering, but should be carefully evaluated in your environment. If there are files with large records (approaching 32K), which are defined recoverable, problems would be introduced by making the change. The additional storage usage is another factor to be evaluated.

User Journals

-LOGSTREAM NAME-	STRUCTURE NAME--	BYT WRITTN BY USERS IXGWRITES	BYT WRITTN TO INTERIM STORAGE	BYT WRITTN TO DASD	#WRITES INVOKED	---# WRITES COMPLETED------ TYPE1	TYPE2	TYPE3	AVERAGE BUFFER SIZE
		BYT DELETD INTERIM ST W/O DASD	# DELETES W/O DASD WRITE	BYT DELETD INTERIM ST W/DASD	# DELETS W/ WRITE	----------EVENT---------- OFF- LOAD	DASD SHFT	STRC FULL	NTRY FULL	STG THLD	STG FULL	RE- BLD
IYOT1.CICS22.DFHLOG	LOG_JG2_20M	357177308	511701760	0	813623	810599	3025		0		438	
		358391348	816443	0	0	64	0	0	0	0	0	0
IYOT1.IYO1.DFHJ03	*DASDONLY*	20472120	68173824	29848810	16645	0		0		0		1229
		0	0	87982080	21480	2	0	0	0	216	0	0

(25)1 Region AKPFREQ 1000 LGDFINT 0 20M structure HIGHOFFLOAD 80 LOWOFFLOAD 50 MAXBUFSIZE 60000 User journal DASD 25K
12/12/02 1:50:00 PM (SMF INTERVAL TIMESTAMP 'B8AAA8983B200000'X)

Type	CICS TCB	DASD I/O	APPL%	# Tran/ Second	Average CPU/Task	Average Resp	Average JC Wait	Average File WT	Logr TCB	Logr SRB	DASD I/O	APPL%	Storage Frames	Idle Frames	Net Frames
CF/D	203.5	4088	83	103	0.0058	0.1028	0.0233	0.0689	0.3	22	56.8	7.5	24210.6	859.08	23351.52

****

(26)1 Region AKPFREQ 1000 LGDFINT 0 20M structure HIGHOFFLOAD 80 LOWOFFLOAD 50 MAXBUFSIZE 60000 User jrnal DASD 2549
12/12/02 4:00:00 PM (SMF INTERVAL TIMESTAMP 'B8AAC5A6E4000000'X)

IYOT1.CICS22.DFHLOG	LOG_JG2_20M	349136595	504647168	0	814210	811392	2818		0		428	
		348285806	812041	0	0	64	0	0	0	0	0	0
IYOT1.IYO1.DFHJ03	*DASDONLY*	19479510	64868352	22010370	15837	0		0		0		1230
		0	0	65695744	16039	7	0	0	0	156	0	0

Type	CICS TCB	DASD I/O	APPL%	# Tran/ Second	Average CPU/Task	Average Resp	Average JC Wait	Average File WT	Logr TCB	Logr SRB	DASD I/O	APPL%	Storage Frames	Idle Frames	Net Frames
CF/D	202.1	3981	82.4	94.44	0.0059	0.1056	0.0229	0.0721	0.2	21.9	53.7	7.4	12520	859.08	11660.92

User Journal notes

Sample 25. 1 Region AKPFREQ 1000 LGDFINT 0 20M structure HIGHOFFLOAD 80 LOWOFFLOAD 50 User journal DASD 25K
12/12/02 1:50:00 PM (SMF INTERVAL TIMESTAMP 'B8AAA8983B200000'X)

```
 IYOT1.CICS22.DFHLOG  LOG_JG2_20M 357177308   511701760          0    813623   810599      3025          0       438
                                  358391348      816443          0         0       64     0    0    0    0    0    0
 IYOT1.IYO1.DFHJ03    *DASDONLY*   20472120    68173824   29848810     16645        0          0          0      1229
                                           0           0   87982080     21480        2     0    0    0  216    0    0

Type CICS  DASD  APPL% # Tran/ Average   Average  Average Average  Logr Logr  DASD  APPL% Storage  Idle     Net
     TCB   I/O         Second  CPU/Task  Resp     JC Wait File WT  TCB  SRB   I/O         Frames   Frames   Frames
CF/D 203.5 4088   83    103    0.0058    0.1028   0.0233  0.0689   0.3  22    56.8  7.5   24210.6  859.08   23351.52
```

In this sample there is a single CICS region, with DFHLOG allocated in a CF structure (LOG_JG2_20M). The region also has a forward recovery log (&applid..&SYSID..DFHJ03) for one of the files (FILEC). The user journal was defined with STG_SIZE of 25000 4K CIs - 102.4M). Compare this sample to sample 19A, which did not include a user journal. This comparison gives an indication of the potential resource impact when adding a user journal. The logger is issuing DASD I/O at the rate of 56.8 per second. The greatest impact to the system is the number of page frames allocated to the logger. 23351.52 amounts to 95.6M.

