C46

Evaluating Performance of the MVS Logger in a CICS Environment

Jim Grauel

Las Vegas, Nevada June 12-16, 2000

Agenda

- Performance related parameters
- Tools
  - DFH0STAT
  - SMF 74 records
  - SMF 88 records

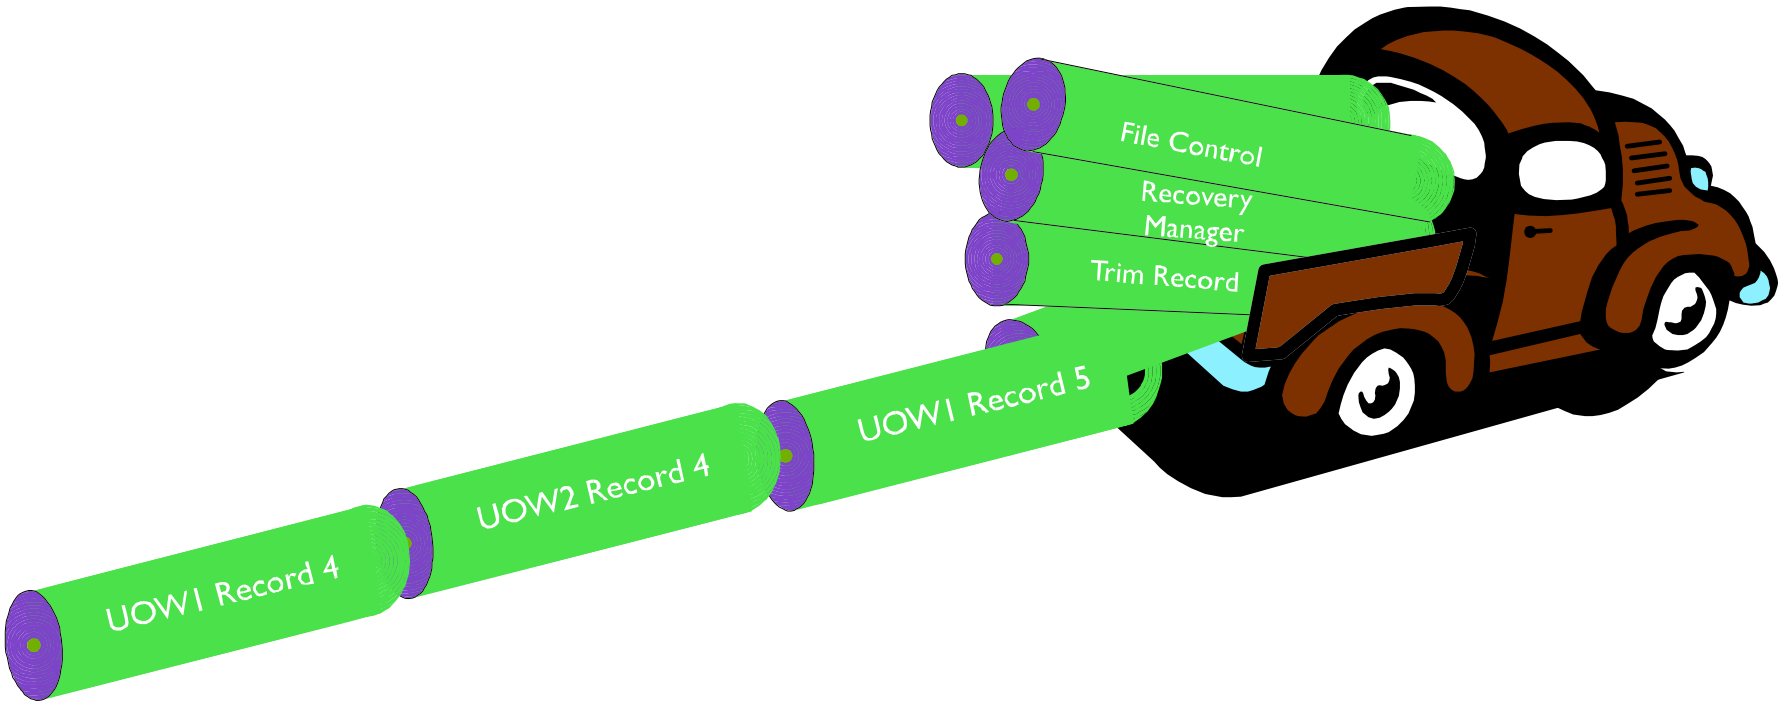

CICS Parameters

- 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 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
  - use CEMT I SYS to view/change

Logstream offloads

- Optimize use of the interim storage
- MVS Offload processing
  - preferred
    - triggered by HIGHOFFLOAD value
    - when the last connector to a log stream disconnects
  - non-preferred
    - having 90% of the CF structure entries in use
    - logstream full
    - staging dataset full
    - structure rebuild
    - recovery processing
  - physically delete any logically deleted records
    - IXGDELETs issued for log tail management on DFHLOG and DFHSHUNT
  - if required, move oldest records to the offload (LOG) dataset until LOWOFFLOAD value is reached
  - perform housekeeping in the dataspace

Define a CF Logstream

```
//DEFSTR JOB CLASS=A,MSGCLASS=A
//POLICY EXEC PGM=IXCMIAPU
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
  DATA TYPE(LOGR) REPORT(YES)
 DEFINE LOGSTREAM NAME(IYOT1.DFHLOG)
 AUTODELETE=NO  *
 HIGHOFFLOAD(80) *

 HLQ(GRAUEL)
 LOWOFFLOAD(40)  *

 LS_DATACLAS(LS10MEG) *
 LS_SIZE(500)  *

 MODEL(NO)
 RETPD=0  *

 STG_DUPLEX(YES) DUPLEXMODE=(COND) *



 STG_SIZE(9000) *
 STRUCTNAME(LOG_JG)  *
  * parms which affect performance
```

Matches name in CICS Journal Model

Always NO for DFHLOG and DFHSHUNT

Threshold, expressed as a % of logstream space, when offloading is to take place

High level qualifier - offload dataset name

Value, expressed as a % of the space used by the logstream, for the offload target

