



SHARE SESSION REPORT

ACKNOWLEDGMENT

The author thanks David Chess and other Yorktown professionals for providing many CAP concepts taken from the VM Prioritizer. Also, the design and implementation efforts of Dave Anderson, Bob Greenwalt, and Don Wagler are appreciated.

SHARE NO.	SESSION NO.	SESSION TITLE	ATTENDANCE
61	M504	SMF and RMF Differences In XA	350
Computer Management & Eval.		Wes Rushton	DCL
PROJECT		SESSION CHAIRMAN	INST. CODE
Datacrown Inc., 650 McNicoll Ave., Willowdale, Ont. 416-499-1012			
SESSION CHAIRMAN'S COMPANY, ADDRESS, AND PHONE NUMBER			

Dr. H.W. Barry Merrill, a recognized leader in the field of SMF and RMF Data Analysis, gave a talk on the differences in SMF and RMF records between MVS/370 systems and MVS/XA. The data presented will be of great interest to those computer performance analysts and capacity planners who make regular use of such data.

A copy of the text and foils used in Dr. Merrill's presentation is attached.

491

SMF AND RMF DIFFERENCES IN XA

H. W. BARRY MERRILL, PhD
 Sun Company
 Information Systems Division
 Box 47803
 Dallas, TX 75247
 214-934-6816

SHARE 61
 Sessions B185/M504
 Wednesday, August 24, 1983
 10:15 a.m., Hilton, West Ballroom

ABSTRACT

This session will present a preliminary discussion of how data used for performance analysis and cost accounting has changed in the MVS/XA implementation of RMF and SMF. Record types and data elements will be addressed in some detail. The new channel I/O measurements (device connect time) will receive preliminary analysis.

MVS/XA has made major changes in SMF and RMF records and their contents. The RMF records are completely re-formatted, the Type 30 records are expanded with new data, several new SMF records exist, and there remains the perennial SMF accounting question: What will happen to CPU timing and EXCP counting and my billing system when we migrate?

The news is that the TCB and SRB CPU timings recorded in the Step records (Type 4 and Type30, subtype 4) appears to have minor increases from MVS/370 to MVS/XA. Preliminary benchmarks with 26 unique steps executed repeatedly as initiators were increased from 5 to 60 in five minute increments showed the following results (with reference to MVS/370):

Metric	Average Ratio	Minimum Ratio	Maximum Ratio
Minimum TCB+SRB	1.03	.88	1.24
Minimum TCB alone	.99	.86	1.18
Average TCB+SRB	1.02	.85	1.24
Average TCB alone	.98	.83	1.17

In addition, the repeatability of these measures within MVS/370 and MVS/XA were measured and compared:

Metric	Repeatability as Percent of Mean For MVS/370			Repeatability as Percent of Mean For MVS/XA		
	Min	Avg	Max	Min	Avg	Max
TCB Time	0	4.5	15	0	3.5	11
TCB+SRB Time	0	5.3	23	0	3.8	12
EXCP Total	0	.5	10	0	0.2	2
Connect Time Total	.	.	.	0	7.4	34

The bad news isn't really that bad, but EXCP counts are increased in MVS/XA. The primary cause for the increase is the re-write of PROGRAM FETCH (the portion of MVS which loads programs into memory). Under MVS/370, FETCH used a combination of EXCP and PCI to read a module from STEPLIB, JOBLIB or LINKLIST into memory, and under MVS/370 PCI were not counted. In MVS/XA, FETCH has been completely re-written and uses only EXCP. Thus the STEPLIB DD segment in the Step Termination records will now contain the true number of text records which were loaded from disk to real memory by PROGRAM FETCH. The count of text records almost always will equal the number of EXCPs which were issued, but if a text record has too many RLD segments, it can take two EXCPs for the text record. For all intents, however, the STEPLIB DD can now be said to contain the true EXCP count to load the module. MVS/XA frequently records 10 EXCP where MVS/370 only counted one STEPLIB EXCP.

There is a much more significant impact on your installation in this re-write of PROGRAM FETCH than the change in STEPLIB EXCPs. Especially for IMS, TSO and CICS environments, and all cases where dynamic LINKS and LOADS are frequent, it will be very prudent to re-block and correct the text record counts in all of your load libraries (or re-Link edit them) to achieve performance equivalent to MVS/370. Without this re-blocking, XA can double device busy and IO rates to program libraries.

