



Highlights From the NaSCOM Interconnect

Special thanks to Richard B. ViPond for providing this month's dialogues.

Surf the 'net with NaSCOM! NaSCOM has full-time, leased-line connections running at 115,000 bps with FTP, Telnet, and Web capabilities. NaSCOM is NaSPA's 24-hour, seven-day-a-week online electronic communications link. It allows members to share technical problems, solutions, and experiences with users of other interconnected bulletin board systems (BBSes) worldwide.

***NaSCOM Access Information:** (414) 768-8002 (300-28,800 bps, 8-N-1, V.42). For more information on Advantis X.25 worldwide, email sysop@nascom.com, Internet Telnet/FTP at nascom.com or access NaSPA's home page at <http://www.nascom.com/>.*

***Note:** If you have forgotten your NaSCOM login ID, it can be found on the mailing label of each issue of **Technical Support** magazine.*

Find out more about NaSCOM or how to join NaSPA by calling (414) 768-8000 Ext. 116.

The following excerpts were extracted from dialogues through NaSCOM and other interconnected BBSes.

CONNECT AND DIAL-IN TCP/IP

IBM tech support walked me through an installation this morning; I am posting the results here for the universal benefit of humanity:

PROBLEM: You can't install OS/2 Warp Connect's networking products (network card not recognized or you connect to the network by "modem-only"; installation program hangs).

NOTE: I have an Intel EtherExpress Pro/10 with flash card installed, but not connected to another machine. I also have a modem on COM 1. The following steps work for a modem; I will find out if they also work for a network card.

SOLUTION: Starting from a clean installation of Connect:

1. Run PRODINST from the command line, select "MTPS (Multi-transport protocol services) to install;
2. Reboot;
3. Run "MTPS" from the command line (PRODINST does not create a desktop object for MTPS);
4. If you only have a modem, select "IBM Parallel port adapter" from the list of devices, click add, select TCP/IP from the list of network protocols, click add, exit, let the installation run its course (i.e., push all the right buttons); and
5. Reboot.

At this point, you should have a TCP/IP folder in the drive you selected for installation as well as a desktop folder for the IBM Internet application. You still have to configure the Dialer app and add HOSTS, RESOLV, and SERVICES to /tcip/etc to be able to connect to your provider (i.e., the Dialer works [once it's configured] and will establish

a SLIP/PPP connection, but you can't Telnet, etc., until you have the other files (HOSTS, etc.) in /tcip/etc. If anyone knows exactly which files to add to /tcip/etc, please post or email me.

NOTE: After you select "IBM Parallel port adapter" in step 4, you may get a error dialog. Click "OK" to get rid of the message, then click "configure ["settings"?]" in the MTPS window. This brings up another dialog box, which you can close by clicking "OK".

NOTE 2: I am guessing that the same steps will work for network adapter cards simply by replacing "IBM parallel port etc." with the name of the card in step 4. Since the EtherExpress Pro/10 requires a separate driver from Intel, I'm anxious to try that installation. Anyone with experience installing the Pro/10 card, please let me know.

NOTE 3: IBM tech support was very helpful. On the whole, and with the exception of this glitch, Connect installs very nicely.

Matthew W. Woodruff (woodruff@panix.com)

MRO QUESTIONS

Before I present my questions, here's a little background on my site:

1. My company has a forecasted throughput rate of 10 transactions per second.
2. I have one TOR interconnected with two AORs.
3. OMEGAMON II for CICS V300 is the monitor.
4. 9121-621 is the box.

Questions:

1. How many send and receive sessions should I define in each connection?
2. How do we measure MRO performance in terms of response time and CPU?

Harman Li (harmanli@asionline.net)

Are the AORs clones of one another or are there different applications running in each? In other words, are you using dynamic or static transaction routing? Also, are you using any STARTs in your applications for terminal-associated transactions? If so, what volume? Generally, assuming you are doing ordinary, dynamic transaction routing, I'd define 10 Send sessions from the TOR to each AOR to allow for an AOR failure and one Receive session.

Why 10 and one? I always work on the theory that you define what you need for your normal rate. The one Receive session is needed for shipping back STARTs, etc., from the AOR to the TOR. I would also use class maximums for the transactions and set the maximum to 11.

Why? Because I work on the theory that you define one more 'slot' than you need so that if you ever hit your peak, you've got one transaction primed and ready to use a Send session when it becomes available. Now this all needs to be tempered by the fact that you will actually have more than one AOR, so this is hardly a definitive answer. I'd need to know more about your setup and application to provide more specific advice. Most monitors have the ability to correlate related regions' information.

