INTRODUCTION TO CPE USING THE MAINFRAME AND PC

by
Bruce L. Green
Senior Systems Programmer
MIB, Inc.

INTRODUCTION

This paper will provide an overview of computer performance evaluation at MIB using both the mainframe and personal computer. MIB is a computer service bureau to member life insurance companies. At MIB, we have two 3083-B CPUs running MVS SP Extended Architecture (XA): a sixteen-megabyte machine dedicated to the "Online" Search system and a sixteen-megabyte batch/TSO/CICS machine. Both systems run JES2 and VTAM. The CPE function is a task of the Systems Programming group. There are approximately 50 application programmers here. MIB's production system is written in assembler language and its data base was written and developed internally. Our MIB Data Base is a multi-volume VSAM cluster that spans eight 3380E volumes.

BACKGROUND

The end of the 1980s has signaled a dramatic increase in the computer equipment acquisition at MIB. In 1980, MIB converted its RCA SPECTRA/70 systems to two IBM 3031s and current configuration is two 3083-Bs. Future plans foresee 3090x systems being installed within three years. The upgrade to 3083-B CPUs (six mip machines with sixteen channels each) began in mid-1987. The decision to obtain 3083 CPUs was based upon performance, expansion capability, reliability and price. Although the 3083 CPUs were no longer manufactured and available from IBM, they were available on the used equipment market at rates which substantially increased the price/performance ratio over newer IBM hardware.

The 3083 upgrade was completed in January, 1988. Associated with the mainframe upgrades were other hardware upgrades: 3380-E (double-density) DASDs (reduced service time), cache memory subsystems (faster access to data), and a 3480 magnetic tape cartridge subsystem (faster backup time). Each of these additions have made a significant impact on the throughput of jobs in the system, and has improved programmer productivity.

The computer equipment configuration, which MIB has employed since 1966, consists of dual CPUs and front-end communications processors (FEP), and a mirror-image of all online data base files on separate DASD. This configuration provides complete backup capability in the event of hardware failures and ensures a continuity of service to our customer base.

Although the primary purpose of the second CPU and FEP is to provide backup to the primary online system, they are used during prime-time for other production jobs and application development.

Operating systems have evolved from RCA TDOS to MVS/XA. In 1985, CICS was acquired to develop VTAM application programs for our member companies. The migration to XA has been very successful. Installed via CBIPO, we had "initial operation complete" within one month. The MVS SP user's single biggest reason for conversion to XA is the virtual storage constraint (VSC) problem. A dramatic rewrite of the Input/Output Supervisor (IOS), under XA, has shown a 60% improvement in throughput in our "Online" search program.

DASD configuration, based on megabytes, have increased more than seven-fold over the last ten years. Mainframe power, in terms of MIPS, have increased more than five-fold. Mainframe storage, based on memory, has increased sixteen-fold. The number of terminals have increased more than twenty-fold.

CPE EVOLUTION AT MIB

The capacity and performance management process is gradual. At the time (1985) we installed MXG ® software (for computer performance evaluation) to create our baseline performance database (PDB), the hardware configuration was two 3031s (each a one "mip" machine) and the 64 3350 DASDs were all shared among both CPUs. It was found that with the restrictions of channel and controller configuration, we had constant DASD contention. Also, as we observed, both the test and batch job CPU utilization had increased significantly. The annualized increase in CPU utilization was predicted around 25% (January, 1985 System A and B average prime-time utilization: 65% and 35%, respectively. January, 1986: 85% and 70%).

Over the past year, the reduced price of PS/2s made an attractive addition to our CPE configuration. We were able to convince upper management that the IBM 3270/PC should be used as a CPE reporting tool. The mainframe's availability of base SAS® software, created low-resolution graphics (connect-the-dots using fiber-tipped pens!). With the addition of a PS/2 came base SAS and SAS/GRAPH® software on the PC(connect-the-dots automatically). So now the methodology:

1) collect SMF/RMF data on the mainframe
2) reduce data to a PDB using MXG
3) download a summarized subset to the PC
4) create standard or "ad hoc" reports

REPORTING

The type of reporting that is done is many and varied, due to the size of our installation. Hardware utilization and projection are some of the reports. The reports have ranged from RPS and cached-DASD analysis to TSO timeout reports. Each of the reports have had a purpose and, in most situations, a different destination.