****

Sample 26 (26)1 Region AKPFREQ 1000 LGDFINT 0 20M structure HIGHOFFLOAD 80 LOWOFFLOAD 50 User jrnal DASD 2549
12/12/02 4:00:00 PM (SMF INTERVAL TIMESTAMP 'B8AAC5A6E4000000'X)

```
 IYOT1.CICS22.DFHLOG  LOG_JG2_20M 349136595   504647168          0    814210   811392      2818          0       428
                                  348285806      812041          0         0       64     0    0    0    0    0    0
 IYOT1.IYO1.DFHJ03    *DASDONLY*   19479510    64868352   22010370     15837        0          0          0      1230
                                           0           0   65695744     16039        7     0    0    0  156    0    0

Type CICS  DASD  APPL% # Tran/ Average   Average  Average Average  Logr Logr  DASD  APPL% Storage  Idle     Net
     TCB   I/O         Second  CPU/Task  Resp     JC Wait File WT  TCB  SRB   I/O         Frames   Frames   Frames
CF/D 202.1 3981  82.4  94.44   0.0059    0.1056   0.0229  0.0721   0.2  21.9  53.7  7.4   12520    859.08   11660.92
```

In this sample, using the formula discussed in the Common Problems presentation, I defined the user journal for an offload rate of approximately once a minute. Comparing sample 26 with sample 25, it is observed there is a decrease in the number of page frames required by the logger (11660.92 vs. 23351.52). This is a reduction of 47.9M storage required for a single logstream. Sample 25 depicts a common error, where the size of a user journal is defined to hold the same amount of data as was held a R410 system. The nature of user journals is they are never accessed by the online CICS region, so all data will be offloaded, i.e. logtail management does not apply to user journals. As such there is little benefit in keeping the data in the interim storage; it's better to reduce the size (STG_SIZE) so offloads are happening on a regular basis, and let the data reside in the offload datasets.

- **Understand the workload**
  - **It's the first step in the resolution of a performance problem**

Summary

If your system is running like a dog,
it might be time to check the log.
OFFLOAD values high and low,
are just a place to stub your toe.
Activity keypoints you must choose,
set it right and you can snooze
Interim storage the place blocks should be,
faster access your data to see.
Sizing, sizing how do you choose,
use Poughkeepsie tools and LSCU.
88 data and 0STAT,
help you get your response time back.
If your throughput is not what you prefer,
better check log defer.
LS_SIZE for the data out back,
needs to be greater than 2 tracks.
Long UOWs must be thinned,
so the DFHLOG can be trimmed.
RMF data and CICS PA,
may be the tools to save the day.
User journals take too much room,
size them small to offload soon.
Define the structure, define the stream,
get it right your system will scream.

References

- *OS/390 MVS Setting Up a Sysplex - GC28-1779*
  - **Lists other useful publications in Chapter 9:**
    - **Finding Information for CICS Log Manager**
    - **Finding Information for OPERLOG Log Stream**
    - **Finding Information for Logrec Log Stream**

- *MVS Diagnosis: Tools and Service Aids - LY28-1085, Chapter 13*

- *OS/390 MVS Assembler Services Guide - GC28-1762*
- *MVS Programming: Authorized Assembler Services Reference,*
  - *Volume 2 - GC28-1765*
  - *Lists return and reason codes and symbols*
    - *For example ---- 08 | xxxx0804 | Equate Symbol: IxgRsnCodeNoBlock*
- *OS/390 Parallel Sysplex Configuration Cookbook,*
  - *Vols. 1-3, SG24-2075, SG24-2076, SG24-2077*
    - *(See Vol. 2, SG24-2076 for Logger info)*

References

- *CICS Transaction Server for OS/390 Version 1 Release 3*
  - *Installation Guide Chapter 20*
- *CICS Transaction Server for OS/390 Version 2 Release 2*
  - *Installation Guide Chapter 24*
  - *Migration Guide*

- *RMF Monitor III - CF Reports*
- *www.IBM.COM/SUPPORT/techdocs*
  - *Flash W9609 MVS/ESA Parallel Sysplex Performance - LPAR Performance Considerations*
  - *Flash W99037 Performance Impacts of Using Shared ICF CPs*

Appendix

- SIZING formulas
- DFHLSCU Sample
- IXCMIAPU Samples
  - logstream definition
  - list logstream
- RMF Samples
- Displays

Sizing .....

