

Benchmarking your AS/400 at Rochester

Contributed by Vincent Leveque
Monday, 31 July 1995
Last Updated Monday, 31 July 1995

Issues of capacity and performance are part of the terrifying unknown any AS/400 manager faces. How do you know that your nightly batch jobs will actually complete overnight six months from now? How will the sudden acquisition of another company affect online response time? Will the new package software that looks so tempting actually run well enough in the real world to keep your users happy? What if you would be the vendor's biggest install to date? These questions all involve estimating (or guessing) the unknown, and they have results that are immediately apparent to your user community.

Capacity management and forecasting is part engineering discipline and part black art. Several tools are readily available for performing capacity planning in-house, including Performance Tools and BEST/1. For short-term extrapolations or straightforward increases in workload, these tools may be all you need.

In many environments, these tools are perfectly appropriate—for example, where software is already in use (or comparable sites are readily available) and where incremental growth over a period of months or years is projected, giving plenty of time to correct an original estimate where the patterns of system use are known and not likely to change dramatically.

On Your Benchmark...

Every so often, a system professional needs an industrial-strength performance benchmark. This may involve a mainframe downsizing, a merger of two sizable concerns onto the same platform, or the introduction of radically new software for which the vendor does not have a comparable site. These are situations that combine high risk with a high level of unknowns. For these situations, a performance benchmark at the IBM Rochester laboratories may give you the assurance you need.

The IBM Rochester benchmark facility (referred to as the Customer Benchmark Center, or CBC) is equipped to test the performance capabilities of your planned production environment, simulated using actual hardware. Located in the same campus as IBM's AS/400 manufacturing plant and development lab, the benchmark facility tests the latest and most powerful products to measure their system performance impact.

Resources available at the benchmark facility include IBM staff members who are highly skilled in benchmark methods, AS/400 systems for tests (including many peripheral devices), and some very sophisticated tools for simulating online workloads under a variety of conditions. The quality and breadth of support cannot be equaled.

Get Set...

A complex interactive benchmark at the Rochester labs is not always easy or cheap, however. It can take weeks of planning and costs into five figures to make a benchmark. Benchmark options available from Rochester include simple batch tests, which span a few days and cost approximately \$3,000-5,000, and many other services covering a range of testing needs. Client/server configurations can be tested using a specialized tool called SET (Solutions Evaluation Tool).

For the purposes of this article, we will focus on complex 5250 interactive benchmarks. Obviously, one must be at a critical point in a cost-benefit curve to even consider this an option. Some factors that weigh in favor of using the CBC include the following:

? The need to test advanced technology not yet in widespread use to solve performance problems. For example, a client of mine used the benchmark facility to evaluate whether OptiConnect (a fiber optic distributed database access technology) would yield performance benefits. A variation of this involves the need to use the expert skills of Rochester

staff in reviewing sizing and performance in a controlled environment.

? The need to model interactive online workloads involving many online concurrent users. The Rochester labs have unequalled facilities for this type of evaluation, including specific software for mimicking the keystrokes of user input at a defined rate. Up to 900 users can be simulated using two banks of specially configured PCs, and with a workstation controller version of interactive simulation, up to 1,400 users have been simulated.

? The need to size a target environment with many unknowns. More traditional sizing tools (such as BEST/1) work best if you have a measured baseline using the target software and transaction mix in order to reliably scale up to your planned environment. This baseline is typically your existing installation or, in some cases, a reference site. Often, this baseline does not exist and cannot be reliably simulated using an analytic modeling tool. As an example-you are acquiring a new AS/400 and a new software package, having decided to migrate from a mainframe-based system. AS/400-based modeling tools cannot simply import performance measurements from your legacy mainframe-you need a different approach to this situation.

? The consequences of failing to deliver a specified level of performance are great, and incremental mitigation measures are not feasible. Your firm has acquired another, equally sized firm. You are given the mandate to run the merged company on an AS/400 platform. A preliminary sizing shows that you may be able to do so, though only with the largest commercially available AS/400 processor. Given the risks and costs of failure, a benchmark in this situation would be an exceedingly prudent move.

Typically, expect a benchmark when there is a major platform change, when installing application software where no comparably sized sites exist, or when a system resizing occurs as a result of organizational restructuring (with all the accompanying unknowns).