FETCH now uses the count of RLD segments for each text record to construct its EXCP. The directory contains the count of RLD segments that follow the first text record, and the last RLD segment contains the number of RLD records following the next text record. FETCH will build an EXCP to read the text record and as many RLD segments as are indicated by the count field. Unfortunately, the count field used by XA was not generally updated by the MVS/370 Linkage Editor, and for a while DFDS created wrong count values. When XA has read in a non-zero count value, and then finds the actual count of RLD segments is not the same, a console message (CSV300I) is created. (More typically, thousands of console messages are created) The message simply advises that an incorrect count of RLDs was encountered. It is nice to see that the IBM designer was watching out for the performance analyst. For instead of ignoring the error, IBM is letting us know of a possible performance degradation with this console message.

The console messages go away when the RLD counts are corrected by using IEBCOPY on the load library. The library can be corrected in place, with the ALTERMOD option to correct the counts, but a 75 cylinder library with 1449 members required 31 minutes elapsed time, 12 seconds TCB and 7 seconds SRB (on a 3081K).

This will still leave PROGRAM FETCH a poor performer, however, until the load library is re-blocked. With RLD counts corrected, and with the load library blocked, it appears PROGRAM FETCH gets back to the same level of IO activity that the previous 370 FETCH used. Load libraries can be blocked with IEBCOPY's COPYMOD options, but it can not be executed in place; one must copy off to tape and back, or from one disk to another. The same 75 cylinder load library running a COPYMOD used 37 minutes elapsed time with 31 seconds TCB and 13 seconds SRB. The load library could also be re-link-edited with the XA LKED (if you still have the source on those ole' modules) to re-block the library.

Even after all this effort, the CSV300I message can occur. There are still some members which IEBCOPY cannot correct. The IEBCOPY listing identifies the members which were not corrected because they either were marked non-Link-Editable (for example, if the ESD records have been stripped by SAS PROC PDSCOPY for performance), or were marked OVERLAY, or were in SCATTER load module format.

The major SMF and RMF changes which will be of interest to performance analysts and capacity planners are identified below.

NEW SMF RECORDS.

.Six additional new records are documented in the XA SMF Manual. The Type 16 record written by ICEMAN sort, a new VSAM Volume Data Set Updated record (Type 60), three new records describing ICF/VSAM Activity (Types 61, 65 and 66), and RMF now creates a Type 78 IO Queueing Activity record.

DATA WHICH NO LONGER EXISTS.

.Type 6 - SMF6IOE (JES3 Control Buffer and Data Buffer Errors) is now reserved.
 .Type 6 - SMF6DFE (JES2 bad record length truncated) is now reserved.
 .Type 71 - PVTSPREC (LPA swap pages reclaimed) was deleted.
 .Type 71 - SMF71LVL (Record Level -Total or Pageable, documented in TNL-GN28-4962) went away.

GENERAL INTEREST ITEMS.

.The Device Number replaces Unit Address in records type 4,8,9,10, 11,14,15,19,21,34,40,64,69,74, and 75.
 .VIO is now identified by Device Number of '7FFF'x (because the old value of '0FFF'x is now a valid device number).
 .MSS Virtual Devices are now identified by a 1 in the first bit of the first nybble of device number (formerly a 1 in the second nybble of Unit Address identified these devices).
 .There are many many typographical errors in the original XA SMF manual (-0). About half of these are corrected in (-1), but it will require a (-2) to clean up the rest.
 .No Device Connect time is recorded in the TYPE 14/15/64 records. While the DD segment for the step using a data set will contain this time, its absence from the actual data set record will effectively prevent analysis of data set usage by Device Connect Time.
 .RMF reports still truncate fractional values, rather than round.
 TYPE 0 SMF Record.
 .SMF0BUF contains x'4000' to indicate a 4096 byte SMF buffer size where one would have expected x'1000'.
 .SMF0VST contains x'FFE00000' for the Virtual Storage size of MVS/XA.

TYPE 22 SMF Record.

.The MSS device segments are documented as through device 0EF, but the length of the data indicates a maximum MSS address of 0DF.
.Real storage pages for MVS/XA appear to be zero.

TYPE 26 SMF Record (JES2 only).

.JES2 has added a new Print Section which contains (or will soon contain) the number of bytes printed/punched. (This segment exists in MVS/370 records also).

TYPE 30 SMF Record.