- **Manual calculations**
  - **1st logstream in a structure**
    - ➤ **SIZE=(300000 + (#entries/highoffload% * (avg. #elements per write * X)))**
  - **subsequent logstreams in the structure**
    - ➤ **SIZE=(#entries/highoffload% * (avg. #elements per write * X))**
      - **this formula is used for DASD only logstreams and staging datasets**

X is 400 bytes when using 256 byte element sizes
X is 800 bytes when using 512 byte element sizes

number of write req(LGSWRITES)/ length of interval in seconds = log writes/second
log writes/sec * (length of longest running task in seconds) = number log writes = number CF entries
number bytes written (LGSBYTES) / LGSWRITES = average bytes per write
Element size is based on MAXBUFSIZE
MAXBUFSIZE 65276 or less -- element size is 256
MAXBUFSIZE greater than 65276 -- element size is 512
Average bytes per write/ element size = average number of elements per write

Divide the final number by 1024 to obtain the value needed for SIZE parameter in the CFRM policy definitions

For STAGING Datasets (DASD Logging)-

if the #bytes per write is <4096 - the number of CIs = #CF entries/highoffload %
if the #bytes per write is >4096 - the number of CIs = (#CF entries/highoffload %)*(#bytes per write/4096)

DFHLSCU - DASDONLY

```
**************** REPORT SUMMARY CONCLUSIONS ****************
The following summary contains the highest workload, based on the number of blocks written:-

  SEGMENT  00000001         DURATION  00000009 seconds
  TIME       6:34.5         DATE  2000.012
  NUMBER OF BLOCKS        : 00000566
  WRITES PER SECOND       : 00000025
  AVERAGE RECORD SIZE     : 00000183
  AKP INTERVAL            : 00000053

  TYPE  QUANTITY    NUMBER OF BYTES     5.1 EQUIVALENT
  FC    00000494    0000000000063232    0000000000110656
  JC    00000000    0000000000000000    0000000000000000
  TD    00000000    0000000000000000    0000000000000000
  TS    00000012    0000000000000600    0000000000002352
  KP    00000030    0000000000042648    0000000000002580
  RM    00000000    0000000000000000    0000000000000000
  SP    00000137    0000000000005316    0000000000007953
  Other 00000006    0000000000001023      *** NONE ***
  Total 00000673

  From this, an AVGBUFSIZE of 00601 was calculated.

************************************************************
 This section applies to DASD-only logstreams:-

 It is recommended to complete the following definition and use it to
 create a suitable logstream:
        DATA TYPE(LOGR) REPORT(NO)
        DEFINE LOGSTREAM NAME(userid.applid.DFHLOG)
               DASDONLY(YES)
               HIGHOFFLOAD(95)
               LOWOFFLOAD(19)
               STG_SIZE(2518)
               MAXBUFSIZE(64000)
```

DFHLSCU - CF Logstream

```
************************************************************
This section applies to CF logstreams:-

It is recommended to complete the following definition and use it to
create a suitable structure for this journal logstream:

       DATA TYPE(LOGR) REPORT(NO)
       DEFINE STRUCTURE NAME(LOG_DFHLOG_nnn) LOGSNUM(10)
              MAXBUFSIZE(64000) AVGBUFSIZE(601)

In addition, the space required within the Coupling Facility by
such a journal can be specified using the following definition:

       DATA TYPE(CFRM) REPORT(NO)
       STRUCTURE NAME(LOG_DFHLOG_nnn)
         INITSIZE(14336) SIZE(21248)
         PREFLIST(cf_name) REBUILDPERCENT(1)

The following is a typical definition of a logstream using some
default values, and some calculated from this utility:

       DATA TYPE(LOGR) REPORT(NO)
       DEFINE LOGSTREAM NAME(userid.applid.DFHLOG)
              STRUCTNAME(LOG_DFHLOG_nnn)
              HIGHOFFLOAD(95)
              LOWOFFLOAD(19)

If staging datasets are to be used for this logstream, the following value
is that calculated for the staging data set size. This assumes the worst
case where only this logstream is actively connected to the structure. If
more logstreams are to be connected in parallel, then this value should
be replaced by one obtained from dividing it by the number of streams.

              STG_SIZE(20864)

************************************************************
```

IXCMIAPU

- **Logstream definition**

```
//MSLDEFIN EXEC PGM=IXCMIAPU
//SYSPRINT DD SYSOUT=H,DCB=RECFM=FBA
//SYSIN    DD *
  DATA TYPE(LOGR) REPORT(YES)
  DEFINE LOGSTREAM NAME(IYOT1.DFHLOG)
  /*STRUCTNAME(LOG_JG)*/
  /*DUPLEXMODE(COND)*/
  /*STG_DUPLEX(YES)*/
  DASDONLY(YES)
  STG_SIZE(3500)
  HIGHOFFLOAD(85)
  LOWOFFLOAD(50)
  LS_DATACLAS(LS10MEG)
  LS_SIZE(500)
  HLQ(GRAUEL)
  DIAG(YES)
  MODEL(NO)
```