SMS data class used for DASD offload

Allocation units for the offload dataset(s) specified in 4K control intervals

Retention period must be zero for DFHLOG and DFHSHUNT

Log writes are to be duplexed to the Staging Dataset if the CF becomes volatile or failure dependent (CF and MVS in same CEC)

No. of 4K control intervals in Staging dataset

Structure that will contain the data

Define CF Logstream Notes

Many parameters affect logstream performance, some more than others.

AUTODELETE and RETPD can have a disastrous effect on the DFHLOG and DFHSHUNT if specified other than AUTODELETE(NO) and RETPD(0). With AUTODELETE(YES) and RETPD>0, even though CICS will attempt log tail management, all data will be offloaded to the offload datasets and held for the number of days specified for RETPD. AUTODELETE(YES) lets the logger (rather than CICS) decide when to delete the data. When a new offload dataset is allocated and AUTODELETE(YES) is specified, the logger will delete the data on the old offload dataset. If CICS needs the data for backout, the result will be an 804 return code and CICS will terminate with a DFHLG0772.

The HIGHOFFLOAD parameter, in conjunction with the size of the logstream, has a major effect on the amount of virtual storage used by the Logger. Before data is written to the Coupling Facility (CF), it is written to a buffer in the logger dataspace. If staging datasets are used, the data is written to the staging dataset rather than the dataspace.

If the HIGHOFFLOAD value is too high, there may not be enough room in the logstream to accommodate data being written to the logstream during offload processing. This can lead to an entry or structure full condition --- which causes log writes to be suspended for 3 seconds.

