

Saturday, March 3. 2007

About reading benchmarks

It's a rule I tend to preach everytime I have to fight against misinterpreted benchmarks: Capacity of the system tells you nothing. It's important how long every user has to wait. A system capable to handle 1000 user with reaction times around one second and a system capable to handle 1000 users with reaction times around 3 seconds are not the same. But, most benchmarks define a maximum reaction time, for example: The maximum capacity is reached, when the answer of the system takes 4 seconds. When you want to make some benchmarking statements, you load the system until you reach the maximum reaction time. Fits the rule of the benchmark? Yes, definitely. Is it a realistic approach for a successful project: No!

Paul Murphy gives an interesting example in *What users want*: Except that when you think about user perception, the Sun's 0.7 second response time would make it seem fast to users while IBM's 3.056 seconds would make them wait long enough to become conscious that they were waiting - and that perceptual difference doesn't count for the Sun so much as it counts against the IBM, because its apparent slowness will frustrate users, thereby dooming your project. So: Even when the IBM is in striking distance to the Sun system benchmarkwise, the systems play in different classes: Fast or dog slow. 2.3 seconds per transaction, 8 hours a day, 5 days a week, roundabout 48 weeks a year are the difference between satisfied users and a lynch mob in front of the IT department.

In the case you don't believe me: I've worked for a unified messaging provider. For every additional second, our systems needed to react our help desk got more work to answer the complaints.

Posted by Joerg Moellenkamp in English, The IT Business at 20:41