- **List logstreams**

```
//LOGLIST  JOB 935112,'CICS530 IYOT',MSGLEVEL=(1,1),
//         CLASS=A,MSGCLASS=H,NOTIFY=GRAUEL
//*
//*ROUTE PRINT WINVMC.GRAUEL
/*JOBPARM SYSAFF=MV55
//MSLDEFIN EXEC PGM=IXCMIAPU
//SYSPRINT DD SYSOUT=H,DCB=RECFM=FBA
//SYSIN    DD *
  DATA TYPE(LOGR) REPORT(NO)
  LIST LOGSTREAM NAME(IYOT*) DETAIL(YES)
  LIST STRUCTURE NAME(LOG_JG) DETAIL(YES)
```

RMF Notes

There are a number of SMF records produced which are helpful in understanding activity relating to logstreams, CF, and the logger address space. Sources of Data:

- SMF 88 Records
  - Single system in scope
- SMF 74.4 - CF Activity
  - Sysplex in scope
- SMF 74.1 - DASD Activity
  - Single or multi-system scope
- SMF 72 -Workload Activity

Using RMF Monitor III (or equivalent), displays can provide a wealth of information. Remember the data gatherer must be active.

Under TSO, once at the RMF Monitor III Primary Menu, selecting the S (Sysplex) option takes you to the RMF Sysplex Report Selection Menu. Reports 5 (CFOVER – Coupling Facility overview), 6 (CFSYS – Coupling Facility systems), and 7 (CFACT – Coupling Facility activity) contain information pertaining to the coupling facilities in the sysplex. Refer to the samples on the pages titled *RMF Monitor III– CF Reports*.

To analyze activity on the Staging and/or offload (log) datasets, display the datasets in ISPF using HLQ.logstream name. This will display both staging datasets, if in use, and offload datasets. There are a number of tools available that provide device activity and response information. One example is to use option 3.3 from the RMF Monitor III main menu. Refer to the sample on the page titled *Staging and Offload datasets*.

Another important area is the amount of storage in use by the logger (IXGLOGR) address space, including its dataspace. The storage used in the dataspace is directly proportional to the amount of data held in all logstreams (i.e. in the coupling facility structures or on the staging dataset for a DASDONLY logstream). The larger the logstreams, the more data must be contained in the dataspace.

The logger address space has one of the highest priorities in the system, so its page reference pattern may preempt other jobs, such as CICS, to the point it's possible to induce response time problems due to paging. Sizing is critical in the operation of a logstream. It can also have a significant impact on the system as a whole, if logstreams are made excessively large. Refer to the sample on the page titled *RMF Storage Display*.

RMF Monitor III - CF Reports

CFOVER Coupling Facility Overview Report (S.5)

```
                    RMF 2.7.0  CF Overview       - PLEXB              Line 1 of 2
 Samples: 100      Systems: 3    Date: 01/30/00  Time: 20.25.00  Range: 100   Sec

 ---- Coupling Facility -----    ---- Processor -----   Request    -- Storage --
 Name       Type    Model Level    Util% Defined Effect   Rate        Size    Avail

 SSCF04     9674      C05     8     3.3       2     1.2      317.0    1012M    909M
 SSCF05     9674      C05     8     0.0       2     1.2               1012M   1011M
```

The CFOVER report (S.5 from the main menu) shows the coupling facilities in the sysplex, the machine type (9674), model and level. The storage in use is the difference between the storage size and the amount available.

Under the processor heading we see there are 2 CPs defined for each of the 2 coupling facilities; however, it is noted the effective processor is only 1.2, because CPs are being shared. If the CPs were dedicated the effective would also be 2. This is a performance consideration because 40% of the time a CP is not available for use. The net result is higher average service times, and high standard deviations. Refer to WSC Flash W9609 *MVS/ESA Parallel Sysplex Performance - LPAR Performance.*

CFSYS Coupling Facility Systems Report (S.6)

```
                    RMF 2.7.0  CF Systems       - PLEXB              Line 1 of 6
 Samples: 100      Systems: 3    Date: 01/30/00  Time: 20.25.00  Range: 100   Sec

 CF Name    System       Subch    -- Paths --    -- Sync ---    ------- Async -------
                         Delay    Avail Delay     Rate  Avg      Rate  Avg  Chng  Del
                           %                      %            Serv          Serv   %     %
 SSCF04     MV55           0.0      4     0.5    167.3   589     20.0  1326  0.0   0.0
            MV56           0.0      4     2.4     78.6   620     17.8  1205  0.1   0.2
            MV57           0.0      4     3.0     14.2   491     19.0  1464  0.0   0.9
 SSCF05     MV55                    4
            MV56                    4
            MV57                    4
```

The CFSYS report shows which MVS images are connected to the CF, and request distribution, service times and any path delays.