HIGHOFFLOAD should be set at 80 - 85% for DFHLOG, DFHSHUNT and user journals.
LOWOFFLOAD should be set between 40 and 60% for DFHLOG and DFHSHUNT, 0 for user journals.

If the extent size of the LS_DATACLAS (or the value specified in LS_SIZE) is too small, frequent DASD shifts, allocation of a new offload dataset, will occur. Frequent DASD shifts have a negative effect on performance.

The DSEXTENT value specified in the LOGR couple dataset defines the number of offload (log) dataset directories which can be used in the sysplex. Each directory can point to a maximum of 168 offload datasets. Prior to OS/390 R1.3, the number of extents is limited to 1, with R1.3 the number is limited only by the amount of DASD available.

LS_SIZE should be specified large enough to contain several offloads, possibly a day's worth. DFHLOG and DFHSHUNT should only offload a minimal amount of data; however, all data is offloaded for user journals.

STG_SIZE and logstream structure size (specified in the structure definition) are specified based on the amount of data to be contained in the logstream. The rule of thumb is it must be large enough to hold the data written during an activity keypoint interval plus the length of time of the longest running unit of work.

Define CF Logstream Notes ...

DFHLSCU should be run against the R410 system journal and the output used as a STARTING POINT for STG_SIZE, and the size of the Coupling Facility structure.

Note - if LS_SIZE is not specified in the logstream definition and an extent size is not specified in the LS_DATACLAS, the value is taken from the ALLOCxx parmlib member or set via an ACS (Automatic Class Selection) routine. The default value in ALLOCxx is 2 tracks. Refer to the *MVS Initialization and Tuning Guide*.

This also applies to the staging dataset size (STG_SIZE). When staging datasets are used with a coupling facility logstream, STG_SIZE must be specified large enough to hold at least as much data as the CF logstream. Otherwise, the offloads will be triggered by the HIGHOFFLOAD percentage of the staging dataset rather than the CF logstream.

Data is written to the staging dataset in 4096 byte increments, regardless of the buffer size.

STG_DUPLEX(YES) with DUPLEXMODE=(COND) means, should the CF become volatile, or exist in a failure dependent configuration, the data will be duplexed to the staging dataset, otherwise it is duplexed to buffers in the logger dataspace. A failure dependent configuration is when the Coupling Facility LPAR and the LPAR running OS/390 reside in the same CEC. Duplexing to the staging dataset means the cost of an I/O will be incurred for each write.

The structure used for the logstream will affect performance. Key factors include the number and characteristics of the logstreams which will share the structure. It is best to limit the number (LOGSNUM specified on the structure definition) to 10. Logstreams in a structure should have like characteristics. For example, TORs, AORs, and FORs typically have different logstream record size and volumes.

It is also important to remember logstream staging and offload (log) datasets are single extent VSAM linear datasets, the shareoptions MUST be specified '3,3'. If the shareoptions are anything other than '3,3' there is a risk the logger will be unable to read offloaded data and post CICS with return code 8, reason code 804 (*IxgRsnCodeNoBlock*). This will cause CICS to abend with a DFHLG0772.

When the offload process is triggered, the offload may be performed by the logger address space on a different MVS image. If the offload requires the movement of data to the log (offload) datasets, and the current dataset fills, a new dataset will be allocated. Unless the shareoptions are specified 3,3 the logger address space on the MVS image where CICS is running may not be able to access the dataset.