Now that you have decided that a benchmark would be a prudent and cost effective move, how would you proceed? The process of reserving a time slot at the CBC and conducting a benchmark is well detailed by IBM in the AS/400 Customer Benchmark Center Benchmark Planning Guide. You may wish to obtain a copy of this guide and any other related information as well. Effective planning of your benchmark will be critical to its success.

A benchmark is a test of the capability of technology as well as a test of your own project management capability. Conditions that strain the ability of AS/400 hardware to support your users will also strain your project management capabilities. Consider the discipline required to conduct a successful benchmark as a test of your project management skills and plan accordingly. If you do not currently have a project management and tracking system, now would be a very good time to start.

IBM does apply some very rigorous criteria you must meet before being allowed to use the center. These criteria emphasize your ability to manage the benchmark process and to provide clear, quantitative goals for the benchmark. IBM enforces these criteria in order to ensure that time and money spent performing a benchmark result in valid and useful conclusions. This should be an additional inducement to adopt good project management techniques and to clearly document your goals.

Go!

To kick off the benchmark process, you first need to submit an application to IBM. This is available electronically through IBMLink under "AEFORMS" or by contacting the CBC directly. The form is sent to your Branch Manager or Business Partner, who reviews and approves it and forwards the application to the CBC in Rochester. You should, of course, discuss your situation with your IBM Marketing Rep or Business Partner first to seek assistance in this process.

A benchmark planning session is the next key step. This session is typically held in Rochester and involves key MIS people, representatives of the software vendor, and IBM staff. The planning session reviews the overall benchmark process and looks at objectives, workload definitions, and key success factors (see 1). The planning session is intended to establish the context and scope for the benchmark.

A benchmark planning session is the next key step. This session is typically held in Rochester and involves key MIS people, representatives of the software vendor, and IBM staff. The planning session reviews the overall benchmark process and looks at objectives, workload definitions, and key success factors (see Figure 1). The planning session is intended to establish the context and scope for the benchmark.

The actual work of the benchmark lies ahead. If not already defined, the performance objectives must be clearly stated and approved by all concerned parties. The statement should specify definite, quantitative goals, as directly measured by the IBM Performance Tools. They are not broad statements of business requirements (such as "enter 5,000 orders within an 8-hour window"), nor is "determine maximum number of transactions an AS/400 model 320 can handle" the right type of goal.

The type of specific benchmark objective you need to define is of the nature of "demonstrate that an AS/400 model 320 is able to process 500 payroll entry transactions per hour with a maximum average online response time of one second." Broader statements of the business objectives are, of course, important as a first step, but these need to be taken further and formulated as a test objective.

Scenarios of business processes and associated AS/400 workloads must be defined. What processing will take place during a run at the CBC? Will you need to run both online and batch processes simultaneously? Will you need to mix different types of processing? What will be the key-to-think time? This requires understanding the users' business processes, the types of software functions a group of users will execute in a typical day, and the expected pacing and overall transaction volume.

Hardware resources required at the CBC will need to be finalized. What model(s) of processor will you need to test? How much DASD will be required? Think about requesting an automated tape library (ATL), especially if you have large amounts of test data. The benchmark process itself will involve multiple saves and restores-being able to run these unattended will be a welcome convenience.

Specific test data will need to be defined and generated. The test data includes parameter and table files required by your software (such as a chart of accounts files), batch transaction data to feed batch processes, and online scripts to define the entry patterns of users. Many times, your existing production data can be used for the benchmark, saving effort in test data generation. However, do not under-estimate the challenge of creating very large files of simulated data for batch processes. Generating good test data that reflects the performance characteristics of your production environment requires thought and work. While some specific benchmarks may be able to use existing data, the generation of test data specifically for a benchmark may take a couple of weeks. (For a quick rundown of the terms you need to know at this point, see the "Terminology" sidebar.)

Online scripts will need to be defined and some initial coding of these scripts will need to be conducted. One of the strengths of the CBC is the availability of the Performance Evaluation Tool Environment (PETE) tool. PETE is a keystroke recorder/playback tool designed for the specific purposes of benchmarking.