Jerry Ozaniec (joz@cix.compulink.co.uk)

My MRO environment is as follows:

1. One TOR and two AORs. Each region has its own application to run.
2. In TOR, 80 percent of the transactions are remote. In each AOR, 60 percent are remote. Any one region can route to the other two. That is, every region can initiate routing. All routing is done by function shipping (EXEC CICS START). No static transaction routing is involved.
3. Estimated throughput rate at peak is 10 transactions per second to one region. Therefore, total peak rate is 20 transactions per second for two AORs.
4. We have 600 users connected to the TOR.
5. We have AMXT at 60 and MXT at 90.
6. All regions are CICS/MVS 212.

I assume:

1. Response time collected in TOR alone is the true MRO response time.
2. The CPU time for one MRO transaction is the total of the CPU second collected in the corresponding TOR and AOR for the same transaction. That is: $CPU_{mro} = CPU_{tor} + CPU_{aor}$.
3. The CPU collected does not include the time spent in DB2 attachment.

Harman Li (harmanli@asionline.net)

PARTITIONING 9121-210

We are a small shop in the middle of installing MVS/ESA 5.1 on a 9121-210. We are currently running MVS/XA 2.2.0. The 9121 has 64 MB of memory. Does anyone know how much memory we will have to allocate to HSA? We were thinking that 32 MB would be used for the XA production region, 8 to 12 MB for the HSA region, leaving 20 to 24 MB for the ESA region.

Any insights into LPARs and this new world of partitioning IBM mainframes would be appreciated.

Mike Sheehan (Sheemicq)

We have an almost identical system to yours—an ES/9000 9121-210 with 64 MB of real memory. My HSA is 6 MB. I run two LPARs—both with MVS/ESA. ESA will not run very well with only 20 to 24 MB. However, you should be able to do some basic stuff, but you probably won't be able to bring up very many subsystems or CICS regions without running real slow or getting messages about memory shortages. I can't seem to convince management how much better the systems would run with only 64 MB more memory.

Rich Durkee (Durkrica)

Mike, I don't think your ESA system will run very well (if at all) with only 24 MB. We are running MVS/ESA 4.3.0 on "big iron" (ES/9021 942 and 821). We have a systems programmer's test LPAR that started out with 32 MB. Then, we wanted to be able to have five or 10 TSO users logged on concurrently, so we had to increase it to 48 MB. Next, we wanted to bring up a few CICS regions and a DB2, so we had to increase it to 64 MB, which is where it stands today.

I don't really know the difference between 4.3.0 memory requirements and 5.1 requirements. I also don't know how much you plan to do. However, I'd be surprised if you can get much past IPLing with only 24 MB!

Rich, the HSA on our ES9021/821 is now up to 48 MB! It used to be 32 MB, but we recently added a Coupling Facility which consumed another 16 MB of HSA!

Rick Hogan (Hogarico)

RECONFIGURABLE CHPIDS/LPAR

I see in the IOCP user guide that you can define CHPIDS as reconfigurable meaning that you can reassign them to another LPAR without doing a power on reset. What is involved in doing this? I know that you enter some command on the hardware console but what I want to know is do you have to IPL the two partitions or can it be done on the fly similar to VM?

Dave Barrows (Barrdavu)

Dave, the LPARS do not need to be IPLed. It is done with the CONFIG command on the different systems. Just make sure the devices are in each I/O CONFIG.

Rick Palmer (Palmrlic4)

Dave, first, the CHPIDs need to be defined as reconfigurable in your IOCP. For example:

```
PATH1C CHPID PATH=1C,TYPE=BL,PART=(MVSPROD,REC)
```

This says that CHPID 1C is initially assigned to MVSPROD LPAR, but is REConfigurable, and can be moved to a different LPAR.

In order to move the CHPID to another LPAR, you first have to take it offline on the LPAR which currently owns it. From an MVS console, you would issue the command: CF CHP(1C),OFFLINE. You can achieve the same thing from the processor service console by first issuing a SETLP command to set the LPAR you want to remove the CHPID from, such as: SETLP MVSPROD. Then use the CHPID command to take the CHPID off from that LPAR: CHPID 1C OFF.