Define Logstream - DASDONLY

```
//DEFSTR JOB CLASS=A,MSGCLASS=A
//POLICY EXEC PGM=IXCMIAPU
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
  DATA TYPE(LOGR) REPORT(YES)
DEFINE LOGSTREAM NAME(IYOT2.DFHLOG)
AUTODELETE=NO  *
DASDONLY(YES)
HIGHOFFLOAD(80)  *

HLQ(GRAUEL)
LOWOFFLOAD(40)  *

LS_DATACLAS(LS10MEG) *
LS_SIZE(4500)  *

MAXBUFSIZE=65532 *
MODEL(NO)
RETPD=0  *

STG_SIZE(9000) *
```

Matches name defined in Journal Model

Always NO for DFHLOG and DFHSHUNT

C. F. structures will not be used

Threshold, expressed as a % of logstream space, when offloading is to take place

High level qualifier - offload dataset name

Value, expressed as a % of the space used by the logstream, of the offload target

SMS dataclass used for DASD offload

Allocation units for the offload dataset(s) specified in 4K control intervals

Max buffer size of blocks written to the log

Retention period must be zero for DFHLOG and DFHSHUNT

The number of 4K control intervals in the Staging Dataset.

* parms which affect performance

Define DASDONLY Logstream Notes

As with CF logstreams, many of the parameters affect logstream performance.

AUTODELETE and RETPD can have a disastrous effect on the DFHLOG and DFHSHUNT if specified other than AUTODELETE(NO) and RETPD(0). With AUTODELETE(YES) and RETPD>0, even though CICS will attempt log tail management, all data will be offloaded to the offload datasets and held for the number of days specified for RETPD. AUTODELETE(YES) lets the logger (rather than CICS) decide when to delete the data. When a new offload dataset is allocated and AUTODELETE(YES) is specified, the logger will delete the data on the old offload dataset. If CICS needs the data for backout, the result will be an 804 return code and CICS will terminate with a DFHLG0772.

The HIGHOFFLOAD parameter in conjunction with the size of the logstream has a major effect on the amount of storage used by the Logger. Before data is written to the logstream, it is written to the logger dataspace.

If the HIGHOFFLOAD value is too high, there will not be enough room in the logstream to accommodate data being written to the logstream during offload processing. This can lead to a staging dataset full condition --- which causes log writes to be suspended. It's very important to factor in the peak times when calculating the logstream size and HIGHOFFLOAD values.

LOWOFFLOAD should be set between 40 and 60% for DFHLOG and DFHSHUNT, 0 for user journals.

If the extent size of the LS_DATACLAS (or the value specified in LS_SIZE) is too small, frequent DASD shifts, allocation of a new offload dataset, will occur. Frequent DASD shifts have a negative effect on performance.

The DSEXTENT value specified in the LOGR couple dataset defines the number of offload (log) dataset directories which can be used in the sysplex. Each directory can point to a maximum of 168 offload datasets. Prior to OS/390 R1.3, the number of extents is limited to 1, with R1.3 the number is limited only by the amount of DASD available.

LS_SIZE should be specified large enough to contain several offloads, possibly a day's worth. DFHLOG and DFHSHUNT should only offload a minimal amount of data; however, all data is offloaded for user journals.

STG_SIZE is specified based on the amount of data to be contained in the logstream. The rule of thumb is it must be large enough to hold the data written during an activity keypoint interval plus the length of time of the longest running unit of work.

Define DASDONLY Logstream Notes ...

DFHLSCU should be run against the R410 system journal and the output used as a STARTING POINT for STG_SIZE.

Note - if LS_SIZE and STG_SIZE are not specified in the logstream definition and an extent size is not specified in the LS_DATACLAS, the value is taken from the ALLOCxx parmlib member or set via an ACS (Automatic Class Selection) routine. The default value in ALLOCxx is 2 tracks. Refer to the *MVS Initialization and Tuning Guide*.

MAXBUFSIZE may be specified for a DASDONLY logstream, it defines the largest block that can be written to the logstream. The default value is 65532.

STG_DUPLEX(YES) with DUPLEXMODE=(COND) are not applicable to DASDONLY logstreams.

With the exception of DFHLOG and DFHSHUNT, DASDONLY logstreams can be shared within the MVS image. User journals and forward recovery logs fall into this category.
Please note, this could have an impact on IOSQ time to the staging datasets.