RMF Monitor III - CF Reports

CFACT Coupling Facility activity (S.7)

```
                    RMF 2.7.0  CF Activity      - PLEXB          Line 1 of 12
Samples: 100     Systems: 3    Date: 01/13/00  Time: 23.21.40  Range: 100   Sec

CF: ALL          Type  ST System      --- Sync ---     --------- Async --------
                                        Rate    Avg       Rate    Avg   Chng   Del
Structure Name                                 Serv              Serv   %      %
DSN510PB_LOCK1   LOCK      *ALL          2.0    151        0.0       0   0.0    0.0
DSN510PB_SCA     LIST      *ALL          2.0    174        0+      267  50.0   50.0
HASPCKPT         LIST      *ALL          4.7    383        7.8     901   0.0   20.7
IEFAUTOS         LIST      *ALL          0.0      0        0.0       0   0.0    0.0
IGWLOCK00        LOCK      *ALL          0.0      0        0.0       0   0.0    0.0
ISGLOCK          LOCK      *ALL         31.6    149        0.0       0   0.0    0.0
ISTGENERIC       LIST      *ALL         18.3    190        0+      225  50.0   25.0
ISTMNPS          LIST      *ALL          0.0      0        0.0       0   0.0    0.0
IXCDEF           LIST      *ALL          0.0      0       45.5     932   0.0    8.0
LOG_JG           LIST      *ALL        114.1    277        3.5    1916   1.4    1.7
LOG_RRS_TEST     LIST      *ALL          0.0      0        6.2    1928   0.0    0.0
SYSIGGCAS_ECS    CACHE     *ALL         11.9    233        0+      733  50.0   50.0
```

The CFACT report shows the activity by structure, showing the rate per second along with the average service times. Please note these are SYSPLEX wide reports.

Staging and Offload datasets

```
Dataset display using ISPF

DSLIST - Data Sets Matching GRAUEL.IYOT1.DFHLOG                    Row 1 of 2

Command - Enter "/" to select action                 Message          Volume
-------------------------------------------------------------------------
         GRAUEL.IYOT1.DFHLOG.A0000002                                 *VSAM*
         GRAUEL.IYOT1.DFHLOG.A0000002.DATA                            PBDA15
****************************** End of Data Set list ********************************
```

To display information about the device, in this case volume PBDA15, select RMF Monitor III option 3.3. Find the VOLSER in the timeframe needed, and put the cursor under the VOLSER, hit enter.

```
                    RMF 2.7.0  Data Set Delays - Volume          Line 1 of 2

Samples: 100     System: MV55  Date: 01/13/00  Time: 23.21.40  Range: 100   Sec

------------------------- Volume PBDA15 Device Data --------------------------
Number:     101F         Active:       1%        Pending:    0%     Average Users
Device:     33903        Connect:      1%        Delay DB:   0%        Delayed
Shared:     Yes          Disconnect:   0%        Delay CU:   0%          0.0
                                                 Delay DP:   0%
-------------- Data Set Name ---------------    Jobname   ASID  DUSG% DDLY%
-- N/A --                                       DRAKEX    0026      1     0
GRAUEL.IYOT1.DFHLOG.A0000002.DATA               IXGLOGR   0021      1     0
```

RMF Storage Displays

From the RMF Monitor III main menu select option 3 (resources) then option 6 (storage)

```
                    RMF 2.7.0  Storage Delays                  Line 13 of 221

Samples: 100     System: MV55  Date: 01/13/00  Time: 23.21.40  Range: 100   Sec
             Service  DLY    ------- % Delayed for ------   -- Working Set --
Jobname   C  Class     %     COMM  LOCL  SWAP  OUTR  OTHR   Central  Expanded

ANTMAIN   S  SYSTEM    0        0     0     0     0     0      4542         0
ANTAS000  S  STCUSER   0        0     0     0     0     0       672         0
OMVS      S  SYSTEM    0        0     0     0     0     0     17811         0
IEFSCHAS  S  SYSTEM    0        0     0     0     0     0        43         0
JESXCF    S  SYSTEM    0        0     0     0     0     0       513         0
ALLOCAS   S  SYSTEM    0        0     0     0     0     0       177         0
IOSAS     S  SYSTEM    0        0     0     0     0     0      1114         0
IXGLOGR   S  SYSTEM    0        0     0     0     0     0      6721         0
LLA       S  STC       0        0     0     0     0     0      2062         0
BPXOINIT  S  SYSTEM    0        0     0     0     0     0       160         0
SMF       S  SYSTEM    0        0     0     0     0     0       922         0
DD55SPAS  S  STCUSER   0        0     0     0     0     0       700         0
```

Neat Displays ..