.The major change in the Type 30 record is the addition of Device Connect Time. Each DD segment in a type 30 is expanded to contain the IO connect time for that DD. In addition, the total IO connect time for the address space for that interval (step, job, or interval) is recorded in a separate field. There are corresponding EXCP counts at the DD Segment level, and the total EXCP count for the address space. By summing the data in the DD segments and comparing with the address space totals, we find three pairs of IO measures for EXCP and for Connect Time: EXCPTOTL, EXCPTODD, EXCPNODD and IOTMTOTL, IOTMTODD, and IOTMNODD. The TOTL counts and times are the address space total values. The TODD counts and times represent the IO which is recorded at the DD level. The NODD counts and times represent the IO not recorded at the DD level, which is primarily the SPOOL I/O activity in JES2, although the NODD data also includes IO to the LINKLIST data sets, Catalog Management IO, and OPEN/CLOSE IO beyond priming and purging the buffer. SPOOL IO activity should predominate the NODD data. Thus, while the JES3 shop is still in the cold, the NODD counts and times should be used in accounting and capacity planning analysis to recover the cost of spool IO activity. Appendix 8-2 of the SMF manual discusses what is and what is not captured in the NODD field and what is captured in the TOTL (address space) field. While that discussion only mentions EXCPs, it applies equally to Device Connect time.

.Under MVS/370, a DD segment existed for each real UCB address through which IO passed enroute to a device. With multiple channel paths to the same actual device, MVS/370 would create a separate segment and count the IOs via each physical path. In MVS/XA, since IO is to a Device Number, which is independent of physical path, these multiple DD segments in the Type 30 are eliminated. This will reduce the size of the Type 30 record, and will make installation banner pages much more readable, since there will be only one EXCP count per DD now.

.SMF30JNM (incorrectly documented in the -0 SMF manual) should contain the JES job identifier (JOBnnnn, STCnnnn, etc.) but for some STC's, it contains the Jobname of the STC.

.A new field, Abend Reason Code is created in MVS/XA.

.A flag is set (SMF30SFL) if IEFUSI was used to change the region size of the extended private area of the task.

.Six new fields exist which describe virtual storage allocations:

SMF30RCB - Region size below 16MB.
SMF30ERG - Region size above 16MB.
SMF30ARB - Maximum LSQA + SWA subpools below 16MB.
SMF30EAR - Maximum LSQA + SWA subpools above 16MB.
SMF30URB - Maximum user subpools below 16MB.
SMF30EUR - Maximum user subpools above 16MB.

TYPE 70 RMF RECORD.

.The version of the CPU has been added. This field identifies which model of 3081 or 3083 is in use. It is needed to determine which model dependent CPU constant was used in converting CPU seconds to service units. In MVS/370 only the CPUTYPE (3081 or 3083) was provided. The model dependent constants should be:

Processor	Version	Service Units per Second	"Speed" Relative to 3081D	"Capacity" Relative to 3081D
3081-D	'03'x	276.3	1.00	1.00
3081-G	'13'x	291.65	1.05	1.05
3081-K	'23'x	368.4	1.33	1.33
3083-E	'01'x	204.8	0.74	0.37
3083-B	'11'x	306.5	1.11	0.55
3083-J	'21'x	388.8	1.41	0.71

However, use these values with caution, since the numbers are changed by PTF. The values are in IRARMCPU or can be determined with the TSOTEST command (as documented in Candle Report Vol 4. No 8. April 18, 1982):

```
TEST 'SYS1.LINKLIB(IEFBR14)'  
L 10.%+25C%+40 F
```

The list command will display

```
HHHHHH NNNNNN
```

(where HHHHHH is the hex address of the constant NNNNNN)

Then,

```
SU_SEC = 16,000,000  
NNNNNN
```

.As discussed in Candle Report Vol 5 No 9, MVS/XA can handle I/O interrupts in "batches" out of the SLIH rather than to always absorb the cost of an interrupt per I/O. The Type 70 record now contains the count of SLIH I/O interrupts handled by each CPU, and the count of the number of I/O interrupts which were handled by the TPI (Test Pending Instruction), or were batched. Since MVS/XA disables processors for I/O interrupts if the percentage of I/O interrupts which are handled by TPI is sufficiently low, as measured by these counts, we may be concerned with the threshold value of TPI interrupts for I/O interrupts.

.SMF70FLA is documented as a field in the SMF manual, but in fact will never contain other than zero. The SMF manual was apparently created from the RMF report program DSECTS, which does document this internally used field]

.The bit map which formerly described the Selectable Units installed no longer exists. It was never reliable anyhow, as bits could be on or off independent of the SU they represented. (Several fine software vendors attempted to use this bit map to chose optimum I/O if SU9 was installed and received program ABENDS for their effort.)

TYPE 71 RMF RECORD.

.In MVS/370, if the processor had more than 16MB of real memory, the Type 71 record used three fields to report the number of pages moved above and below this line, as MVS/370 tried to optimize the movement constraints (not all pages could be placed above the 16MB line). In MVS/XA the constraints are lifted so that the three fields which individually identified page movement are replaced by a single count of pages moved. However, the designers took advantage of the 12 bytes already defined in the record, and they now contain the average, minimum and maximum number of frames below 16MB real which are fixed.