EXCEPTION ANALYSIS

Closely related to our CPE function is the day-to-day computer operation. The stand-alone OMEGAMON monitors, in both the Systems and Operations area clearly aid in the area of exception analysis (as well as many other MVS "areas").

Several problems have been resolved with the use of MXG/PDB. There have been several occasions to execute the "BUILDPDB" against "live" SMF data. Incidentally, even with the ACF2 software, I have been asked to determine who scratched a data set. With the MXG software, I would execute SAS with the member "TYPE1718" as input against either archived or live SMF data, and then issue a "PROC PRINT".

JOB ACCOUNTING

Job Accounting, a term which needs no explanation, has surfaced at this installation. Over the past decade, the price of hardware has dropped dramatically, while the costs associated with software development has now become the dominant portion of a total system budget. This rapid rise has sparked management's interest in determining computer resource accountability. In early 1989, we investigated both in-house development and purchased systems. All of the independent software suppliers are expensive. My recommendation was to use the MXG software to gather resource data and send it into the General Ledger, then let the already existing corporate charge-back system handle data processing as well. Many organizations, however, believe that they should manage and handle billing of data processing as a part of the DP function.

We are currently pursuing the in-house development of a chargeback system. The SAS/CPE Starter Set version 3.1, is available to all MXG users. This new version has many enhancements, including the new Cost Analysis option. This SAS/AF® software application runs on the base SAS mainframe and is absolutely free. SAS/FSP® software is required to run the "Cost Analysis" option (SAS/AF® is required if you would like to customize the panels). We have found that, being a chargeback "fledgling", this system is sufficient for what it has to offer (and the price is right). It contains many fine examples of reports produced directly from the MXG data bases. Basically, for Job Accounting purposes, the CPE starter set receives as input the PDB.JOBS, after one has edited a "rate table", and produces a Cost Data Base (CDB).

REPORTING DETAILS

There are two major types of reporting that are being done - performance and capacity planning. All programming, whether original, or creatively plagiarized is written in SAS. A regular series of reports (daily, weekly, and monthly) are, for the most part, complete and in production. Since capacity management deals with the computer resources for supporting a prescribed workload, the total measurement of computer resource utilization is a direct description of the capacity and an indirect description of the workload characteristics. The major computer resources are CPU memories, controllers, channels, networks, disk drives, tape drives and printers. Due to our sophisticated operating system with multiprogramming and multisystem accesses, each resource element of the system is not independent of other elements.

The most important question in capacity planning is "What kind of computer configuration do we need to support our peak load hours in terms of all resource elements?" The final configuration must include the capacity to handle the effective maximum of all resources (CPUs, memories, controllers, channels, tape, disks, network and printers). In addition, we must compute the effective resource rate with consideration given to interference, overhead (approximately 7%), scheduled maintenance, and unscheduled outages.

CPU

CPU utilization reports come in several different "flavors". Who the end-user is depends on the report's content. We have defined three different shifts (prime, night, and weekend). But, more importantly is the "peak-time" (prime-time less lunch-time) utilization. I have written reports that break down each weekday into its normal first, second and third shifts.

Current CPU prime-time utilization is approximately 40% on each system. The increasing use of CICS has had the most impact on computer resources. Conversely, the increased use of 3270/PCs has offloaded a small percentage of both CPU and DASD resource consumption.