D CF

The display coupling facility can be used to show key information, such as the amount of storage in use by structures, free space, channel paths and microcode level. However, one of the most important bits of information is whether or not the CF is volatile. If the CF is volatile, the logger will allocate staging datasets for the duplexed copy of the data, rather than using the dataspace.

```
COUPLING FACILITY 009674.IBM.51.000000068441
                  PARTITION: 4  CPCID: 00
                  CONTROL UNIT ID: 0900
NAMED SSCF04
COUPLING FACILITY SPACE UTILIZATION
 ALLOCATED SPACE                 DUMP SPACE UTILIZATION
  STRUCTURES:     119040 K        STRUCTURE DUMP TABLES:        0 K
  DUMP SPACE:      12032 K                  TABLE COUNT:        0
 FREE SPACE:      904704 K       FREE DUMP SPACE:           12032 K
TOTAL SPACE:     1035776 K      TOTAL DUMP SPACE:           12032 K
                               MAX REQUESTED DUMP SPACE:        0 K
   VOLATILE:           NO        STORAGE INCREMENT SIZE:      256 K
    CFLEVEL:            8

COUPLING FACILITY SPACE CONFIGURATION
                          IN USE          FREE          TOTAL
CONTROL SPACE:          131072 K      904704 K     1035776 K
NON-CONTROL SPACE:           0 K           0 K           0 K

SENDER PATH       PHYSICAL            LOGICAL       CHANNEL TYPE
        09         ONLINE              ONLINE        CFS
        85         ONLINE              ONLINE        CFS
        A5         ONLINE              ONLINE        CFS
        DD         ONLINE              ONLINE        CFS
COUPLING FACILITY DEVICE     SUBCHANNEL     STATUS
                  FF00         05C6         OPERATIONAL/IN USE
                  FF01         05C7         OPERATIONAL/IN USE
                  FF02         05C8         OPERATIONAL/IN USE
                  FF03         05C9         OPERATIONAL/IN USE
                  FF04         05CA         OPERATIONAL/IN USE
                  FF05         05CB         OPERATIONAL/IN USE
                  FF06         05CC         OPERATIONAL/IN USE
                  FF07         05CD         OPERATIONAL/IN USE

The above data is repeated for each CF connected to the MVS image
```

Neat Displays ...

```
D LOGGER,C,LSNAME=IYOT1.DFHLOG,D

IXG601I   19.59.31  LOGGER DISPLAY 791
CONNECTION INFORMATION BY LOGSTREAM FOR SYSTEM MV55
LOGSTREAM                  STRUCTURE        #CONN  STATUS
IYOT1.DFHLOG               LOG_JG           000001 IN USE
STG DS: NO
    JOBNAME: IYOT1     ASID: 0073
      R/W CONN: 000000 / 000001
      RES MGR./CONNECTED: *NONE*    / NO
      IMPORT CONNECT: NO

D XCF,STRUCTURE,STRNAME=LOG_JG

 IXC360I  19.25.28  DISPLAY XCF 517
STRNAME: LOG_JG
 STATUS: ALLOCATED
  POLICY SIZE     : 35000 K
  POLICY INITSIZE: 20000 K
  REBUILD PERCENT: 1
  DUPLEX          : DISABLED
  PREFERENCE LIST: SSCF04
  ENFORCEORDER    : NO
  EXCLUSION LIST IS EMPTY
 ACTIVE STRUCTURE
 ALLOCATION TIME: 02/10/2000 14:55:36
  CFNAME          : SSCF04
  COUPLING FACILITY: 009674.IBM.51.000000068441
                     PARTITION: 4   CPCID: 00
  ACTUAL SIZE     : 20224 K
  STORAGE INCREMENT SIZE: 256 K
  PHYSICAL VERSION: B3942B1D ED66AF02
  LOGICAL  VERSION: B3942B1D ED66AF02
  SYSTEM-MANAGED PROCESS LEVEL: 8
  DISPOSITION     : DELETE
  ACCESS TIME     : 0
  MAX CONNECTIONS: 32
  # CONNECTIONS   : 2
 CONNECTION NAME  ID VERSION  SYSNAME  JOBNAME  ASID STATE
 IXGLOGR_MV55     01 00010081 MV55     IXGLOGR  0015 ACTIVE
 IXGLOGR_MV56     02 00020017 MV56     IXGLOGR  0015 ACTIVE
```

D LOGGER,C,LSNAME=IYOT1.DFHLOG,D is requesting a detailed display about logstream IYOT1.DFHLOG. The response indicates the logstream is in use, with one connection, IYOT1, ASID 73, in MV55. The interim storage is contained in a coupling facility structure (LOG_JG) and there no are no staging datasets at this time.

A D XCF,STRUCTURE,STRNAME=LOG_JG command is requesting information about a given coupling facility structure, in this case LOG_JG. The response shows the initial size is 20M expandable to 35M. It is located in a CF which is a 9674 having 4 partitions. The maximum connections is 32, i.e. 32 MVS images. There are currently 2 connections the logger address space on MVS 55 (ASID 15) and the logger address space on MVS 56 (ASID 15).