TYPE 72 RMF RECORD.

.I/O Service Units can now be captured based on EXCP count (the only case in MVS/370, and the default in MVS/XA), or the new Device Connect Time (DCTI) can be used as the basis for I/O Service Units. The choice is made by setting the IOSRVC parameter in IEAIP5xx. Unfortunately, there is no bit set in any SMF/RMF record which indicates whether I/O service units were based on EXCP or on Connect Time. This will create measurement problems until a flag is added to the record. (The IPS Name is in the Type 72, but using it would be shaky at best.)

.If Device connect time is used for I/O Service Units, the I/O Service Unit is defined as 65 Connect Time Units. One Connect Time Unit is 128 microseconds. Thus an I/O Service Unit based on connect time represents 8.32 milliseconds of connect time, which is just about one half revolution of a 3600 rpm disk drive] Thus if all your I/O were at half track blocking, I/O Service Units based on EXCP or Connect should be equivalent.

.A nit, but the last two bytes of the ICSNAME field contain a value of x'6060' when the ICSNAME is not specified.

.SMF72FLA and SMF72FLG are also RMF Report fields like SMF70FLA, described above.

TYPE 73 RMF RECORD.

.Since MVS/XA no longer has Logical Channels, the logical channel queuing data has disappeared from the Type 73. In fact, the Type 73 now contains only physical channel busy percentage; there is not even a count of STO (or now, SSCH) activity on the physical channel. While this data loss may affect some, it results from RMF saving overhead. In MVS/370, RMF acquired physical channel data by issuing Test Channel instructions and actually counted channel SIO. In MVS/XA RMF 'steals' the results of the STCPS instruction which is issued by the SRM to manage I/O loading. The SRM attempts to schedule STCPS every 200 milliseconds, but can be pre-empted from sampling. Thus the Type 73 record also contains the count of STCPS instructions which were 'stolen' and used in the RMF calculation of physical channel busy. This is the first case wherein the sampling rate of RMF data is not set by RMF.

TYPE 74 RMF Record.

.Ensure UZ58785 is installed, its pre- and co-recs, and an associated microcode level. Without these fixes, the device segment in each Type 74 record is eight bytes short, and RMF reports may show negative connect times. This cluster of fixes are referred to as the "Channel Measurement Overflow PTF).

.The four connect time measurements, Connected, Active, Disconnected, and Pending are separately measured for each device. No longer is the RMF Device busy measurement made by sampling UCB bits at the RMF sampling interval. Rather, the percent of connect and disconnect time is added to the percent of time the device was blocked from another system to measure the device 'utilization'. This definition of utilization is good in that it reflects when the device was unavailable to another IO request.

TYPE 78 RMF Record.

.In all prior RMF records, the offsets were based on the beginning of the logical record. In this record, however, R781CPDS is an offset from the start of the IO Queuing Control Section. Additionally, R781CPDN, the count of configuration data sections contains 4, even when there are only 2 sections with data (4 sections exist, the last two are all hex zero).

.The logical queuing statistics, which were maintained in the Logical channel segments of the MVS/370 record, are not recorded at the device level in the Type 78 record. New fields report the percent of IOs which were delayed because All Paths were busy, because Control Unit was busy, or because the Device was busy, and the percent of successful IOs to the device can also be calculated from the data.

SUMMARY.

All in all, the performance analyst and capacity planner will be very pleased with the improvements made in SMF and RMF data with MVS/XA. As with any major architectural change, there will be minor errors to be corrected and new measures to be understood, but the data added by MVS/XA can only serve to provide better measures of resources and service.

SMF AND RMF DIFFERENCES IN XA

H. W. BARRY MERRILL, PhD
SUN COMPANY
INFORMATION SYSTEMS DIVISION
BOX 47803
DALLAS, TX 75247
214-934-6816

SHARE 61
SESSIONS B185/M504
WEDNESDAY, AUGUST 24, 1983
10:15 A.M., HILTON, WEST BALLROOM

MVS/XA CPU TIME
 DIVIDED BY
 MVS/370 CPU TIME

<u>METRIC</u>	<u>AVERAGE RATIO</u>	<u>MINIMUM RATIO</u>	<u>MAXIMUM RATIO</u>
MINIMUM TCB+SRB	1.03	.88	1.24
MINIMUM TCB ALONE	.99	.86	1.18
AVERAGE TCB+SRB	1.02	.85	1.24
AVERAGE TCB ALONE	.98	.83	1.17