CPU PRIME-SHIFT UTILIZATION ANALYSIS

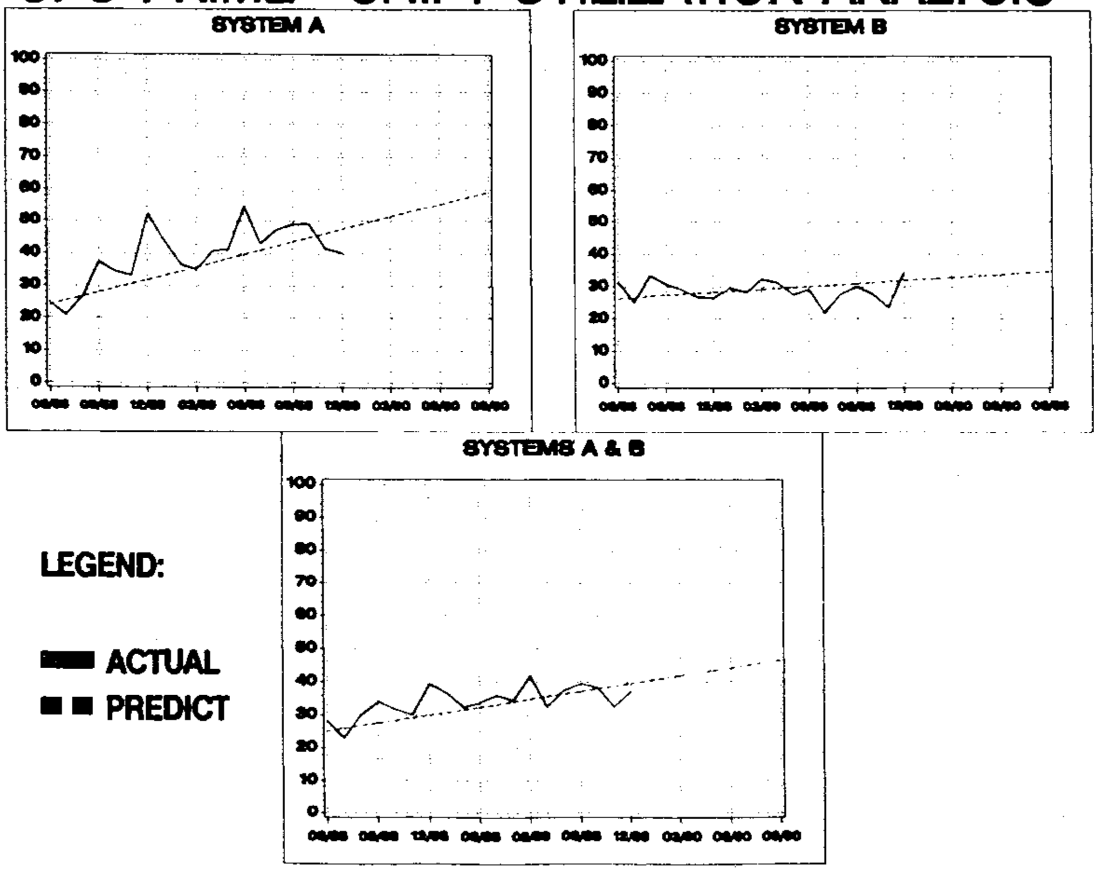

DASD

DASD utilization has become a problem, with "test" device utilization around the 90% mark. It became apparent that a trend analysis had to be performed. Twice a week, a job scans the "test" VTOCs, groups them, and performs a "PROC APPEND" to add the current observations to the history file.

FDR/ABR is capable of generating a wide range of reports for the DASD Manager. Also supported is a generalized report writer ("FDREPORT"). FDREPORT functions as a data extractor for use with the VTOC. What has given us the most information, is FDREPORT's ability to report on information within the ICF VSAM catalog and VVDS data set.

RPS TECHNOLOGY

XA is clearly an architecture designed specifically to move the management and control of the system and peripheral components as close to those components as possible. This is done through extensive use of microcode and offloading of overhead.