Once the CHPID is offline, you can allocate it to another LPAR. From an MVS console on the system you want to allocate the CHPID to you can issue: CF CHP(1C),ONLINE. Or, from the service console, you could set the LPAR you want to allocate the CHPID to; for example: SETLP MVSDEVL, and then use the CHPID command: CHPID 1C ON.

That's basically all there is to it. No need to IPL the systems. Just issue the right commands... unless of course you need to change your IOCP to make the CHPIDs reconfigurable. Then you might have to IPL. But once that's done, you're all set!

If you have ESCON channels, and your processor supports EMIF, you can define channels as SHARED, where one channel can be in use by multiple LPARs at the SAME TIME!

Rick Hogan (Hogarico)

CAPTURING DFHSM HQUERY OUTPUT

We are on MVS/ESA 3.1 with DFP 3.1.1 and TSO 2.X and DFHSM 2.6. I wish to submit a batch TSO job that issues an HQUERY REQUEST command through a CLIST and then to trap the HQUERY output so that I can process it. However, the system insists on sending the output back to my interactive TSO session terminal screen. For example:

Case 1: From session A, submit batch TSO job B, which issues HQUERY. The HQUERY output comes back to A.

Case 2: From A submit B which submits batch TSO job C which issues the HQUERY (hoping the output might go to SYSTERM or SYSTSPRT of B). However, output comes back to A.

Case 3: From A submit B with TYPRUN=HOLD, then log off from A and release B from the console (experiment hoping that the output will be accessible by the user). However, the output seems to get lost.

Can anyone suggest a way to trap the HQUERY output in B? To further compound this, we are trying to recover thousands of backed up

data sets to temporary new names so they can be transmitted to a UNIX system (we are migrating). We want to be able to throttle the recovery process and not overfill the queue of requests waiting to be processed.
(urjlew@uncmvs.oit.unc)

Not knowing exactly what you are querying, I would suggest that you write a program to view the CDS and get your information from there. It will be in a more readable format than the output from a query. You will find all of the records in the CDSes documented in the diagnosis manual (you will probably want the 2.5 version). Also, depending on the query, you could direct your output to a data set and read the information from there.

Paul (prmayhew@mobility.com)

A solution is to use REXX instead of CLIST to issue the commands in the background. REXX offers a facility called OUTTRAP with which it traps the output of a given command into variables which you can EXECIO to a data set.

Gerard Blokdijk (bloksijk@xs4all.nl)

CLIST has SYSOUTTRAP which works the same way as REXX OUTTRAP, using the same facilities. Neither facility can trap the output from that category of HSM commands because the response comes back via TPUT ASID, and IBM designed OUTTRAP to only be able to trap output from PUTLINE, not TPUT, and no consideration was given to TPUT ASID.

Do an HQUERY which results in lots of output to your screen. Next, try to break out of this session. You can't, because the output is not being generated on your TSO session. Caveat: Fortunately, it's been a decade since I worked with HSM, so (a) I may not be remembering everything exactly, and (b) IBM could have fixed some of these things since then.

My distant recollection is that there is an option on some of these commands to have the command put its output in a data set. Unfortunately, you have no way (as far as I remember) to know when it's complete.

(ldw@lindy.stanford.edu)


Here's the answer I received from IBM. It's a case of BAD (trade-marked). But I was hoping for a pleasant surprise; that someone on the list might know of a workaround. It appears that we will have to submit about 100 recover requests then keep checking back, and when 50 are completed (i.e., queue should be down to around 50), release 50 more requests, and continue. I'm still hoping that someone comes through with a brilliant suggestion.

(urjlew@uncmvs.oit.unc)

Use the UNIX system to issue the commands via Telnet and have a script (or "expect") drive them. Then, the UNIX system can stay logged on and can get the responses. (I'd guess you could do something similar with a VTAM secondary program, but it would probably be harder.)
(csysmas@mvs.oac.ucla.edu)

I'm no HSM expert either, but when one way to attack a problem does not pan out, look for other techniques. I don't know exactly what information you are looking for, but I think the original post had something to do with obtaining information from HSM about backed up data sets.

There is a program named ARCUTIL that can write a sequential data set containing information about backed up and migrated data sets. The program can be invoked from a user-written program. Therefore, you can call ARCUTIL, request information about backed up data sets and read the created data set to find the data set information you are interested in. Hope this helps.

Chuck Arney (arney@metronet.com) 



NaSPA member Richard B. ViPond is a consultant with DTA Technical Services, Inc., Oak Brook, Ill. If you have any questions or comments, you can contact Rich through NaSCOM ID Viporicf.

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