It also important to remember logstream staging and offload (log) datasets are single extent VSAM linear datasets, the shareoptions MUST be specified '3,3'. If the shareoptions are anything other than '3,3', there is a risk the logger will be unable to read offloaded data and post CICS with return code 8, reason code 804 (IxgRsnCodeNoBlock). This will cause CICS to abend with a DFHLG0772.

When the offload process is triggered, the offload may be performed by the logger address space on a different MVS image. If the offload requires the movement of data to the log (offload) datasets, and the current dataset fills, a new dataset will be allocated. Unless the shareoptions are specified 3,3 the logger address space on the MVS image where CICS is running may not be able to access the dataset.

Another important point. If you have a DASDONLY user journal which is shared between multiple CICS regions in the same image, you must ensure all the CICS regions sharing the journal always run on the same MVS image. If one CICS is moved to another image, only the first CICS region to connect to the logstream will be successful.

Define Structure -LOGR Policy

Sample job to define the CF structure to be used for logstreams.

```
//DEFSTR JOB CLASS=A,MSGCLASS=A
//POLICY EXEC PGM=IXCMIAPU
//SYSPRINT DD SYSOUT=A
//SYSIN DD *

DATA TYPE(LOGR) REPORT(YES)

DEFINE STRUCTURE(LOG_JG)          LOG_JG is the structure name
LOGSNUM(10)   *                   up to 10 logstreams can connect
AVGBUFSIZE(400)  *                 size of the starting 'average' buffer
                                    - monitor the 'effective average buffer'
                                      using IXCMIAPU
MAXBUFSIZE(64000)  *              determines the size of each log buffer
                                   - also defines the CF element size
                                     > 65276   - element size is 512,
                                     = <65276 - element size is 256
```

* parms which affect performance

Define Structure Notes

The size of a STRUCTURE is specified in the CFRM policy. Each STRUCTURE is divided into ENTRIES and ELEMENTs. Each write uses 1 ENTRY and 1 or more ELEMENTs based on the length of data written.

Each STRUCTURE is divided EQUALLY between the connected log streams. When another log stream connects to the structure, the space is dynamically redistributed.

MAXBUFSIZE is returned to CICS at connect time, this is the value used for internal log buffers.
For user journals where the application does not use the wait option, it may be advantageous to specify a smaller size as the buffer is flushed when filled.

MAXBUFSIZE in conjunction with AVGBUFSIZE is used to determine the CF structure ENTRY/ELEMENT ratio. When data is written to the CF, it's written in increments equal to ELEMENT size. A MAXBUFSIZE greater than 65276 gives an element size of 512, a MAXBUFSIZE equal to or less than 65276 results in an element size of 256. For example: MAXBUFSIZE(65532) AVGBUFSIZE(1100)

CF element size is therefore 512 bytes, 3 elements are used for an average write. The Entry/Element ratio is 1:3 . Beginning with OS/390 R1.3, the Entry/Element ratio is dynamically adjusted.

When ENTRY space becomes 90% full, all logstreams in the structure are offloaded to DASD.

Monitor the 'effective average buffer' using IXCMIAPU

```
LIST STRUCTURE NAME(LOG_JG) DETAIL(YES)
     STRUCTURE NAME(LOG_JG) LOGSNUM(10)
     MAXBUFSIZE(64000) AVGBUFSIZE(400)
     EFFECTIVE AVERAGE BUFFER SIZE(768)
```

Monitor via RMF Post Processor.