REPEATABILITY OF CPU TIME

<u>METRIC</u>	<u>REPEATABILITY AS PERCENT OF MEAN FOR MVS/370</u>			<u>REPEATABILITY AS PERCENT OF MEAN FOR MVS/XA</u>		
	<u>MIN</u>	<u>AVG</u>	<u>MAX</u>	<u>MIN</u>	<u>AVG</u>	<u>MAX</u>
TCB TIME	0	4.5	15	0	3.5	11
TCB+SRB TIME	0	5.3	23	0	3.8	12
EXCP TOTAL	0	.5	10	0	0.2	2
CONNECT TIME TOTAL	.	.	.	0	7.4	34

437

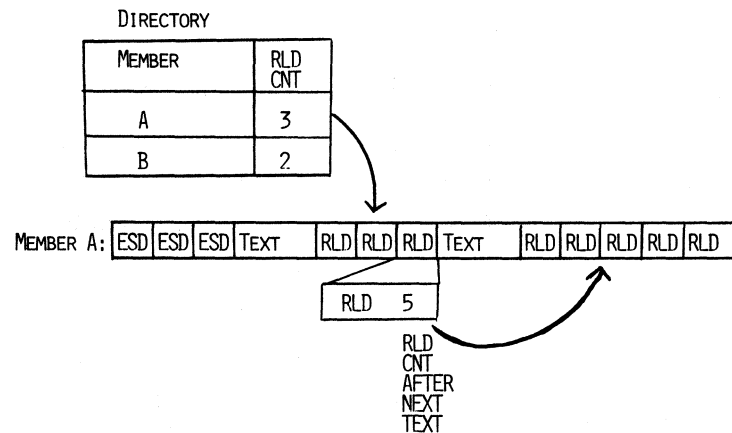
STEPLIB EXCP COUNT INCREASES

- MVS/370 PROGRAM FETCH
 - USES EXCP + PCI TO LOAD MODULES.
 - ONLY EXCPs COUNTED IN STEPLIB
 - TYPICALLY COUNTED 1-2 EXCPs TO LOAD PROGRAM
- MVS/XA PROGRAM FETCH
 - REWRITTEN, USES ONLY EXCP
 - STEPLIB COUNT IS COUNT OF TEXT RECORDS
 - USUALLY 1 EXCP PER TEXT RECORD
 - MODULE WITH TOO MANY RLDS WILL TAKE MORE THAN 1 EXCP PER TEXT RECORD
 - TYPICALLY 10-12 "EXCPs" COUNTED NOW

IMPACT OF REWRITE OF PROGRAM FETCH

- MUST REBLOCK LOAD LIBRARIES
- ESPECIALLY CRITICAL FOR TSO/IMS/CICS LIBRARIES

498



- READS RLD COUNT "N" IN DIRECTORY
- BUILDS EXCP FOR FIRST TEXT AND "N" RLDs
- GETS COUNT OF RLDs FOLLOWING NEXT TEXT IN LAST RLD
- BUILDS EXCP FOR SECOND TEXT AND ITS RLDs
- IF RLD COUNT IS NON-ZERO AND WRONG
 - CONSOLE MESSAGE CSV300I WRITTEN
 - ONE EXCP PER RLD PLUS ONE FOR TEXT

FOIL 6

WRONG RLD COUNT

- WHY COUNTS WRONG
 - DFDS 1.4 CREATED WRONG VALUES
 - OLD MODULES/LKED MAY HAVE NOT COUNTED CORRECTLY
- FIX COUNT
 - IEBCOPY ALTERMOD WILL CORRECT RLD COUNTS
 - CAN RUN IN PLACE
 - 75 CYL 1449 MEMBERS ALTERMOD REQUIRED
 - 31 MINUTES ELAPSED
 - 19 CPU SECONDS (12 TCB 7 SRB)
 - 3081K

FOIL 7

500

- FETCH STILL POOR UNTIL LIBRARY IS REBLOCKED

IBM NUMBERS:

IBM ANALYSIS (FROM G622-9325-00)

LIBRARY TYPE	MEASURE	MVS/370	MVS/XA NO CHARGE	MVS/XA ALTERMOD (RLD COUNTS FIXED)	MVS/XA COPYMOD (RLD COUNTS PLUS REBLOCK)
PROGRAM	I/O RATE:	8,6	14,9	13,9	8,9
PRODUCT	DEVICE				
LIBRARY	BUSY %	21,9	42,4	36,7	23,8
SYS1.LINKLIB	I/O RATE:	10,0	15,3	14,5	8,8
AND	DEVICE				
SYS1.CMDLIB	BUSY %	29,7	36,6	34,8	24,4