Dynamic Path Reconnects (DPR) for 3380s is a tremendous function. This occurs because although MVS SP allows a start I/O to be initiated on a primary or alternate path, at RPS (rotational positional sensing) reconnect time it must finish the I/O operation on the same path it started on. RPS on DASD is a software-assisted feature that lets the software tell the DASD hardware where it expects the requested record to be on the

```
                        3380 "TEST" DASD UTILIZATION BY GROUP ANALYSIS                          1

                                          GROUP=ALP

WEEK OF      #TRACKS    MIN                                                                  MAX
                        302                                                                19874
                        *-------------------------------------------------------------------*
06FEB89         8296    | 4-a3----S-------------T                           |               |
13FEB89         8267    | 4-a3----S-------------T                           |               |
21FEB89         8591    |  4a3----S--------------T                          |               |
27FEB89         8434    |  4a1----S-------------T                           |               |
02MAR89         9342    | 4--a-----S---------------T                        |               |
06MAR89         9636    |  4-a-----S----------------T                       |               |
09MAR89        11807    | 4-21-3---------S-----------------T                |               |
13MAR89        10372    | 4-2--a---S------------------T                     |               |
16MAR89        10419    | 4--231--S-------------------T                     |               |
20MAR89        11019    |  4-2-a---S---------------------T                  |               |
23MAR89        14646    |      a-3---------S--------------------------|-T               |
27MAR89        16290    |     24-3--1-----S--------------------------|-------T          |
30MAR89        15075    |    4---32-1S-------------------------------|--T               |
03APR89        12927    |  4-32-----a-------------------------T       |                  |
06APR89        15036    |  4-32-----1----S------------------------|--T                   |
10APR89        12776    |  a-3------1-S-------------------------T     |                  |
13APR89        15128    |      a3----1-S---------------------------|--T                 |
18APR89        14687    |    2a-----1-----S------------------------|-T                  |
20APR89        14938    |  4-23-----1----S-------------------------|--T                 |
24APR89        12043    |  42--3--S-1----------------------T          |                  |
27APR89        13240    |  42---3--S1---------------------------T     |                  |
01MAY89        11410    |  423----S-1-------------------T             |                  |
04MAY89        13798    |  4--3--2S-1------------------------------T                     |
08MAY89        13727    |   43---2S-1-----------------------------T|                     |
11MAY89        16462    |   4--3---1-S2----------------------------|-------T            |
15MAY89        14522    |  4-3----S12------------------------------|-T                  |
18MAY89        16095    |   4--3---1S2-----------------------------|------T             |
22MAY89        14673    |   43-----aS------------------------------|-T                  |
25MAY89        17987    |   4--3----1-2-S--------------------------|------------T       |
30MAY89        14769    |   .a----1S2------------------------------|-T                  |
01JUN89        17016    |   4---3--1S--2---------------------------|---------T          |
05JUN89        15350    | 4---3----a--2----------------------------|---T                |
08JUN89        16920    |  4-------aS2-----------------------------|---------T          |
12JUN89        16547    | 4---------aS-----------------------------|-------T            |
15JUN89        18310    | 4------3---a---1-------------------------|-------------T      |
19JUN89        17271    |4-------3-2S----1-------------------------|----------T         |
22JUN89        19874    | 4----------aS-----1----------------------|----------------T|
26JUN89        14566    |    4--S--3a------------------------------|-T                  |
29JUN89        15898    |   4----S--a-2----------------------------|-----T              |
```

track, so that it can disconnect from the channel until that record comes under the read/write head. Without this software assistance, the hardware must remain connected to the channel while it searches for the requested record (looking at each record as it comes under the read/write head to see if it is the correct one). Under XA the RPS reconnect can proceed on either the primary or alternate path and can even be serviced by a different CPU. In this way XA does away with CPU and path affinity and the price to be paid in extra overhead for this increased efficiency is wisely offloaded to the channels. The result is only a modest increase in IOS path length, up from 3600 instructions under MVS SP to around 4000 under XA. The channels on 3080s (and 3090s), because they are separate processors, do the rest of the work.