Define Structure - CFRM policy

Sample job to define the CF structure to be used for logstreams.

```
//DEFSTR JOB CLASS=A,MSGCLASS=A
//POLICY EXEC PGM=IXCMIAPU
//SYSPRINT DD SYSOUT=A
//SYSIN DD *

DATA TYPE(CFRM) REPORT(YES)

DEFINE STRUCTURE(LOG_JG)

 INITSIZE(20000)  *
 SIZE(35000)  *

 PREFLIST(CF01,CF02)

  REBUILDPERCENT(1)
```

LOG_JG is the structure name

INITIAL CF size

if larger than INITSIZE, allows for a rebuild up to the SIZE value

specifies coupling facilities preference selection order for this structure

specify this value low so structures are rebuilt in the event of connectivity failure

- a value of 1 indicates when 1% of the systems lose connectivity to a structure, MVS will initiate a rebuild

Duplex data

- Duplexing
  - Each time a unit of work writes a log block to a log stream, system logger automatically makes a duplicate copy of the data.
    - the IXGWRITE calls are issued under the QR TCB, hence the CPU time is charged to the transaction
- Interim Storage
  - where data can be accessed quickly - w/o I/O to long term DASD
  - for CF logstreams - the list structure in the coupling facility
  - for DASDONLY - the local buffers in the logger dataspace
- Log or Offload datasets
  - data offloaded from interim storage at HIGHOFFLOAD
  - hardened for longer term access
  - HSM can archive
- Staging Dataset
  - for DASDONLY logstreams, contains the duplexed copy
  - for CF logstreams, used for duplexing if the CF becomes volatile or is failure dependent e.g. the CF LPAR and the MVS LPAR are in the same physical box
    - STG_DUPLEX(YES) DUPLEXMODE=(COND)

HSM Considerations

As noted under Log or Offload datasets, Hierarchical Storage Manager (HSM) may be used to manage these datasets. However, caution should be used when allowing HSM to manage any logger datasets. Staging datasets and the offload datasets for DFHLOG and DFHSHUNT should NOT be under the control of HSM.

If SETSYS INTERVALMIGRATION has been specified in the ARCCMDxx, HSM will on a regular interval (default is 90 minutes) examine the managed volumes to see if the allocated space is over the high threshold. If the high threshold has been reached it will begin the migration of datasets to reach the low threshold value.

When HSM is about to migrate (or restore) a dataset it holds the serialization lock (ENQ) for the file. If the logger attempts to allocate or connect to the file , Media manager (a component of DFSMS) will return an error. The error as presented does not indicate a temporary condition so it is treated as though the file is inaccessible and a gap-type error or an offload failure is returned to CICS based upon which operation is in progress.
Examples would be return code 8, reason code 84A, or return code 4 with reason code 403.

CICS treats this as a log integrity error and terminates with message DFHLG0772.

If this type of error is being encountered, examine the CICS joblog and MVS console log for messages issued by the underlying infrastructure. For example it is common to see message IEC161I 052(015)-084,IEESYSAS,IXGLOGR,.....
IEC161I 052(015)-084 for a serialization type problem.

- 052 indicates an open failure
- (015) sub function code - indicates a problem opening the dataset with share options (2,x) dataset for output, an exclusive ENQ request has been rejected - indicating another ACB has the dataset open for output
- 084 indicates a problem sharing the dataset

Notice the jobname is IXGLOGR.

HSM Considerations ...

Another common example is IEC161I 052(009)-084,IEESYSAS,IXGLOGR,.....,
(009) sub function code indicates the open failed for a shareoption (1,x) dataset for output, an exclusive ENQ request has been rejected - indicating another ACB has the dataset open

Shareoptions for staging and log datasets must be set to 3,3 .

Either an IDCAMS LISTCAT or ISMF displays may be used to verify the SHAREOPTIONS are set to 3,3.

An interesting side note here is the first offload dataset is allocated as part of logstream definition, but it is not actually opened until needed for an offload or possibly during the heartbeat health check. Based on logging activity and how well interim storage is tuned, it might be several hours or longer, but the offload dataset is needed.

Interactive Storage Management Facility (ISMF) can be used to display the HSM variables by dataset and group.
Refer to examples in the appendix.

