



When Performance/Capacity Becomes a
Performance/Capacity Issue

Or
The Emperor has no Clothes



Copyrights

- The author acknowledges any and all copyrights for any product herein mentioned



Caveat

- The opinions expressed herein are the author's alone and do not necessarily in whole or in part reflect the views of any company with which said author may be associated



Some History - 1983

- MVS/XA Arrives – wow – 24 bit addressing
 - And there are probably still 16 bit applications running
- MXG arrives soon thereafter
- DB2 is still years away
- MIPS are measured in 1 or 2 digits



Some History - 1998

- OS/390
- WLM has arrived
- 40 GB of SMF data to process is described by Barry Merrill and Chuck Hopf at SHARE in Anaheim
- DB2 is beginning to be the big dog in the SMF world
- MIPS are 3 digits



Finally – 2011

- zOS v12 is being installed
- DB2 is running rampant
- CICS is dwindling?
- 40 GB is perhaps an hour or less at some large shops
- Both CICS and DB2 SMF data can be compressed
- MIPS are now 5 digits topping out over 50000



If you cant measure it you cant manage it

- Thus spake CME (now EWCP)
- For 30 years there has been an ongoing struggle to process all of the data available within a reasonable time frame
- It has now reached the point where for some shops, SMF post-processing runs 7X24!!!
- If there is a failure along the way catching up is ugly



Enough!





An Example

- A relatively small shop
- 2098-T04
- 2 LPARs – 1 is the ‘sandbox’ with very little activity
- In a single day – 17GB of SMF data




The SMF Data

SMFID / SUBTYPE	TOTAL	# RECORDS	% RECS	BYTES	% BYTES
6.000		1129	0.01	606K	0.00
21.000		3580	0.03	307K	0.00
26.000		25033	0.23	10M	0.06
30.001		25032	0.23	9M	0.05
30.002		28915	0.26	86M	0.48
30.003		97268	0.88	128M	0.71
30.004		98009	0.89	132M	0.73
30.005		25078	0.23	48M	0.27
30.006		1585	0.01	1572K	0.01
70.001		192	0.00	736K	0.00
71.001		192	0.00	355K	0.00
72.003		19968	0.18	25M	0.14
72.004		192	0.00	3620K	0.02
73.001		192	0.00	3797K	0.02
74.001		1824	0.02	54M	0.30
74.002		96	0.00	796K	0.00
74.005		3648	0.03	70M	0.39
74.008		96	0.00	846K	0.00
75.001		1248	0.01	321K	0.00
77.001		96	0.00	44K	0.00
78.003		96	0.00	1221K	0.01
100.000		384	0.00	916K	0.00
100.001		384	0.00	1821K	0.01
100.002		384	0.00	278K	0.00
100.004		384	0.00	197K	0.00
101.000		4370062	39.54	9409M	51.98
101.001		4367326	39.52	3274M	18.09
102.105		2642	0.02	5555K	0.03
102.106		384	0.00	1065K	0.01
102.172		538	0.00	689K	0.00
102.191		2	0.00	1372	0.00
102.196		3	0.00	1848	0.00
102.258		1437	0.01	378K	0.00
102.337		1	0.00	562	0.00
110.001		110411	1.00	3375M	18.65
110.002		41516	0.38	331M	1.83
TOTAL		11051464	100.00	17G	100.00