PETE allows a single micro-channel PS/2 to simulate up to 24 online users and permits varying important parameters, such as key-to-think time. PETE permits cycling through a transaction update, using a different case for each update. A certain element of randomness may be programmed into the scripts to more accurately reproduce the use patterns of end users. PETE scripts can be produced at the customer site. Due to the requirements of the PETE software, the debugging and finalization of PETE programs must be done at Rochester. See the "Meet PETE" sidebar for a more thorough examination of PETE.

Earning a Benchmark of Excellence

The specific plan of activity you will undertake at your Rochester benchmark must be defined. This will include time for PETE training, script entry, and finalizing test data. Each specific benchmark run should be described, including any data reset procedures before and after, and the specific performance measurements the run will grant. This plan, also referred to as a run calendar, will map your activities day by day. The CBC staff will assist in the initial planning of the run calendar at the planning session, and will help with its finalization at or prior to the Readiness Review.

All of the above steps need to be substantially complete before a pre-benchmark Readiness Review is conducted. Your benchmark workplans will be submitted to the CBC prior to the Readiness Review for Rochester staff to assess. The Readiness Review is typically conducted as a conference call in which you discuss critical items required to conduct a successful benchmark.

Several specific items must be satisfactorily completed before the benchmark may be conducted. These are taken from the Benchmark Readiness Checklist:

? All data and program conversion needed to execute the environment at the Benchmark Center are complete.

? All applications have been tested and debugged. (Note: a scripting tool such as PETE is very intolerant of application failures!)

? All legal agreements necessary to use the software at the CBC have been signed.

? If PETE is required, all run scripts have been specified in detail.

? A detailed benchmark plan, scheduling all activities to be done at the CBC, has been produced and is realistic.

? The team required for the benchmark is adequately staffed and prepared to assume its responsibilities.

A satisfactory Readiness Review means you can purchase those tickets to Rochester. Expect to spend two to three weeks onsite conducting the benchmark activities. Your first several days will be spent restoring data files, installing software, and performing other configuration activities. In parallel with these activities will be the final development and debugging of the PETE scripts.

With the environment set up, several test runs will be conducted. This is both for final debugging of the PETE scripts and for tuning the systems for optimum performance. Any number of production runs may follow for gathering the data necessary to support (or contradict) your objectives. Despite all the careful planning, you will find it necessary to improvise. Some runs may be canceled, and others you did not anticipate will become necessary. Data from each run should be reviewed on conclusion of the run, and any problems or ambiguities noted and corrected. To see a sample of other customers' benchmark experiences, see 2.

With the environment set up, several test runs will be conducted. This is both for final debugging of the PETE scripts and for tuning the systems for optimum performance. Any number of production runs may follow for gathering the data necessary to support (or contradict) your objectives. Despite all the careful planning, you will find it necessary to improvise. Some runs may be canceled, and others you did not anticipate will become necessary. Data from each run should be reviewed on conclusion of the run, and any problems or ambiguities noted and corrected. To see a sample of other customers' benchmark experiences, see Figure 2.

You May Approach the Benchmark

After completing the runs, you will spend a day or two at the CBC, reviewing your results with IBM staff and formulating your report. By supporting or contradicting your test objectives, the benchmark will contain important information regarding your business objectives. What model AS/400 will it take to run the new software? Given the length of the nightly batch process, what window will you have for backup and how will this affect your recovery strategy? What technical features will you need in order to get the performance required? How will your system run on multiple AS/400s? Any number of these can be established at a benchmark, and the risks and costs of your cutting edge AS/400 strategy reduced substantially.

Performance data from the benchmark may also be saved to magnetic media and used for ongoing analysis after the conclusion of the benchmark. Benchmark data can be a good baseline for tools such as BEST/1, which can permit you to

analytically model transaction loads and system characteristics you were not able to test at Rochester.

Leaving the Rochester CBC with a well-thought-out results document is not the end of the benchmark process. By conducting a benchmark, you have learned a lot about your project management process, the nature of your software, and how your users will interact with it. You have taken the first step toward an ongoing system capacity plan, something all large AS/400 shops should have.

A benchmark provides many valuable lessons beyond the capacity and performance conclusions. Applying these lessons will provide additional value from the benchmark and help make your AS/400 project a success.