- IEBCOPY COPYMOD WILL CORRECT RLD AND REBLOCK
- CANNOT RUN IN PLACE - MUST DISK-TO-TAPE-TO-DISK
- SAME LIBRARY COPYMOD REQUIRED

37 MINUTES ELAPSED
 44 CPU SECONDS (31 TCB 13 SRB)
 3081K

- NEW SMF/RMF RECORDS
 - TYPE 16 - ICEMAN
 - TYPE 60 - VSAM VOLUME DATA SET UPDATED
 - TYPE 61 - INTEGRATED CATALOG FACILITY DEFINE
 - TYPE 65 - ICF DELETE
 - TYPE 66 - ICF ALTER
 - TYPE 78 - I/O QUEUING AND CONFIGURATION
- OLD FIELDS NOW RESERVED
 - TYPE 6 - SMF6IOE (JES3 BUFFER ERRORS)
 - - SMF6DFE (JES2 BAD RECORD LENGTH)
 - TYPE 71 - PVTSPREC (LPA SWAP PAGES RECLAIMED)
 - - SMF71LVL (TOTAL OR PAGEABLE RECORD FLAG)

GENERAL CHANGES

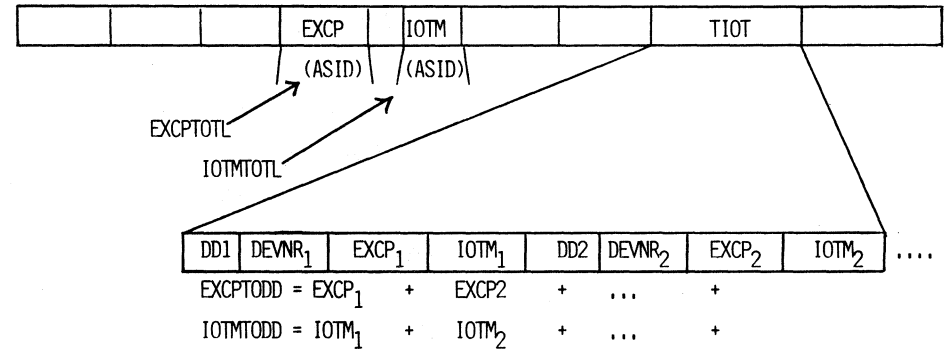
- DEVICE NUMBER REPLACES UNIT ADDRESS
 - RECORDS

4	8	9	10	11	14	15
19	21	34	40	64	69	74 75
 - CAN HAVE 4080 DEVICES (0000 THROUGH 00FF)
 - UNIQUE DEVICE NUMBER FOR EACH DEVICE
 - UNIT ADDRESS WAS CHANNEL + UNIT
 - I/O-RELATED CONTROL BLOCKS KEPT LAST USED CHANNEL
 - I/O TO 181/381 UNIT ADDR HAD BOTH ADDRESSES
 - REQUIRED TWO-CHANNEL RELATIONSHIP IN SMF SOFTWARE
 - REDUCES SIZE OF TIOT AND RECORDS CONTAINING TIOT
- VIO IDENTIFIED BY DEVICE NUMBER '7FFF'X
 - OLD VALUE WAS 'OFFF'X
 - OLD VALUE IS NOW VALID DEVICE NUMBER
- MASS STORAGE VIRTUAL VOLUMES IDENTIFIED
 - WAS 1ST BIT OF SECOND NIBBLE OF UNIT ADDRESS
 - NOW 1ST BIT OF FIRST NIBBLE OF DEVICE NUMBER

501

- TYPE 30 CHANGES - STEP/JOB TERMINATION/INTERVAL DATA

- DEVICE CONNECT TIME - IOTM



- THEN
- EXCPNODD = EXCPTOTL - EXCPTODD
- IOTMNODD = IOTMTOTL - IOTMTODD
- THE NODD MEASURES COUNT/RECORD ACTIVITY WHICH IS NOT CAPTURED AT THE DD LEVEL:
 - LIBRARY SEARCHES AND FETCHES FROM LINKLIST DATA SETS
 - I/O TO THE JES2 SPOOL FROM THE ASID
 - CATALOG MANAGEMENT I/O
 - OPEN/CLOSE I/O BEYOND PRIMING AND PURGING THE BUFFER