SMF Records

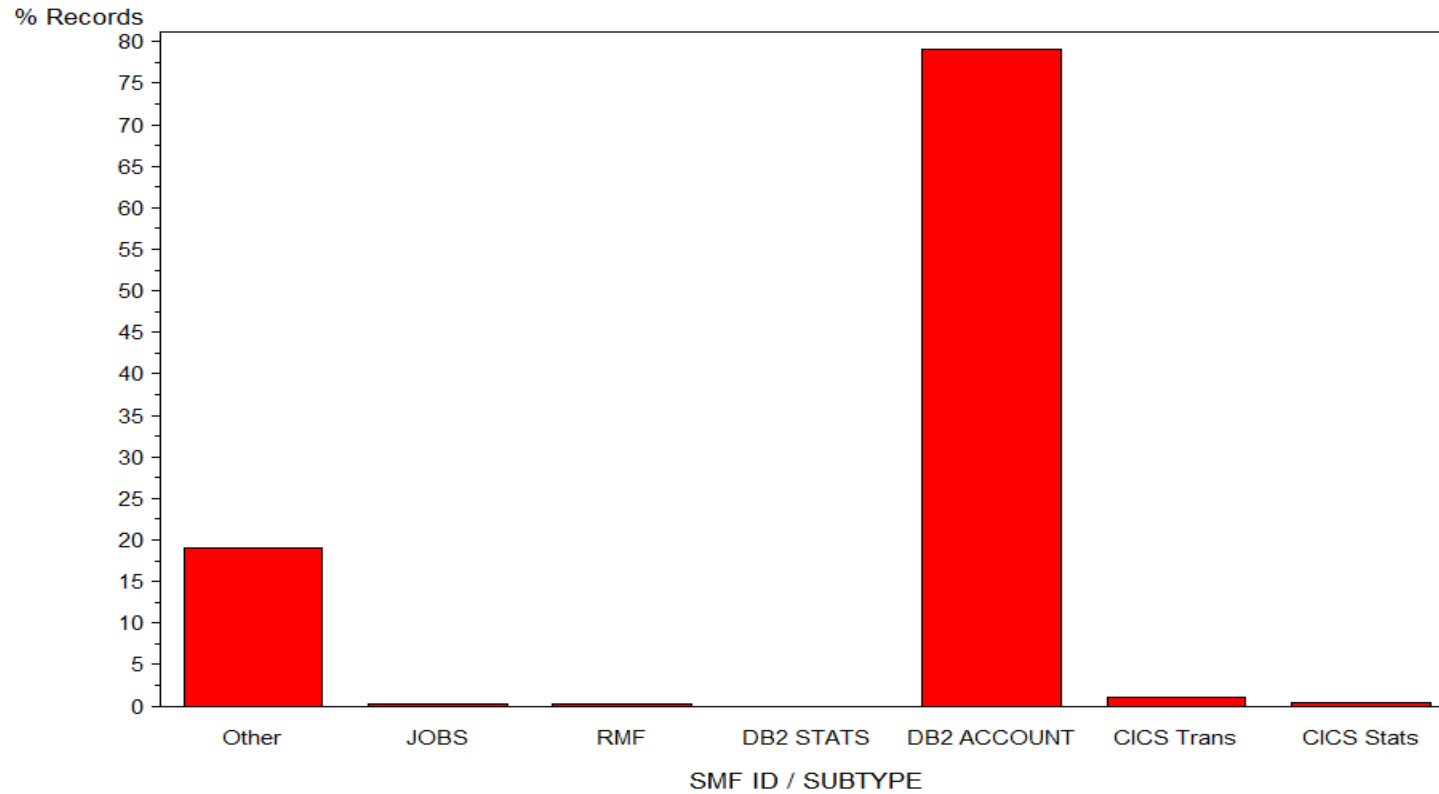
Percent of SMF Records by Type of Data





SMF Bytes

Percent of SMF Records by Type of Data





Daily Post-Processing

Job	CPU Time	Elapsed Time	EXCP
Base PDB	4.76	9.50	64513
DB2/CICS	44.99	102.20	1641472
Reports	1.57	3.30	31112

Three production jobs to process data
Begin at 5AM normally finish about 6:30




A Larger Shop

Type of Data	# records	% Records	# Bytes	% Bytes
CICS Statistics	205522		0.19 4383M	1.07
CICS Transactions	6487774		5.93 203G	50.69
DB2 Accounting	51722487		47.28 85G	21.32
DB2 Statistics	3		06K	0
JOB	981079		0.93411M	0.83
MQ	12962363		11.85 67G	16.74
Other	22972431		21 17G	4.48
RACF	1231262		1.13 321M	0.08
RMF	590654		0.54 8321M	2.02
WLM type 99	12234117		11.18 11G	2.77
Total	109387692		100 401G	100



A Larger Shop


Large Site SMF Distribution





A Larger Shop

Large Site SMF Distribution






Testing Some Options

- Test 1 – Baseline – process all data
- Test 2 – Suppress processing of CICS data
- Test 3 – Suppress processing of DB2 data
- Test 4 – Suppress processing of type 74 data
- Test 5 – Suppress DB2 and CICS
- Test 6 - Suppress DB2, CICS, and type 74
- Test 7 - Suppress processing of DB2 accounting data – type 101
- Test 8 – Suppress processing of DB2 acctg and CICS transaction data
- Test 9 – extract 1 hour of CICS and DB2 transaction data



CPU Time per Test


Tests Run





Elapsed Time


Tests Run





IO Time


Tests Run





EXCP Counts

Tests Run





And What Does it Mean?