An RPS miss is classified as: when the device is ready to transmit the designated block of data, the attempt to reconnect fails, due to another I/O. The I/O request is queued for retry, but in this case the retry cannot occur until the disk completes a rotation and is again positioned to read the desired block of data.

CACHE ANALYSIS

The introduction of cached storage controllers introduced a new function of the IOS. MIB installed two 3880-23 storage controllers in a dual-frame configuration, each with sixteen megabytes. The MXG/PDB captures Cache RMF statisics. The caching of attached disk volumes was fairly straightforward. Of the sixteen 3380-E volumes (connected to the cache controller), eleven are currently "cached." Tracking the performance characteristics of the cache controllers, using both RMF type 74 records (i.e., PDB.TYPE74) and the RMF Cache DASD Reporter FDP SMF records (i.e., PDB.CACHE) is a simple report to review [1] (the SAS MXG/PDB source code, self-authored, is available on the MXG distribution media).

M E D I C A L   I N F O R M A T I O N   B U R E A U 1

SAS PERFORMANCE DATA BASE

MONTHLY CACHED DASD STATISTICS FOR THE MONTH OF: JANUARY, 1990

REPORTING PERIOD: 8:15:00 TO 15:45:03

LCHAN=1

DEV	VOLUME	I/O RATE	RESP TIME (MS)	READ REQS /SEC	READ HITS /SEC	READ HIT% /SEC	RD/WT RATIO	WRITE REQ/SEC	VOLUME'S SHARE OF SIO	VOLUME'S SHARE OF READ HITS	B/SEC DASD CACHE	B/SEC DASD CHAN	B/SEC DASD BOTH	B/SEC CACHE DASD	B/SEC CACHE CHAN
120	DBTS04*	0.91	10	0.60	0.55	92.76	1.96	0.30	6.68	4.48	1145	94	433	0	1211
121	DBTS09*	0.19	14	0.14	0.13	93.85	3.48	0.04	1.40	1.04	689	940	45	0	652
122	DBTS05*	1.16	10	0.78	0.70	89.48	2.08	0.37	8.50	5.63	2039	46	1666	0	6261
123	DBTS10*	0.32	15	0.21	0.18	86.50	2.19	0.10	2.37	1.49	669	237	466	0	738
124	DBPROO	1.78	21	.	.	.	.	.	13.03	.	.	.	.	.	.
125	DBTS01*	1.59	8	1.25	1.19	95.14	3.81	0.33	11.63	9.59	1739	83	627	0	2475
126	NCPRIA*	2.73	33	1.51	1.29	85.17	1.66	0.91	20.01	10.39	4879	55	1842	0	9559
127	DBTSC8*	0.38	8	0.33	0.29	90.14	7.36	0.04	2.77	2.38	1650	14	226	0	1552
130	MVSRSA*	1.33	6	1.18	1.12	95.27	9.84	0.12	9.72	9.05	841	183	432	0	5037
131	MVSCTA	1.36	19	.	.	.	.	.	9.94	.	.	.	.	.	.
132	D8WRK2	1.70	27	.	.	.	.	.	12.48	.	.	.	.	.	.
133	DBIAFL	0.20	14	.	.	.	.	.	1.47	.	.	.	.	.	.
LCHAN		13.65		5.99	5.45			2.21	100	44.04	13652	1652	5757	1	27684