Vincent LeVeque is a manager in KPMG Peat Marwick's Information Risk Management practice. He has more than 10 years experience with the IBM midrange, starting as a programmer with the S/38. He can be reached at 213-955-8921 or via Internet E-mail at pp002390@interramp.com.

Mr. LeVeque wrote this article with the generous assistance of the IBM Customer Benchmark Center and would like to thank them for their help.

References

AS/400 Customer Benchmark Center Benchmark Planning Guide (this is an internal IBM document with no publication number assigned).

Performance Benchmarking for the AS/400 (Redbook GG24-4030, CD-ROM 66244030).

Benchmarking your AS/400 at Rochester
Terminology

Confusion over terminology can slow down your benchmark planning process and result in long discussions lacking in content and direction. Common terms used in the benchmark process often have meanings different from the same terms defined in IBM's BEST/1 documentation.

AS/400 Transaction For a traditional 5250 application, any processing that occurs between the presses of the Enter key.

Business Transaction A discrete operation within a scenario (what your users might understand as a transaction).

Run A discrete test at the Rochester CBC, typically lasting half a day.

Scenario A specific business process performed by a specific user group.

Script The file used by the PETE software to simulate an end user's online activities.

Workload The processing that takes place on the AS/400 during a run.

Benchmarking your AS/400 at Rochester
Meet PETE

PETE is a keystroke recording/playback tool specifically designed for performance testing. PETE simulates "plain vanilla" 5250 users on the system under test.

The simulators are PS/2 workstations acting as 5250 workstations to the AS/400. The simulators actually execute the keystrokes sent to the AS/400. Each workstation is connected by twinaxial cable to the AS/400. PETE creates no CPU, memory, or disk overhead on the AS/400. As far as the AS/400 is concerned, the attached workstations could well be 5250 terminals with users doing their work.

Key features of PETE include the ability to capture and replay keystrokes and the ability to control keying rates and key-to-think time. PETE provides a programmable scripting language to enable tight control over the characteristics of a script. PETE simulates the realism of the actual user environment and provides the ability to repeat specific workloads.

PETE additionally permits measuring end-to-end response time at the workstation. PETE measures the actual response time at the workstation, rather than at the workstation controller, as the Performance Tools do. The PETE response time reporting capability additionally gives finer specificity of interactive response time. PETE actually permits measurement of the response time resulting from a specific Enter or function key.

Benchmarking your AS/400 at Rochester

Figure 1: Summary of Steps to a Successful Benchmark

1. Discuss your plans to benchmark with your IBM representative and your software vendors.
2. Ensure the nomination form is completed.
3. Attend a benchmark planning session to review the specifics of your benchmark.
4. Define your benchmark team and members' specific roles.
6. Define specific, measurable objectives for the benchmark.
7. Define the specific hardware your benchmark will require.
8. Define the specific runs you will need to execute to test your objectives.
9. Define business-critical scenarios.
10. Prepare test data.
11. Specify online scripts.
12. Obtain copyright agreements and sign any necessary disclosure agreements for proprietary software or any other non-IBM product used in the benchmark.
13. Conduct Readiness Review.
14. Conduct benchmark.
15. Develop and present benchmark findings and recommendations.
16. Implement long-term capacity planning process to ensure adequate system capacity is available to support your business.

Benchmarking your AS/400 at Rochester
Figure 2: Some Actual Customer Experiences

Here is a sampling of some recent benchmarks at the CBC, giving some idea of situations in which IBM customers seek to use benchmarking:

Name Type Duration Comments

Customer A Interactive & Batch 2 weeks Software release-to-release testing.

Customer B Client/Server 3 weeks New application stress test.

Customer C Interactive 2 weeks Determine performance limits and architectural limits of application software.

Customer D Batch 4 days Large growth due to acquisition/merger.

Customer E Interactive 2 weeks Testing of OptiConnect.

Customer F Interactive & Batch 3 weeks Migration from mainframe environment.

Customer G Interactive 2 weeks Benchmark of AS/400 performance, to compare with prior benchmark of competitive equipment.

Customer H Batch 5 days Consolidating multiple systems into one.