- If DB2 is 80% of the data, it will likely be 80% of the processing time. Same for CICS.
- CICS and DB2 statistics are miniscule in terms of processing time
- Type 74 may be larger in shops with 10s of thousands of devices – this is not one and the type 74 data really does not matter here
- Shop does not run MQ and that can also be large
- Best run time eliminates the processing of CICS and DB2 transaction data



So how do we fix it?

- It depends...
 - There are three kinds of data here (in my mind)
 - Accounting – if you are doing detailed chargeback or have a group of surly auditors this can be all encompassing
 - Tactical – data needed for problem solving
 - Strategic – data for longer range planning



Accounting Data

- Types 6 26 30
- Possibly DB2 accounting and CICS transaction data
- Possibly HSM data and DCOLLECT data
- Possibly tape management data
- If there is detailed chargeback this may need to be retained for long periods



Tactical Data

- May encompass most of the accounting data
- RMF
- Tape Mount Monitor
- Other monitors



Strategic Data

- All of the above but highly summarized with only the variables that are needed
 - Week and shift (If shift is important)
 - Some variables may be archaic – anyone still have a 3350 or a 3380? EXCP counts for those devices are fundamentally useless since they don't translate well and are almost certainly missing in any case



What Are the Options?

- Forego the processing of DB2 and CICS transaction data every day but process as needed for problem solving
- Just as it was in 1998, one option is to split the data into bite sized pieces for processing in separate tasks though some of the bites are fairly huge
- Also now an option to 'outsource' the processing of SMF data to a UNIX or Windows platform



Option 1 – No DB2/CICS/MQ

- If detailed chargeback is being done – not an option
- Reporting of transaction volumes, CPU consumption, response time can (in some cases) be done from RMF type 72




Reporting for DDF

- Report granularity is dependent on how well (or badly) the queries are identified
- USERID (QWHCAID) may not be adequate
- PLAN (QWHCPLN) will always be something like DISTSERV
- Developers may not want to take the time to properly identify the work
- WLM only sees the first action for a query



DDF Reporting


DDF Resource Consumption





DDF Reporting


DDF Resource Consumption





DDF Reporting


DDF Resource Consumption





DDF Reporting

DDF Resource Consumption






CICS Reporting

- If response time goals are being used, something similar to DDF can be done by breaking the transactions down into report classes
- CPU time must come from the base service class since it does not exist in the transaction report classes
- Transaction count and response time from the base service class have to be ignored



