

Re: Seeing VERSIONINFO under Vista?

Source: <http://www.tech-archive.net/Archive/VC/microsoft.public.vc.mfc/2007-05/msg02239.html>

- *From:* Daniel James <wastebasket@xxxxxxxxxxxxxxxxxxx>
 - *Date:* Wed, 30 May 2007 12:16:49 +0100
-

In article news:<63hp53pjml817qj8utalm3gsf83037t7mg@xxxxxxx>, Joseph M. Newcomer wrote:

Not at all. This very interesting side-discussion started when you said "Don't blame IBM". I'm just explaining why I do blame IBM.

For failing to have a crystal ball? That's what I'm objecting to!

No, not for failing to have a crystal ball — IBM is (and was) a big business. It has marketing and research departments that are surely equal to any but the best of crystal balls. The problem was that IBM — partly through arrogance and complacency, and partly because of a lack of coherent strategy throughout the organization — didn't make use of the information that was available.

It's not that they didn't *have* a crystal ball, just that they didn't look into it.

"Blame" implies a deliberate decision or set of decisions to build something that had self-limiting problems.

I disagree. Blame can be ascribed for negligence, not just for deliberate acts.

However, I'm talking about "blame" because that's the word you used. When I say that I "blame" IBM I don't mean to suggest that I, personally, have suffered in any material way from their actions. I suppose I'm just saying that a number of decisions made by various departments at IBM at and just after the time when the first IBM PC was being designed and delivered have shaped the landscape of computing for more than two decades ... and the responsibility for that must lie with IBM.

There was no "different department" issue; there was ONE SMALL GROUP IN BOCA RATON with a well-defined goal to produce an affordable personal computer ...

Re: Seeing VERSIONINFO under Vista?

The "different departments" to which I alluded were the strategic planners and marketing people who failed to pay proper attention to the PC because their heads were up their corporate mainframes rather than looking at how the market was actually moving. They didn't need a crystal ball to see that PCs were going to be important, they just needed to open their eyes.

Whether they would have had the courage to bring a PC to market at all, if they had looked at the market and understood what they saw I don't know -- but it's pointless to speculate on that. If IBM hadn't made the PC someone else would have, and history would be different.

... just assigning the blame where I see it as being due.

Due to whom? This is where I think you are completely wrong. The concept that there is even someone to "blame" is already a flawed concept, so anything based on that fundamentally flawed assumption is nonsensical.

Lighten up, Joe! You're reading entirely too much into "blame".

This whole thing started because you wrote "Don't blame IBM" -- your remark established the ground rule that the concept of "blame" is meaningful for the purposes of this discussion. Don't worry, no-one here is going to nuke Poughkeepsie as a result of it!

Even so, when the PC was new many people asked why the 8088 was run at "only" 4.77MHz (IIRC intel sold it as a 6MHz part).

Oh, the reason for that is simple: the only crystal oscillator that was available at the price and quantity required was one used for TVs, which used a 4.77MHz oscillator!

Understood. I wasn't really looking for a reason, just pointing out that 4.77MHz wasn't really such a big deal.

Clueless, certainly. The only ineptness was in not having an open-collector/open drain interrupt line so that interrupts could not be shared. But in that sense, nobody else of that era had a clue either, so they were not alone.

I meant inept at marketing and corporate strategy -- or "crystal ball gazing" if you prefer. <smile>

Re: Seeing VERSIONINFO under Vista?

Absolutely. Motorola clearly screwed up; had they bid the 68000 to spec, they would be the semiconductor powerhouse today, and Intel might not even exist.

OK, I blame IBM *and* Motorola. Happier?

Note that comparing the 8088 to a Core2 or an XScale ARM is sort of silly, since those are MUCH later architectures ...

You started it. I was responding to your:

What, by the way, is "crappy" about the chip? Right now, Intel produces the fastest consumer chip on the planet.

I agree that intel NOW produces some very good chips, but we had been talking about the 8088 ... which I still maintain should be judged as "crappy" by any technical yardstick of its day.

The PC XT/370 was ...

[snip]

Yes, I knew about that.

I thought you probably would have done.

One of the problems in comparing PCs to mainframes is that mainframes tend to run massively-I/O-bound code, and until recently NO x86-class architecture could come within a couple orders of magnitude of mainframe I/O bandwidth. The NUMA x64 machines, however, have essentially destroyed the advantage mainframes held there, and a machine equipped with a couple \$3000 Adaptec cards is now in the same performance range as many mainframes, but is a fraction of their cost. So having a 68K-speed CPU that ran mainframe programs made no sense, because a single-bus PC (even a Microchannel architecture) simply did not have the I/O bandwidth to run most mainframe apps.

A number of interesting points, there.

Yes, mainframes have much better i/o handling than PCs ... but N PCs handling one task each can outperform 1 mainframe running N tasks with N sets of i/o (for

Re: Seeing VERSIONINFO under Vista?

Re: Seeing VERSIONINFO under Vista?

suitably large values of N). Where the break-even point comes depends a lot on the size of the mainframe and the nature of the tasks being run — my old University bought an ICL 2900 mainframe (the government insisted that they buy British, so there wasn't much choice) on the understanding that it could support over 100 users working interactively at one time. It turned out that ICL's understanding of "working interactively" meant "using terminals attached to one monolithic application" — as in data entry tasks) while the University's requirement was 100 users each running sizeable programs in separate VMs. Understandably, the system crawled. Eventually the government and ICL coughed up a million or so pounds for the University to buy something else to handle the heavy interactive use, and the University got a couple of VAX-11/780s. While the VAXes couldn't crunch numbers as fast as the 2900 they gave a much more responsive service when 100 people were online at once, and half of them were running compilers.

As for a desktop machine running mainframe software making no sense ... it depends what the intended use is. Most mainframe software only has high i/o bandwidth requirements because it is doing i/o for a lot of users at once. A desktop running mainframe software is likely to be being used by just one user and won't have the bandwidth requirements. Something like the XT/370 strikes me as ideal for a small consultancy offering software development and maintenance services for mainframe code, but not having the need to process the volumes of data that a live installation (on a real mainframe) would have to handle. It may also be useful for small companies that want to process the same sort of data in the same sort of way and with the same software as much larger companies, but who don't have the volume of data (or the capital) to buy their own mainframes (in this case it might be cheaper to buy time on a real mainframe, but the cost dynamics probably vary from case to case).

And a FORTRAN programmer really doesn't care what architecture is executing FORTRAN.

As long as all interaction with the real world is abstracted by the runtime libraries that's probably true, but as soon as he has to use some platform-specific API he's going to need to test it on the right platform.

Supercomputers do have their uses, but there aren't very many of them. They're still expensive, and you can really only justify them if you have massive budgets and massive computing need. Otherwise, arrays of PCs will do the job.

Well, it's certainly true that you can't justify them if you can't afford them. The interesting question comes when you /can/ afford it, and is about the position of the break-even point between supercomputers and clusters of big workstations — and that all depends on the sort of work you need them to do.

IBM still offer mainframes, and one of the purposes that they offer them for is

Re: Seeing VERSIONINFO under Vista?

running large numbers of VMs hosting alien OSes (such as linux). In some cases the mainframe can work out more economical than discrete workstations (especially when the workstations are typically not all used at once, so load-balancing is possible) and in others the workstations work out cheaper.

Then again: A friend of mine works for IBM, and a lot of his work is at the ECMWF where they have clusters of supercomputers ... I don't see that lot being replaced by a couple of PCs and bit of wet string any time soon!

<http://www.ecmwf.int/services/computing/overview/index.html>

—

Cheers,

Daniel.

.