Neat Displays ....

```
d xcf,cf,cfname=sscf04

 RESPONSE=MV55
  IXC362I  19.30.27  DISPLAY XCF 528
  CFNAME: SSCF04
    COUPLING FACILITY      :   009674.IBM.51.000000068441
                               PARTITION: 4   CPCID: 00
    POLICY DUMP SPACE SIZE:    12000 K
    ACTUAL DUMP SPACE SIZE:    12032 K
    STORAGE INCREMENT SIZE:    256 K

    CONNECTED SYSTEMS:
    MV55      MV56      MV57

    STRUCTURES:
    DSN510PB_LOCK1    DSN510PB_SCA      HASPCKPT     IEFAUTOS
    IGWLOCK00         ISGLOCK           ISTGENERIC   ISTMNPS
    IXCDEF            LOG_DFHLOG_001    LOG_JG       LOG_SYSTEST_001
    SYSIGGCAS_ECS
```

The D XCF,CF,CFNAME=SSCF04 asks for a display of coupling facility SSCF04. The response shows the CF is a 9674 S/N 68441. The are 4 partitions and 3 connected MVS images (MV55, MV56, MV57). A list of the active structures is provided.

```
d xcf,couple,type=logr     (OS/390 R2.8 and up)

IXC358I  20.01.21  DISPLAY XCF 804
LOGR COUPLE DATA SETS
PRIMARY     DSN: SYS1.SYSPLEXB.PLOGR
            VOLSER: PBXCF1      DEVN: 1008
            FORMAT TOD          MAXSYSTEM
            06/03/1998 10:32:49         8
            ADDITIONAL INFORMATION:
             LOGR COUPLE DATA SET FORMAT LEVEL: HBB6603
             LSR(4000) LSTRR(50) DSEXTENT(100)
ALTERNATE   DSN: SYS1.SYSPLEXB.ALOGR
            VOLSER: PBXCF2      DEVN: 1028
            FORMAT TOD          MAXSYSTEM
            06/03/1998 10:37:41         8
            ADDITIONAL INFORMATION:
             LOGR COUPLE DATA SET FORMAT LEVEL: HBB6603
             LSR(4000) LSTRR(50) DSEXTENT(100)
LOGR IN USE BY ALL SYSTEMS
```

The D XCF,COUPLE,TYPE=LOGR command available with OS/390 R2.8 and above, provides information about the logger couple dataset. It gives the primary and alternate dataset names and the format level HBB6603 (OS/390 R1.3).

ISMF Displays

```
 ------------------------------------------------  DATA SET LIST OPTION MENU
OPTION  ===> 1

   1  ISMF     - Interactive Storage Management Facility
   2  DAF      - Data Access Facility
   3  DSLIST   - PDF Dataset List
   4  NETVFTP  - Netview FTP


                  ISMF PRIMARY OPTION MENU - DFSMS/MVS 1.5

Select one of the following options and press Enter:

0  ISMF Profile            - Change ISMF User Profile
1  Data Set                - Perform Functions Against Data Sets
2  Volume                  - Perform Functions Against Volumes
3  Management Class        - Specify Data Set Backup and Migration Criteria
4  Data Class              - Specify Data Set Allocation Parameters
5  Storage Class           - Specify Data Set Performance and Availability
9  Aggregate Group         - Specify Data Set Recovery Parameters
L  List                    - Perform Functions Against Saved ISMF Lists
R  Removable Media Manager - Perform Functions Against Removable Media
X  Exit                    - Terminate ISMF


Enter Selection or Command ===> 1
```

ISMF Displays ...

Option 1 - Dataset list Example showing the offload dataset for logstream IYOT2.DFHLOG

```
                                        DATA SET LIST
                                                                        Entries 1-1 of 1
Enter Line Operators below:                                      Data Columns 3-12 of 39
                        ALLOC    ALLOC    % NOT  COMPRESSED  % USER DATA  NUM  ALLOC   SEC      DS     REC
       DATA SET NAME    SPACE    USED     USED   FORMAT      REDUCTION    EXT  UNIT    ALLOC    ORG    FMT
------------(2)-------- --(3)--  --(4)--  -(5)-  ---(6)----  ----(7)----  (8)  -(9)-  --(10)--  (11)   (12)-
GRAUEL.IYOT2.DFHLOG.     -------  -------    ---  ---                  --   ---    ---    --------  ---    -----
  A0000000
-------  ------  --------  --------- ------------  BOTTOM  OF  DATA  ------------  ---------  -------  ----
 Enter Line Operators below:                                       Data Columns 16-24 of 39
                       BLOCK    VOLUME  MULT  DEVICE     CREATE      EXPIRE      LAST REF   LAST BACKUP  CHG
      DATA SET NAME    UNUSED   SERIAL  VOL   TYPE       DATE        DATE        DATE       DATE         IND
------------(2)------- -(16)--  -(17)-  (18)  -(19)--  ---(20)---  ---(21)---  ---(22)---  ---(23)----  (24)
GRAUEL.IYOT2.DFHLOG.    -------  MIGRAT  NO    3490     2000/03/20  0000/00/00  ----------  ----------   ---
  A0000000
------  ------  --------  --------- ------------  BOTTOM  OF  DATA  ------------  ---------  --------  ----

Enter Line Operators below:                                                  Data Columns 25-32 of 39
                        DATA      MANAGEMENT   STORAGE                DATA SET     DATA SET   NUM OF   ENTRY
     DATA SET NAME      CLASS NAME CLASS NAME  CLASS NAME   OWNER     ENVIRONMENT  NAME TYPE  STRIPES  TYPE
  --------(2)-------  ---(25)---  ---(26)---  ---(27)---  --(28)--  ---(29)----  --(30)---  -(31)--  --(32)--
GRAUEL.IYOT2.DFHLOG.    VSAM100      STANDARD    STANDARD    --------  MANAGED      OTHERS           --  NONVSAM
  A0000000
```