CICS Reporting


CICS Resource Consumption





CICS Reporting


CICS Resource Consumption





CICS Reporting

CICS Resource Consumption






CICS Reporting

CICS Resource Consumption

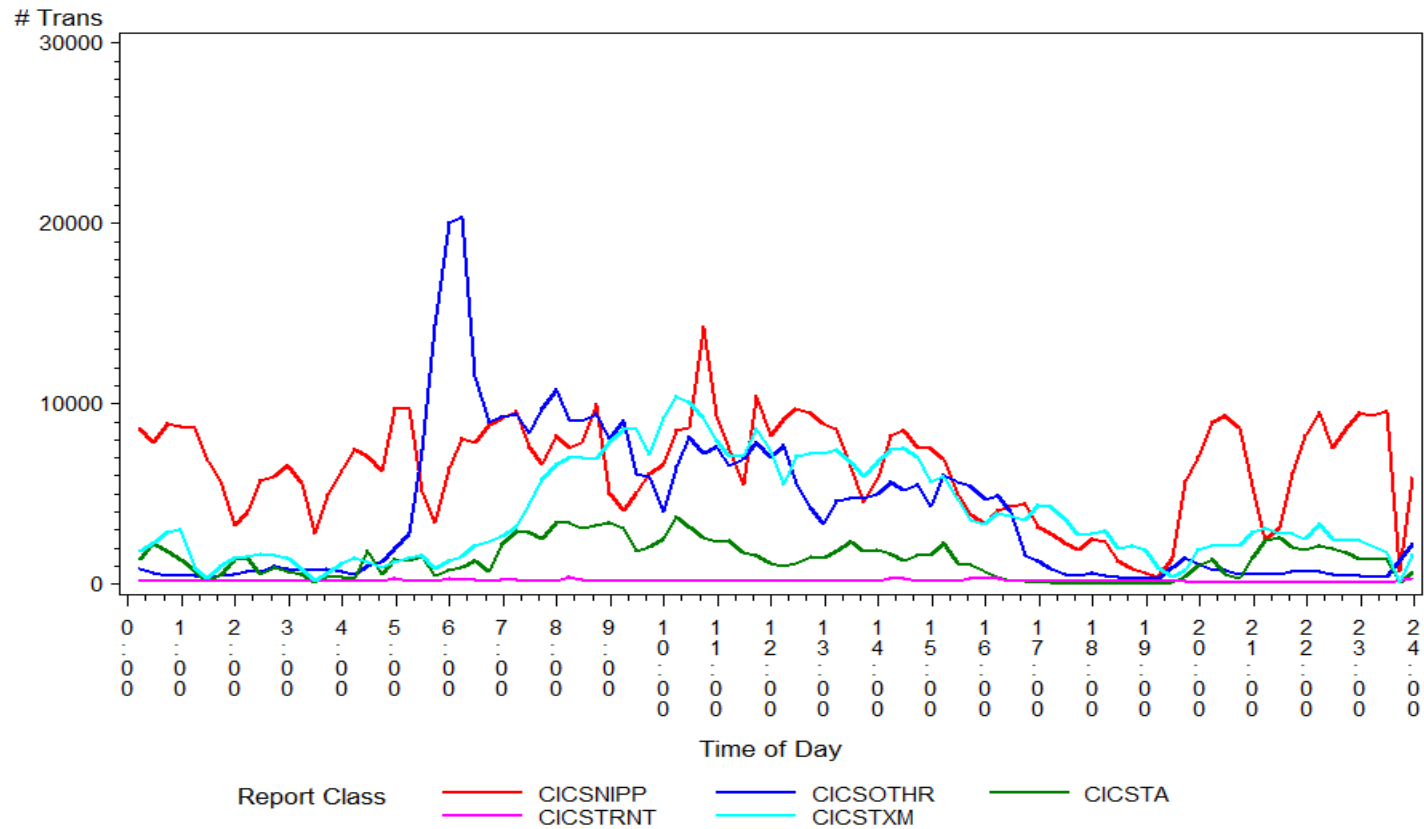
Total zIIP Time





CICS Reporting

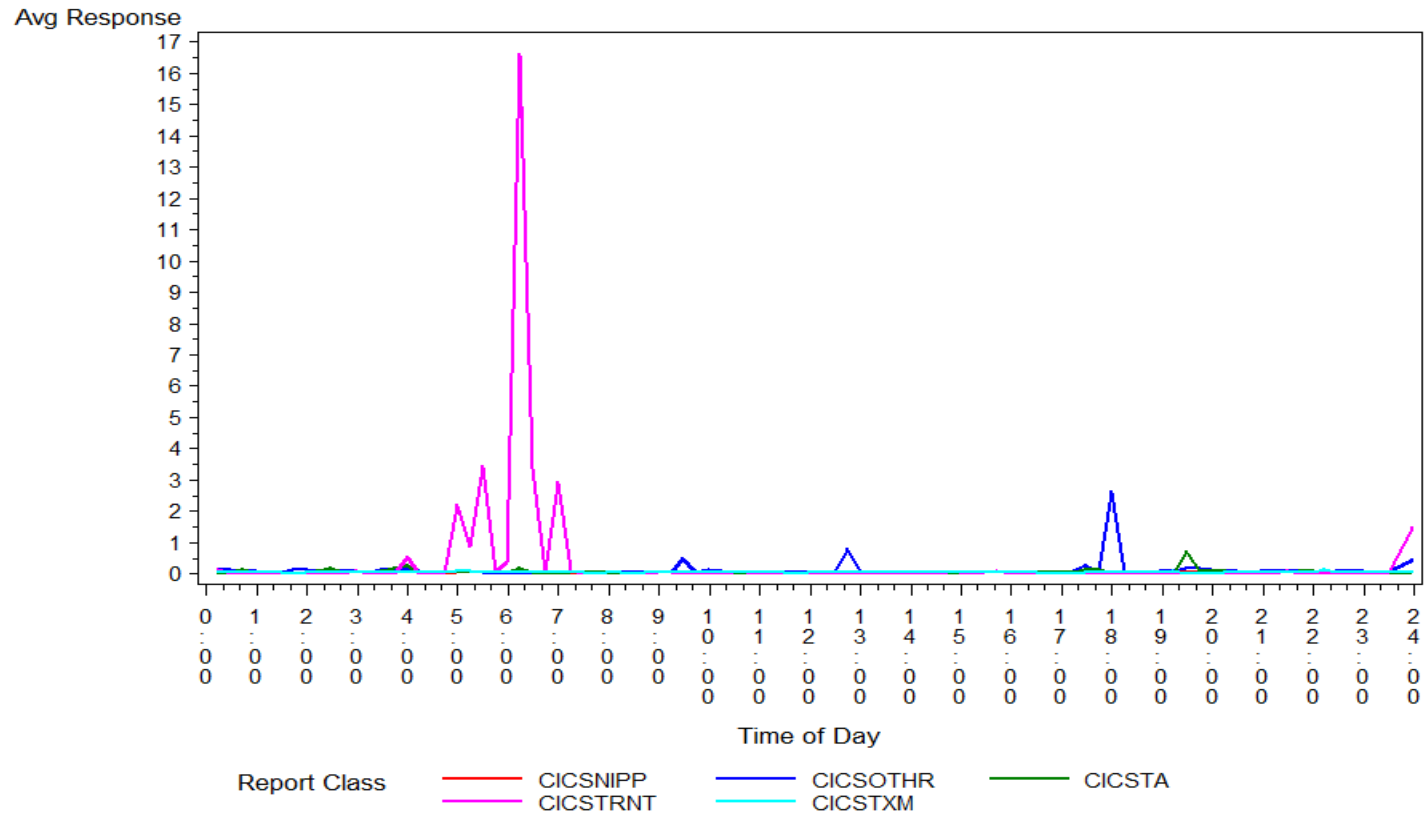
CICS Resource Consumption





CICS Reporting

CICS Resource Consumption





What about planning?

- Report classes can be as granular as may be needed but in the case of CICS, CPU time will not be captured at the transaction level
 - Sample 1 hour per day to build a baseline?
 - 1 hour of CICS and DB2 can be extracted and processed relatively quickly (test 9)
 - Use the samples to project CPU by transaction



What about other stuff?

- DB2 – will most likely be batch, DDF, or CICS
 - Batch will be in the type 30 but the DB2 time and resources will not be separated
 - DDF is covered
 - CICS like batch will not have the DB2 resources broken out except in the samples taken
- Other work will be in the TYPE 72 data as it is now