Reference material :

OS/390 MVS Setting up a Sysplex GC28-1779
- 9.4.5.7 Set Up the SMS Environment for DASD Data Sets
- 9.4.9.3 Specifying SMS Data Set Characteristics for DASD Data Sets

DFSMS/MVS DFSMSdfp Storage Administration for SMS class and group considerations:

1. Storage groups identify volumes to be used for data set allocation.
2. Storage classes describe the performance characteristics for the data sets.
3. Data classes specify the data set characteristics and size of the data sets.
4. Management classes specify the migration characteristics of the data sets.
5. Automatic class selection (ACS) routines to assign the storage groups, classes, data classes and management classes, based on data set name when the DASD data sets are allocated.

HSM Considerations ....

Cross-region SHAREOPTIONS 1: The data set can be shared by any number of users for read processing, or the data set can be accessed by only one user for read and write processing. With this option, VSAM ensures complete data integrity for the data set. This setting does not allow any type of non-RLS access when the data set is already open for RLS processing.

Cross-region SHAREOPTIONS 2: If the data set has not already been opened for record-level sharing (RLS) processing, the data set can be accessed by any number of non-RLS users for read processing and it can also be accessed by one non-RLS user for write processing. With this option, VSAM ensures write integrity by obtaining exclusive control for a control interval when it is to be updated.

Cross-region SHAREOPTIONS 3: The data set can be fully shared by any number of users. With this option, each user is responsible for maintaining both read and write integrity for the data the program accesses. This setting does not allow any type of non-RLS access when the data set is already open for RLS processing. This option requires that the user's program use ENQ/DEQ to maintain data integrity while sharing the data set, including the OPEN and CLOSE processing.
User programs that ignore the write integrity guidelines can cause VSAM program checks, lost or inaccessible records, uncorrectable data set failures, and other unpredictable results. This option places responsibility on each user sharing the data set.

Cross-region SHAREOPTIONS 4: The data set can be fully shared by any number of users, and buffers used for direct processing are refreshed for each request. This setting does not allow any type of non-RLS access when the data set is already open for RLS processing. With this option, as in SHAREOPTIONS 3, each user is responsible for maintaining both read and write integrity for the data the program accesses. Refer to the description of SHAREOPTIONS 3 for ENQ/DEQ and warning information that applies equally to SHAREOPTIONS 4.

Cross-system SHAREOPTION 3: The data set can be fully shared. With this option, the access method uses the control block update facility (CBUF) to maintain integrity. As in cross region SHAREOPTIONS 3, each user is responsible for maintaining both read and write integrity for the data accesses by the program. User programs that ignore write integrity guidelines can cause VSAM program checks, uncorrectable data set failures, and other unpredictable results. This option places heavy responsibility on each user sharing the data set. The RESERVE and DEQ macros are required with this option to maintain data set integrity.

Cross-system SHAREOPTION 4: The data set can be fully shared, buffers used for direct processing are refreshed for each request.

LSN, data set names

- **Log Data Set name**
  - **often referred to as the offload data sets**
  - **HLQ.LSN.sequence#**

```
GRAUEL.IYOT1.DFHLOG.A0000003
```

- **Staging data set name -- used with coupling facility**
  - **HLQ.LSN.system_name**
    - ➤ system_name may be found in IEASYSxx, IEASYMxx or Loadxx parmlib members

```
GRAUEL.IYOT1.DFHLOG.MV55
```

- **Staging data set name -- used with DASD logging**
  - **HLQ.LSN.sysplex_name**
    - ➤ sysplex_name is specified as SYSPLEX() in the COUPLExx parmlib member

```
GRAUEL.IYOT1.DFHLOG.PLEXB
```