- MORE TYPE 30 CHANGES

- FEWER DD SEGMENTS IN TYPE 30 (AND TYPE 4, 40)
 - ONLY ONE DEVICE NUMBER PER DEVICE
 - MVS/370 HAS ONE DD SEGMENT FOR EACH PHYSICAL ADDRESS USED (MULTI-CHANNEL SWITCHES)
- ABEND REASON CODE IS A NEW FIELD
- FLAG SET IF IEFUSI CHANGED REGION SIZE OF TASK
- SIX NEW FIELDS DESCRIBE VIRTUAL STORAGE SIZES:
 - SMF30RGB - REGION SIZE BELOW 16MB
 - SMF30ERG - REGION SIZE ABOVE 16MB
 - SMF30ARB - MAXIMUM LSQA + SWA SUBPOOLS BELOW 16MB
 - SMF30EAR - MAXIMUM LSQA + SWA SUBPOOLS ABOVE 16MB
 - SMF30URB - MAXIMUM USER SUBPOOLS BELOW 16MB
 - SMF30EUR - MAXIMUM USER SUBPOOLS ABOVE 16MB

502

TYPE 70 CHANGES: RMF CPU ACTIVITY

- CPU VERSION HAS BEEN ADDED

PROCESSOR	VERSION	SERVICE UNITS PER SECOND	"SPEED" RELATIVE TO 3081D	"CAPACITY" RELATIVE TO 3081D
3081-D	'03'x	276.3	1.00	1.00
3081-G	'13'x	291.65	1.05	1.05
3081-K	'23'x	368.4	1.33	1.33
3083-E	'01'x	204.8	0.74	0.37
3083-B	'11'x	306.5	1.11	0.55
3083-J	'21'x	388.8	1.41	0.71

- VALUE OF SERVICE UNITS PER SECOND (SU_SEC) IS ACTUALLY SET BY NIP
- THE CONSTANT IS CHANGED BY PTFs
- TO ACTUALLY DETERMINE THE SU_SEC VALUE IN EFFECT, UNDER TSO

```
TEST 'SYS1.LINKLIB(IEFBR14)'
L 10.%+25C%+40 F
```

THE LIST COMMAND WILL DISPLAY

```
HHHHH NNNNN
```

(WHERE HHHHH IS THE HEX ADDRESS OF THE CONSTANT NNNNN)

THEN,

$$SU_SEC = \frac{16,000,000}{NNNNN}$$

MORE TYPE 70 CHANGES:

- MVS/XA CAN HANDLE IO INTERRUPTS IN "BATCHES" VIA TEST PENDING INSTRUCTION (TPI)
- THIS AVOIDS COST OF ONE INTERRUPT (ONE FULL SLIH ENTRY PER IO)
- MVS/XA WILL DISABLE PROCESSOR(S) FOR IO:
 - WHEN THE PERCENTAGE OF IO INTERRUPTS "BATCHED" IS SUFFICIENTLY LOW
 - PROVIDES MORE THRUPT FOR PROCESSORS WHICH ARE DISABLED FOR IO INTERRUPTS
 - DISABLING IS DYNAMIC WITH LOAD
- TYPE 70 RECORDS, FOR EACH PROCESSOR
 - IO INTERRUPTS HANDLED BY SLIH
 - IO INTERRUPTS HANDLED BY TPI

TYPE 71 CHANGES - RMF PAGING AND MEMORY

FOR PROCESSORS WITH MORE THAN 16MB REAL MEMORY

- MVS/370 TRACKED PAGE MOVEMENT IN THREE FIELDS
 - PAGES MOVED BELOW 16MB REAL (BAD FOR PERFORMANCE)
 - PAGES MOVED UP, NOT LIKELY TO BE MOVED DOWN
 - PAGES MOVED UP, LIKELY TO BE MOVED DOWN AGAIN
- MVS/XA HAS EASED THE CONSTRAINTS ON PHYSICAL PAGE LOCATION
- TYPE 71 NOW TRACKS ONLY ONE FIELD
 - TOTAL PAGES MOVED
- THESE FIELDS ARE INCORRECTLY DOCUMENTED IN -O SMF MANUAL

603

TYPE 72 CHANGES - RMF WORK ACTIVITY

- ° IO SERVICE UNITS
 - ° MVS/370 - BASED ONLY ON EXCP
 - ° MVS/XA - CAN USE EITHER EXCP OR IOTM (DCTI)
 - ° CHOICE SPECIFIED BY IOSRVC PARM IN IPS
 - ° UNFORTUNATELY, NO EXPLICIT FLAG TO TELL WHICH
 - ° MIGHT USE IPSNAME, BUT NOT SAFE NOR GENERALIZED
- ° IO SERVICE UNIT CALCULATION
 - ° EXCP - ONE "SU" PER EXCP
 - ° IOTM - ONE "SU" PER 65 CONNECT TIME UNITS
 - ° ONE CONNECT TIME UNIT IS 128 MICROSECONDS
 - ° $65 * 128 \text{ MICROSECS} = 8,32 \text{ MILLISECONDS}$
 - ° $8,32 \text{ MILLISECONDS} = \text{ONE HALF REVOLUTION (3600 RPM)}$
 - ° THUS, IOTM "SU" EQUALS EXCP "SU" IF HALF TRACK IO
 - ° ACTUAL IO SERVICE UNITS IS
IOC COEFFICIENT * "SU"
- ° NITS
ICSNAME FIELD CONTAINS X'6060' AS LAST TWO BYTES WHEN NO ICS IS SPECIFIED