Option 2 – Divide and Conquer

- Starting with MXG 29.04 sample JCL and code is provided to split the processing of SMF data into parts
 - JCLSPSMA – Read CICS transaction data
 - JCLSPSMB – Read DB2 accounting data
 - JCLSPSMC – Read IO related records
 - JCLSPSMD – Read MQ data
 - JCLSPSME – Read all the rest



Option 2

- Samples no longer use a special PROC for MXG
- Problem have been encountered when the NLS options in SAS get crossways with the MXG PROC
- SOURCLIB/LIBRARY concatenations built dynamically
- Code is largely based on UTILBLDP



Option 2 - JCLSPSMA

```
//S1 EXEC SAS,  
// CONFIG='UXMCBH.MXG.SOURCLIB.V2903(CONFIMXG)'  
//MXGNAMES DD *  
%LET MXGSOURC=MXG.SOURCLIB;  
%LET MXGFORMAT=MXG.FORMATS;  
%LET MXGUSER1=MXG.USERID.SOURCLIB;  
%LET MXGUSER2=;  
%LET MXGUSER3=;  
//WORK DD UNIT=(SYSDA,16),SPACE=(CYL,(500,500))  
//CICSTRAN DD DSN=MXG.DAILY.CICSTRAN(+1),  
// UNIT=TAPE ESOTERIC,DISP=(,CATLG,DELETE)  
//CICSBAD DD DSN=MXG.DAILY.CICSBAD(+1),  
// SPACE=(CYL,(5,5)),DISP=(,CATLG,DELETE)  
//SMF DD DSN=YOUR.DAILY.SMF.CICS(0),DISP=SHR
```