Using option 1 (Dataset List), if a data class is associated with the dataset, it is shown in column 25.
The management class is given in column 26, with the storage class listed in column 27.
If the dataset is HSM managed, it is noted in column 29. The location of the dataset is given in column 17.

ISMF Displays....

```
 OPTION 4 -        DATA CLASS LIST          Example showing class VSAM100
                                                                  Entries 1-1 of 1
                                                                  Data Columns 3-14 of 39
CDS Name : ACTIVE

Enter Line Operators below:

LINE       DATACLAS                                                AVG    SPACE    SPACE      SPACE      RETPD OR
OPERATOR   NAME      RECORG  RECFM  LRECL  KEYLEN  KEYOFF  AVGREC  VALUE  PRIMARY  SECONDARY  DIRECTORY  EXPDT
--(1)----  --(2)---  -(3)--  -(4)-  -(5)-  -(6)--  -(7)--  -(8)--  -(9)-  -(10)--  --(11)---  --(12)---  ---(13)---
           VSAM100   LS      ----   -----    ---   -----  U        6160      100          10     ------  ----------
---------  --------  --------  ----------  ----------  BOTTOM  OF  DATA --------  ----------  ----------  --------


 LINE       DATACLAS VOLUME  ADDITIONAL                        CISIZE  % FREE    % FREE    SHARE     SHARE
 OPERATOR   NAME     COUNT   VOLUME AMT  IMBED  REPLICATE      DATA    SPACE CA  SPACE CI  XREGION   XSYSTEM
---(1)----  --(2)--- -(14)-  ---(15)---  (16)-  --(17)---      -(18)-  --(19)--  --(20)--  -(21)--   -(22)--
           VSAM100      --  ---------  ---    ---            -----       ---       ---         3         3
----------  --------  --------  ----------  BOTTOM  OF  DATA --------  ----------  ----------  --------

Using option 4 (Data Class list), the SHAREOPTIONs for all datasets in the class are shown in columns 21 and 22.


OPTION 3 -        MANAGEMENT CLASS LIST     Example showing management class STANDARD
                                                                  Entries 1-1 of 1
                                                                  Data Columns 3-12 of 40
CDS Name : ACTIVE

Enter Line Operators below:

LINE       MGMTCLAS EXPIRE     EXPIRE        RET    PARTIAL      PRIMARY   LEVEL 1  CMD/AUTO  # GDG ON  ROLLED-OFF
OPERATOR   NAME     NON-USAGE  DATE/DAYS     LIMIT  RELEASE      DAYS      DAYS     MIGRATE   PRIMARY   GDS ACTION
--(1)----  --(2)--- ---(3)---  ---(4)----  --(5)--  ----(6)----  ---(7)--  --(8)--  --(9)---  --(10)--  ---(11)---
           STANDARD       720     NOLIMIT        0  NO               100         0  BOTH              1  EXPIRE
---------  --------  --------  ----------  ----------  BOTTOM  OF  DATA --------  ----------  ----------  --------

Option 3 (Management Class list), column 9 indicates if the dataset may be migrated either automatically or via
                                  a manual command.
```

SMF 88 Heading Guide

This page can be used as a 'heading' guide when examining IXGRPT1 reports.

```
                                BYT WRITTN  BYT WRITTN  BYT WRITTN                                          AVERAGE
                                 BY USERS   TO INTERIM    TO DASD    #WRITES  ---# WRITES COMPLETED------   BUFFER
-LOGSTREAM NAME- STRUCTURE NAME-- IXGWRITES   STORAGE                 INVOKED   TYPE1     TYPE2     TYPE3    SIZE

                                BYT DELETD  # DELETES   BYT DELETD  # DELETS   ---------------EVENT---------------
                                INTERIM ST   W/O DASD   INTERIM ST     W/     OFF- DASD STRC NTRY STG  STG   RE-
                                W/O DASD     WRITE      W/DASD       WRITE    LOAD SHFT FULL FULL THLD FULL  BLD
```