- **USING ISPF**

```
                DSLIST - Data Sets Matching GRAUEL.IYOT1.DFHLOG

               Command - Enter "/" to select action          Message           Volume
               ------------------------------------------------------------------------
OFFLOAD dataset  GRAUEL.IYOT1.DFHLOG.A0000003                                  *VSAM*
                 GRAUEL.IYOT1.DFHLOG.A0000003.DATA                             PBDA02
STAGING dataset  GRAUEL.IYOT1.DFHLOG.PLEXB                                     *VSAM*
                 GRAUEL.IYOT1.DFHLOG.PLEXB.DATA                                PBDA06
```

LSN = LogStream Name
HLQ = High level qualifier - specified on logstream definition

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

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
  - provides an easy to use interface to calculate the structure sizes based on minimum input
  - information from the CICS statistics is used as input

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 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)**
- **BYT Written to DASD (invoked) (SMF88LDB)**
- **# 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 entries 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 the HIGHOFFLOAD percentage reached in the staging dataset
  - ST Full (SMF88ETF)
    - number of times staging dataset full
  - Rebuild (SMF88ER)
    - number of structure rebuilds

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. 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 is more costly than synchronous requests.

SMF 88 notes

Using the SMF 88 report produced via IXGRPT1:

NOTE: OW36423 should be applied, - it increases several counters 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 BY USERS IXGWRITES". 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 message DFHLG0743 message being issued for DFHSHUNT indicates units of work, which had been in active, have completed. APARS PQ22563 and PQ14796 reduce data recorded on DFHSHUNT
   - 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.
     > DASD-only, the allocation size of the staging dataset may be too small for CF logstreams
     > incorrect entry : element ratio can happen when unlike logstreams are in the same structure, the ratio is based on the worst (largest) buffsize, dynamic changes will happen on no less than 30 minute intervals
   - 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 an offload is in progress
   TYPE3 -- writes issued when the number of entries in use for a logstream has reached 90%
   - 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.
   - 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	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	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	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	0
IYOT3.DFHLOG	LOG_JG	4340339	6758144	2167961	11238	10788	428	22	386
		2040630	6883	2011881	3902	30	2 0	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	0
IYOT3.DFHLOG	LOG_JG	116	256	0	1	1	0	0	116
		0	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) 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.

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 entries for the structure 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) 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.

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.

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.

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.

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, reference 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.

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'.
```

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	REASON	DELAYED REQUESTS # 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 requests 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), round the answer 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 std 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 us. 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

                                      % OF                    % OF      AVG    LST/DIR  DATA      LOCK      DIR REC/
        STRUCTURE              ALLOC  CF                 #     ALL       REQ/   ENTRIES  ELEMENTS  ENTRIES   DIR REC
TYPE NAME   STATUS    CHG     SIZE   STORAGE      REQ      REQ       SEC    TOT/CUR  TOT/CUR   TOT/CUR   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
          # REQ    -------------- REQUESTS -------------    -------------- DELAYED REQUESTS -------------
SYSTEM    TOTAL                 #     % OF  -SERV TIME(MIC)-    REASON     #    % OF  ---- AVG TIME(MIC) -----
NAME      AVG/SEC               REQ    ALL     AVG    STD_DEV             REQ    REQ    /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 requests 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 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.

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.
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 in topic 9.3.1.**
    - **Finding Information for OPERLOG Log Stream in topic 9.3.2.**
    - **Finding Information for Logrec Log Stream in topic 9.3.3**

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

- *CICS Transaction Server for OS/390 Version 1 Release 2 Implementation Guide -(Redbook) - SC24-2234*

- *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) / number log writes = 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 is 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 logstreams interim storage (i.e. in the coupling facility structure or on the staging dataset for a DASDONLY logstream). The larger the interim storage the more data must be contained in the dataspace.

The logger address space is one of the highest priority in the system, so it's 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 processors 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 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 3 -        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 3 (Data Class list), the SHAREOPTIONs for all datasets in the class are shown in columns 21 and 22.


OPTION 4 -        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 4 (Management Class list), column 9 indicates if the dataset may be migrated either automatically or via
                                  a manual command.
```