Option 2 - JCLSPSMA

```
//SYSIN DD *  
%LET MACKEEP=%QUOTE(  
  _N110  
  MACRO _S110 %  
  MACRO _WCICTRN CICSTRAN %  
  MACRO _LCICTRN CICSTRAN.CICSTRAN %  
  MACRO _WCICBAD CICSBAD.CICSBAD %  
  MACRO _LCICBAD CICSBAD.CICSBAD %  
);  
%INCLUDE SOURCLIB(TYPE110);
```



Option 2 - JCLSPSME

```
Option 2 - JCLSPSMW //S1 EXEC SAS,  
// CONFIG='UXMCBH.MXG.SOURCLIB.V2903(CONFIMXG)'  
//MXGNAMES DD *  
%LET MXGSOURC=MXG.SOURCLIB;  
%LET MXGFORMAT=MXG.FORMATS;  
%LET MXGUSER1=MXG.USERID.SOURCLIB;  
%LET MXGUSER2=;  
%LET MXGUSER3=;  
//WORK DD UNIT=(SYSDA,16),SPACE=(CYL,(500,500))  
//PDB DD DSN=MXG.DAILY.PDB(+1),DISP=(,CATLG,DELETE),  
// SPACE=(CYL,(500,500))  
//SPININ DD DSN=MXG.DAILY.SPIN(+0),DISP=SHR  
//SPIN DD DSN=MXG.DAILY.SPIN(+1),DISP=(,CATLG,DELETE),  
// SPACE=(CYL,(50,50))  
//SMF DD DSN=YOUR.DAILY.SMF.SPLITPDB(0),DISP=SHR
```



Option 2 - JCLSPSME

```
%LET SPININ=SPININ;
%LET MACKEEP=%QUOTE(
  MACRO _WCICTRN _NULL_ %
  MACRO _WCICBAD _NULL_ %
  MACRO _WDB2ACC _NULL_ %
  MACRO _WDB2ACP _NULL_ %
  MACRO _WDB2ACB _NULL_ %
  MACRO _WDB2ACG _NULL_ %
  MACRO _WDB2ACR _NULL_ %
  MACRO _WDB2ACW _NULL_ %
  MACRO _SDB2ACP %
  MACRO _SDB2ACB %
  MACRO _SDB2ACG %
  MACRO _SDB2ACR %
  MACRO _SDB2ACW %
);
%UTILBLDP(BUILDpdb=YES,
  SUPPRESS=74 115 116,
  MXGINCL=ASUM70PR ASUMTAPE ASUMTMNT ASUMTALO,
  OUTFILE=INSTREAM);
%INCLUDE INSTREAM;
```




Option 2

- JCLSPUOW – combines CICS, MQ, DB2 data by unit-of work
- JCLSPCPY – copies summary CICS/DB2 datasets into base PDB
- JCLSPWEK – weekly job
- JCLSPMTH – monthly job



Option 3 - Outsourcing

- Not to a foreign country – only to a foreign operating system/platform
- Can be Windows or UNIX (or anywhere else you can run SAS)
- Same set of jobs as in option 2 but members start with BLD
- Uses a 'pseudo-GDG' structure



Option 3 – Pseudo-GDG

- Directories are built and managed dynamically based on user parameters for how long to keep them and where to place them
 - Dddmmyy – daily
 - Wddmmyy – weekly
 - Mddmmyy – monthly
 - Tddmmyy – trend
 - Sddmmyy – spin
 - CICSddmmyy – CICSTRAN
 - DB2ddmmyy – DB2ACCT



Option 3 - BLDSPSMA

```
%LET MACKEEP=%QUOTE(  
  _N110  
  MACRO _S110 %  
  MACRO _WCICTRN CICSTRAN %  
  MACRO _LCICTRN CICSTRAN.CICSTRAN %  
  MACRO _WCICBAD PDB.CICSBAD %  
  MACRO _LCICBAD PDB.CICSBAD %  
);  
%VMXGALOC(BASEDIR=C:\MXG);  
%INCLUDE SOURCLIB(TYPE110);  
RUN;
```