LCHAN=2

DEV	VOLUME	I/O RATE	RESP TIME (MS)	READ REQS /SEC	READ HITS /SEC	READ HIT% /SEC	RD/WT RATIO	WRITE REQ/SEC	VOLUME'S SHARE OF SIO	VOLUME'S SHARE OF READ HITS	B/SEC DASD CACHE	B/SEC DASD CHAN	B/SEC DASD BOTH	B/SEC CACHE DASD	B/SEC CACHE CHAN
240	CICRIA*	0.29	8	0.21	0.20	96.71	2.72	0.08	1.89	1.63	187	14	39	0	665
241	STG001*	2.66	7	2.16	2.09	96.56	4.48	0.48	17.05	16.85	1494	19	797	0	11434
242	DBPR20	0.75	21	.	.	.	.	.	4.79	.	.	.	.	.	.
243	DBTS06*	1.16	10	0.85	0.78	91.06	2.87	0.30	7.45	6.28	2137	174	1338	0	3403
244	DBTS07*	0.46	9	0.33	0.29	88.03	2.60	0.13	2.97	2.35	1173	44	344	0	1987
245	DBTS02*	1.48	9	1.34	1.28	95.58	4.05	0.33	10.79	10.34	1946	28	978	0	3508
246	DBPR30	0.57	21	.	.	.	.	.	3.67	.	.	.	.	.	.
247	DBTS03*	2.48	9	1.84	1.80	97.55	2.93	0.63	15.88	14.52	1303	37	584	0	8099
250	MVSRSB*	0.56	6	0.52	0.49	94.20	21.66	0.02	3.59	3.98	546	82	198	0	2205
251	D8WRK1	1.94	38	.	.	.	.	.	12.42	.	.	.	.	.	.
252	MVSCTB	1.70	22	.	.	.	.	.	10.92	.	.	.	.	.	.
253	D8WRK3	1.34	23	.	.	.	.	.	8.59	.	.	.	.	.	.
LCHAN		15.59		7.26	6.93			1.97	100	55.96	8784	398	4279	0	31301
		29.24		13.25	12.38			4.18	200	100	22438	2050	10036	1	58985

[1]The Computer Measurement Group 1987 Proceedings (pp. 432ff), "A Method for Reporting Cached I/O Subsystem Performance", by Nancy Nearing, Washington Consulting Group

HISTORICAL ANALYSIS

CPU ANALYSIS

MIB's 3031 processors ran at full capacity for two years (9/85 to 9/87) prior to obtaining the 3083 processors. The installation of the 3083s clearly illustrates a resultant decrease in percentage of utilization. A similar amount of work was being performed, but the increased internal speed of the 3083 lowered the percentage of utilization. Multi-programming capability increased significantly. This allowed for increased programmer productivity to support a growing application development work schedule.

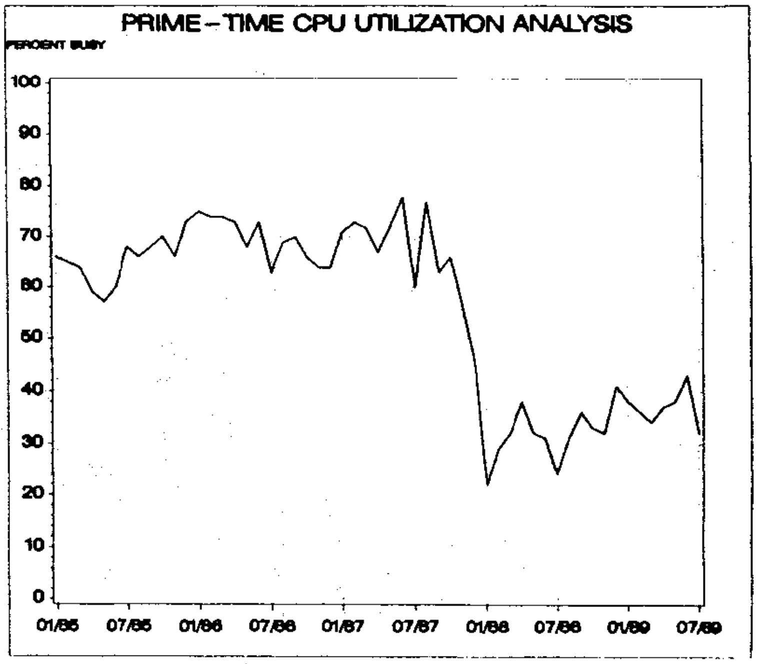

SERVICE UNIT ANALYSIS

Another important measuring criteria MIB utilizes is "Service Units". This measuring technique was devised by IBM due to the inherent problems in measuring resource utilization in shops with multiple CPUs of varying speeds and the inaccuracy of measuring processors in MIPS. The variable is a combination of CPU, I/O, and real storage multiplied by a CPU dependent constant. It is an ubiquious indicator of total amount of work being performed.

