



This article examines the pros and cons of the P/390 system from the technical users' point of view.

BY STEVE THOMPSON AND MARY MARTINELLI

IBM's PC Server 500 S/390: A User Experience

In 1995, IBM announced the PC Server 500 S/390 (P/390). There was much discussion about this machine by IBM at the MVS Vendor Enablement meeting in June, and at SHARE in August 1995, among other venues.

Affiliated Computer Services (ACS) obtained a P/390 in October 1995 specifically for offloading software development and support. Since this system would be primarily used by developers and/or systems programmers, we can tell you what many of the pros and cons of the P/390 system are from a technical point of view. The caveats that must be stated up front include: IBM's literature appears to be based on work for the P/390 using VM. At ACS, we've done all of our work using MVS. So, our descriptions of the system will be from an MVS software developer's and systems programmer's viewpoint, not someone who is doing development via VM.

P/390 CONS

This machine is a single CPU (S/390) with a single channel. It cannot be connected to external 33xx DASD, and currently can't emulate ECKD devices (e.g., 3380K or 3390). Knowing this, it's easy to see that it is not possible to do race condition testing of software, nor the work that would be directed specifically against ECKD devices.

Setting up and operating a P/390 requires knowledge of OS/2, PC architecture, and MVS tuning. The reasons are:

- OS/2 is used as the service processor;
- OS/2 is used as the channel processor;
- DASD (3380/3375) is emulated via OS/2's file system; and
- OS/2 Communications Manager is used to emulate 3270 consoles.

This means that large quantities of SCSI DASD drives are necessary to satisfy the needs of an MVS I/O subsystem. IBM sells different sizes of SCSI drives for this machine, and pushes for the use of RAID 5. The DASD that is set up by OS/2 can be thought of as mini-disks (for those of you who are familiar with this terminology from VM environments).

Another caveat is that the 2.25GB SCSI hot swap drives that IBM sells for this machine are not 2.25GB after formatting. They are actually 2.15GB, and after OS/2 is installed and the High Performance File System (HPFS) is set up, the size decreases to approximately 2GB per drive. When RAID 5 is run, one of these drives per RAID array is lost for error recovery purposes. (While this is expensive, it's not as expensive as RAID 1 where there is one-for-one mirroring.)

Based on our experience, it is our contention that it is necessary to obtain a minimum of 20GB of SCSI DASD to be able to provide for a minimum configuration of DASD for a production MVS system. To do this, it is necessary to obtain 11 of the 2.25GB drives (total). Then, RAID 5 can be run and still have enough space to operate the system.

Now, the major tuning problems we encountered were caused by the way I/O is handled. Part of the I/O bottleneck is caused by the way the system emulates 3375 and 3380 type DASD. The OS/2 drivers for these CKD devices emulate all of the overhead associated with a CKD device (IRG [InterRecord block Gap], Count, Key, Data) as stated in the manuals for each of these devices.

Because this machine has only one channel, avoiding I/O is the best policy. Remember that PAGE data sets are sharing the channel with every other type of I/O, so this system is very sensitive to paging rates. IBM will tell you that this system can sustain 50 I/Os per second with bursts that exceed 100. What this translates to is this: TSO users get hit heavily with swap in/out and page trim. Batch jobs with large working sets will become bogged down because of page I/O and their own I/O needs.

P/390 PROS

Two of the pros of the P/390 are that it can be used for doing destructive testing of software and the MVS system can be completely recovered within four hours maximum (this includes a complete reload of MVS from the "EXPRESS" and restoring software to other volumes, assuming that Standalone Restore is being used). Furthermore, due to file locking, MVS needs to be shutdown to use an OS/2 program to copy the MVS disks. This brings up an interesting point. MVS can be used to back up all of the MVS known DASD, or OS/2 can be used to back up the files. However, an OS/2 backup can only be restored by OS/2 (not to imply OS/2 can restore MVS backups).

Because the OS/2 DASD support behaves similarly to VM, the emulated 3375/3380 DASD volume(s) only has to be defined as large as needed. For example, a volume called "PAGE00" can be a 3375 or 3380, and defined as having the number of cylinders of space necessary for the PAGESPACE plus one cylinder for the label, VTOC, VVDS, and index (optional). This means that it is not necessary to dedicate a whole 885 cylinders of 3380 single density DASD to a PAGE data set or any other high usage system data set.

Another benefit of this box is that MVS consoles can be LAN-based because the OS/2 3270 device drivers allow 3270 emulation over a LAN. The initial location of the Master Console still has to be on the P/390, so an emulated DFT terminal session under OS/2 is typically used.

The P/390 supports client/server application development directly. A user can logon to TSO or do an RLOGIN from an OS/2 window on the P/390 system and run the client application from another OS/2 window. Or, a user can run on a LAN, with 3270 emulation, and do the client server work all from his/her workstation.

These P/390 systems have the power to make a departmental processing system. We have figured that one of these systems would be able to handle all of the POS processing for a large department store with a VTAM connection to a central site for doing file transfers. Or, it would be possible to run all of payroll on the system, and use NJE for sending the output (checks, ledgers, etc.) to a central site running the "large" printers.

The configuration that we are using here at ACS is 128MB of storage for the S/390 (maximum Main Store available), and 32MB RAM for OS/2. We have the LAN Streamer Token-Ring card which allows us to connect via TCP/IP to our main MVS systems (downstairs from us). This, in turn, allows us to use NDM and NJE with our production MVS environments.

Our P/390 configuration allows us to reload OS/2, Communications Manager/2 1.11, LAN support, and MVS-JES2 SP5.2 from the CD-ROM unit in less than eight hours (from a truly destructive test). If we are only restoring the MVS environment, it takes less than two hours from the CD-ROM.

P/390 FACTS OF LIFE


The P/390 systems use Main Store basically the same way as the Amdahl 5890 systems. E-Store is taken from C-Store. With the 128MB main store systems, the maximum you can assign to E-Store is half of the main store, or 64MB.

The horse power of this system is roughly equivalent to a 4381-11 with only one channel. (It would be interesting to see what the numbers are from some benchmarking systems.)

The MVS CD-ROM that IBM sent to the S/390 Development

Association members who use this system is similar to an old MVS Express system. It must be scattered to take advantage of the I/O "sub-system". This means that someone must play systems programmer so that the MVS scatter gets done correctly. Just remember, if you break something in SYS1.PARMLIB, the recovery process is exactly the same as if you were in a single "CPU" shop. You don't have the ability to connect to another MVS system and share DASD to recover (barring having two P/390s connected to each other).

This machine is overly sensitive to DASD/C-Store tuning compared to a system having multiple channels. So, the better the blocking, the better the system runs. Remember, on the P/390, any I/O you don't do is the least costly. We found that with 16MB E-Store our IPL time was cut by two minutes. We also found that we needed this to avoid some paging, and to let various subsystems make use of Data/Hiperspaces.

Authors' Note: Since this article was written, IBM has been working on various performance improvements and problem resolutions to the P/390 system. An upcoming article will address these issues. 

NaSPA member Steve Thompson has more than 15 years of experience in the computer industry. He currently works for ACS as both a developer and manager of the WYLBUR support group.

Mary Martinelli has seven years of experience in the computer industry and recently joined ACS as a developer in the WYLBUR group. Prior to ACS, she worked for IBM in support of DFP.

©1996 Technical Enterprises, Inc. Reprinted with permission of **Technical Support** magazine. For subscription information, email mbrship@nasp.net or call 414-768-8000, Ext. 116.