Option 3 - BLDSPSME

```
%LET MACKEEP=%QUOTE(  
  MACRO _WCICTRN _NULL_ %  
  MACRO _LCICTRN _NULL_ %  
  MACRO _WCICBAD _NULL_ %  
  MACRO _WDB2ACC _NULL_ %  
  MACRO _LDB2ACC _NULL_ %  
  MACRO _WDB2ACP _NULL_ %  
  MACRO _WDB2ACB _NULL_ %  
  MACRO _WDB2ACG _NULL_ %  
  MACRO _WDB2ACR _NULL_ %  
  MACRO _WDB2ACW _NULL_ %  
  MACRO _SDB2ACP %  
  MACRO _SDB2ACB %  
  MACRO _SDB2ACG %  
  MACRO _SDB2ACR %  
  MACRO _SDB2ACW %  
  _N74  
  MACRO _S74 %  
);  
%UTILBLDP(BUILD PDB=YES,  
  USERADD=TMNT/238,  
  SUPPRESS=74 115 116,  
  MXGINCL=ASUM70PR ASUMCACH ASUMTAPE ASUMTMNT ASUMTALO,  
  OUTFILE=INSTREAM);  
%BLDSMPDB(  
  AUTOALOC=YES,  
  BASEDIR=C:\MXGTEST,  
  ERASEPDB=NO,  
  RUNDAY=YES,  
  BUILD PDB=INSTREAM,  
  RUNWEEK=NO,  
  RUNMNTH=NO  
);  
RUN;
```



Option 3 – A Test

Intel dual core 2.2Ghz

Win 7 Ultimate 32 bit

SAS 9.2

	User CPU	System CPU	Run Time	Memory	Total Elapsed	% Reduced Elapsed Time	% Reduction in CPU Time
BLDSIMPL	0:07:31	0:02:23	0:34:59	423165k			
BLDSPSMA – CICS	0:00:14	0:00:05	0:04:30	26496k	0:27:30	21.38%	14.52%
BLDSPSMB - DB2	0:04:29	0:01:28	0:24:38	425984k			
BLDSPSMC – IO	0:00:23	0:00:12	0:05:30	148452k			
BLDSPSME – SPLIT	0:01:04	0:00:38	0:06:45	98072k			
BLDSPUOW	0:00:12	0:00:04	0:02:52	39700k			

2098-T04 zOS 1.10

SAS 9.1.3

	User CPU	System CPU	Run Time	Memory	Total Elapsed	% Reduced Elapsed Time	% Reduction in CPU Time
JCLSIMPL	0:16:42		0:19:54	127837K			
JCLSPSMA - CICS	0:00:31		0:00:42	31991k	0:13:36	31.66%	14.27%
JCLSPSMB - DB2	0:09:54		0:12:42	33471K			
JCLSPSMC - IO	0:00:54		0:01:36	44401k			
JCLSPSME - SPLIT	0:02:30		0:03:42	88933k			
JCLSPUOW - UOW	0:00:30		0:00:54	22927K			

SMF Data 3738MB

Records 2353851

DB2ACCT OBS 917607

CICSTRAN OBS 255522



Some Caveats

- All of these jobs A-E are designed to run concurrently but...
 - On zOS you must have separate datasets
 - On ASCII the same directories are used for all jobs
- Locking on ASCII is at the level of the individual SAS dataset but on zOS it is at the level of the SAS data library (unless you happen to have SAS/SHARE).
- Jobs on ASCII might run faster spread across multiple platforms



So, What to do?

It Depends!!!!



Not Doing Chargeback?

- Any of the options will work
 - Not processing DB2/CICS is a management decision
 - If management is unhappy with the cost of SAS on zOS, ASCII might be a better choice



Got Chargeback?

- May not have a choice other than running on zOS or ASCII
- Unless there is enough granularity in the DDF and CICS report classes and management will buy into charging for CICS based on sampling transactions and applying it to the counted transactions



What to Keep and How Long?

SMF ID / SUBTYPE	Process Daily	Retention	Note
6	Yes	Ask Auditing	
21	Yes	Ask Auditing	
26	Yes	Ask Auditing	
30	Yes	Ask Auditing	
70	Yes	3-4 years	
71	Yes	3-4 years	
72	Yes	3-4 years	
72	Yes	3-4 years	
73	Yes	3-4 years	
74	Yes	3-4 years	
75	Yes	3-4 years	
77	Yes	3-4 years	
78	Yes	3-4 years	
99	No	10 days	
100	Yes	3-4 years	
101	No	2 weeks	Accounting Data
102	Maybe	Ask Auditing	Some accounting data may be 102
110	No	2 weeks	Transaction data
110	Yes	3-4 years	Statistics
217	Maybe	Ask Auditing	HSM
218	Maybe	Ask Auditing	HSM
238	Maybe	Ask Auditing	MXG Mount Monitor
OTHER	Maybe	Ask Auditing	