

October 26, 2001 01:01 PM PST

Waiting for Windows XP

P.J. Connolly with Randall C. Kennedy

HOPELESS OPTIMISM must be a fundamental part of human nature, because we want to believe that new operating systems truly represent an improvement on their predecessors. It's easy to point to certain features in a new OS as examples of progress, but end-users often find that a new OS performs like molasses compared to the version they were using. As a result, CTOs wanting to capitalize on the benefits of a new OS may find that new hardware investments are necessary -- and expensive -- requirements.

Unfortunately, Microsoft's Windows XP appears to be maintaining that tradition, as indicated by results of independent testing performed by CSA Research and confirmed by our work in the InfoWorld Test Center. Our tests of the multitasking capabilities of Windows XP and Windows 2000 demonstrated that under the same heavy load on identical hardware, Windows 2000 significantly outperformed Windows XP. In the most extreme scenario, our Windows XP system took nearly twice as long to complete a workload as did the Windows 2000 client. Our testing also suggests that companies determined to deploy Windows XP should consider ordering desktop systems with dual CPUs to get the most out of the new OS.

XP by the numbers

In every test we performed on systems with a single CPU, OfficeBench ran in less time under Windows 2000 than under Windows XP. The differences ranged from slight to dramatic, depending on the hardware configuration, but XP was always slower. For example, our initial baselining using Office XP and an optimized UI on the Pentium 4 system indicated a lag of a modest 11 percent, but things went downhill from there. The baseline reading for the Pentium III system showed that under Windows XP and Office XP, OfficeBench took 27 percent more time to execute than under Windows 2000 and Office XP.

Generally, Windows XP proved increasingly slower than Windows 2000 as load increased, with a few rare exceptions. For example, in the first multitasking scenario (scenario 1), using light database, messaging, and multimedia workloads, we got mixed results using Office XP on our Pentium 4 client. Under the default UI, Windows XP with Office XP narrowed the performance gap to 24 percent, compared to a gap of 35 percent in the baseline scenario.

Except for a few instances, Windows XP increasingly ate the dust of Windows 2000 as load ramped up, regardless of machine specs or Office version. When the Pentium 4 client with Office XP was tested, script execution generally took between a quarter and a third longer with Windows XP as with Windows 2000, and as much as half again as long with the heaviest load and a stock UI. The Pentium III client fared even worse. Running Windows XP with our heaviest workload and the default UI raised script execution time to more than twice that of Windows 2000. Optimizing the UI helped Windows XP to narrow that gap, taking 1.6 times as long as Windows 2000 to process the workload.

Windows XP stayed closer to Windows 2000 when we tested the software on fast hardware using an optimized UI and Office 2000 instead of Office XP, but speed differences were still as great as 18 percent on the dual-CPU Pentium III and 25 percent on the Pentium 4. Not surprisingly, Windows XP posted the slowest times and the greatest deltas using the stock UI and Office XP on our Pentium III client, taking 58 percent longer than Windows 2000 needed to perform the relatively undemanding workload.

SMP makes the difference

Our tests on a dual-CPU system indicate that both Windows XP and Windows 2000 run better on an SMP (symmetric multiprocessing) configuration with relatively slow CPUs than on a single-CPU system with a screamingly fast processor. As we added more and more load, the benefits of a dual-processor configuration became more apparent. Both OSes (using Office 2000 and optimized UIs)

handled the heaviest workload (scenario 3) nearly 40 percent faster on the SMP client machine than on the single-CPU Pentium 4.

Finally, our cross-generational testing, which measured the performance of Windows XP and Office XP directly against that of Windows 2000 and Office 2000, found that once again, newer means slower. In every one of our scenarios the combination of Windows XP and Office XP took noticeably longer -- from 35 percent to 68 percent longer -- to complete the script than Windows 2000 and Office 2000.

Overall we are quite disappointed with Windows XP's ability to pull serious weight when compared to Windows 2000. We are not certain where the problem lies. Our follow-up testing indicates that the additional database and multimedia workloads are breaking the proverbial camel's back. Microsoft claims it's been unable to duplicate our results, but hasn't supplied us with a better explanation or identified a major flaw in our testing. Whatever the cause, until the problem behind Windows XP performance is resolved, we can't recommend Windows XP as a client for serious database crunching.

In fact, until 2GHz desktop PCs become commonplace, we have a hard time recommending widespread adoption of Windows XP at all. Granted, it appears that for light-duty service on the newest hardware, Windows XP with Office XP is an acceptable choice -- if an 11 percent performance hit, or 53 minutes added to an 8-hour day, is acceptable. But beware of this combination in more demanding environments, whether the workload is greater or the equipment is older.

Barring the need for Windows XP-specific features, such as the remote-control and management options, IT departments should take advantage of license downgrade provisions and continue to press forward with Windows 2000 deployments until the installed hardware base catches up with XP. Shops lured by XP features should weigh their options carefully. In many cases, these features may not be compelling enough to justify saddling your end-users with a slower OS. Although differences between Windows XP and Windows 2000 can be measured in seconds, what business can afford to put a 20 percent or greater bite on worker productivity?

How we tested

For our evaluation of Windows XP's performance, we used CSA Research's Benchmark Studio Professional, a suite of benchmark tests that are especially well-suited for evaluating the performance of both PC hardware and the Windows environment.

At our direction, CSA Research carried out four sets of tests to compare Windows XP and Windows 2000 performance: one set using Microsoft Office XP on a PC equipped with a 1.5GHz Pentium 4 CPU; a second set using Office XP on a PC with a 733MHz Pentium III CPU; a third set using Microsoft Office 2000 on a PC with two 1GHz Pentium III CPUs and on a PC with a Pentium 4; and finally, a fourth set on a 1.5GHz Pentium 4 PC running Office 2000 over Windows 2000 and Office XP over Windows XP.

Each set of tests involved measuring the execution times of four different workload scenarios, dubbed baseline, scenario 1, scenario 2, and scenario 3. First we used Benchmark Studio's OfficeBench 3.0, a linear script that performs tasks in Microsoft Word, Excel, PowerPoint, and Internet Explorer, to provide a baseline for the system's performance under each OS. We then added increasingly heavy workload simulations to OfficeBench to create the other three scenarios. Performance was measured under both default and optimized versions of the Windows user interface. For optimized testing, the animation and font-smoothing features of both OSes were disabled, as was Windows XP's System Restore feature.

Senior Analyst P.J. Connolly (pj_connolly@infoworld.com) covers groupware, messaging, networking, operating systems, and security for the Test Center. Randall C. Kennedy (rck@csaresearch.com) is the Director of Research and co-founder of Competitive Systems Analysis.