TYPE 73 CHANGES - RMF CHANNEL ACTIVITY

- ° NO MORE LOGICAL CHANNELS IN MVS/XA
 - ° IO IS TO DEVICE
- ° HENCE, NO MORE LOGICAL QUEUING DATA IN TYPE 73
 - ° QUEUING NOW IN NEW TYPE 78
- ° NO MORE SIO COUNT IN PHYSICAL CHANNEL DATA
 - ° MVS/370 RMF ACTUALLY COUNTED SIO AND ISSUED TEST CHANNEL
 - ° MVS/XA "STEALS" RESULTS OF SRM'S STCPS INSTRUCTION TO SAMPLE CHANNEL BUSY
 - ° STCPS ISSUED EVERY 200 MILLISECONDS BY SRM
 - ° CAN BE PRE-EMPTED
 - ° COUNT OF STCPS PROVIDED IN TYPE 73
 - ° ONLY CASE WHERE RMF SAMPLE RATE (CYCLE) IS NOT USED FOR SAMPLED DATA

TYPE 74 CHANGES - RMF DEVICE ACTIVITY

- ENSURE UZ58785 INSTALLED
 - Co-RECS AND PRE-RECS, TOO
 - ASSOCIATED MICROCODE LEVEL REQUIRED
 - REFERRED TO AS "CHANNEL MEASUREMENT OVERFLOW" PTF
- IMPACT WITHOUT UZ58785
 - TYPE 74 DEVICE SEGMENT 8 BYTES TOO SHORT
 - RMF REPORTS MAY SHOW NEGATIVE CONNECT TIME
- FOUR NEW MEASURES OF IO ACTIVITY SEPARATELY MEASURED
 - CONNECTED
 - DISCONNECTED
 - ACTIVE
 - PENDING
- DEVICE BUSY NO LONGER JUST SAMPLED FROM UCB BITS
 - BUSY = CONNECT + DISCONNECT + BLOCKED BY OTHER SYSTEM
- REFLECTS TIME DEVICE WAS UNAVAILABLE TO ANOTHER IO REQUEST

TYPE 78 - IO CONFIGURATION AND QUEUING

- IDENTIFIES CONTROL UNIT AND DEVICE CONNECTIVITY IN CONFIG SECTION
- CONTAINS LOGICAL QUEUING STATISTICS
 - PERCENT OF IOs WHICH WERE SUCCESSFUL
 - FOR IOs WHICH WERE DELAYED:
 - PERCENT DELAYED BECAUSE ALL PATHS BUSY
 - PERCENT DELAYED BECAUSE CONTROL UNIT BUSY
 - PERCENT DELAYED BECAUSE DEVICE BUSY
- OFFSET (R781CPDS) IS NOT FROM START OF RECORD (AS ALL OTHER SMF/RMF OFFSETS), BUT IS THE OFFSET FROM BEGINNING OF IO QUEUING CONTROL SECTION
- COUNT OF CONFIGURATION DATA SECTIONS CONTAINS VALUE OF 4, EVEN WHEN THERE ARE ONLY 2 SECTIONS WITH DATA. LAST 2 SECTIONS CONTAIN HEX ZEROS.

SUMMARY

- ° IO CONNECT TIME WILL BE BETTER MEASURE OF IO
- ° CPU MEASURES SEEM MORE STABLE
- ° REBLOCK LOAD LIBRARIES
- ° PERFORMANCE IMPROVEMENT BEYOND VSCR PROBABLY NOT IN XA UNTIL ALL DEVICES ARE 3380S (TO TAKE ADVANTAGE OF DPS)

° SERIOUS PROBLEMS

- ° DESIGN OMISSION
 - ° DEVICE CONNECT TIME IS NOT RECORDED IN TYPE 14/15/64 (DATA SET) RECORDS
 - ° CANNOT IDENTIFY DATA SET USAGE OF CONNECT TIME
- ° DEVICE CONNECT TIME IN TYPE 30 DD SEGMENTS IS INCORRECT
 - ° AT IBM LEVEL 2
 - ° VALUE IN DD SEGMENTS IS GREATER THAN IN ASID TOTAL
 - ° ASID TOTAL AND RMF TOTALS APPEAR CORRECT