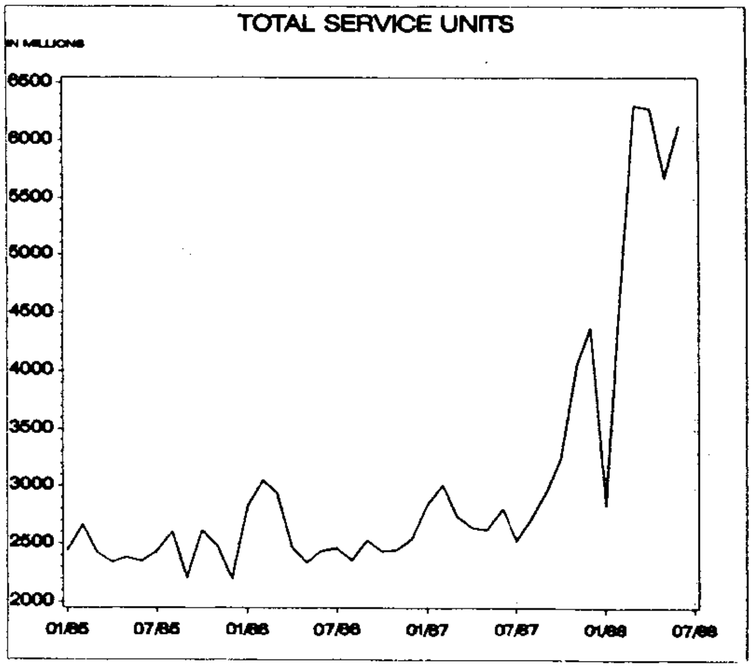

SUMMARY

Without spending large sums of money, we are able to have an efficient resource utilization and performance database system using the SAS System on both the mainframe and personal computer. MXG software is extremely cost-effective in that it helps with all phases of our performance effort. A wide range of management questions can be easily answered with data in the performance data base. As a result, by combining PCs and mainframes for the CPE reporting function, we can reduce load on our mainframe and provide a self-contained responsive application.

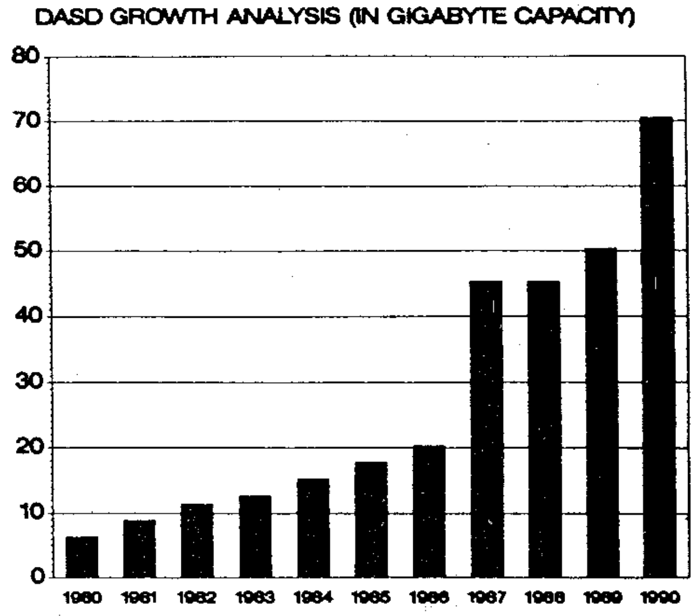

MXG IS A REGISTERED TRADEMARK OF MERRILL CONSULTANTS OF DALLAS, TEXAS, USA.
SAS, SAS/AF, SAS/FSP, AND SAS/GRAPH ARE REGISTERED TRADEMARKS OF SAS INSTITUTE, INC, CARY, NC, USA.