



STATS Investment for HP and Intel: billions of dollars and more than 10 years

Itanium: “Itanic” or full steam ahead?

↘ In June 1994, Intel and Hewlett-Packard publicly unveiled their relationship and CPU-development plans for a compelling 32- to 64-bit transition. HP was seeking a successor for its aging PA-RISC line, and Intel was aiming for an architecture that would eventually replace the full gamut of its x86 products—simultaneously and conveniently obsoleting competitors’ x86-compatible CPUs in the process.

The EPIC (explicitly-parallel-instruction-computing) moniker appeared in 1997, reflecting the companies’ belief that intelligent code and compiler design could yield instruction-level parallelism and, consequently, the extraction of high IPC (instructions-per-clock) ratings. The first Itanium processor, Merced, finally ramped into production in May 2001. Its floating-point-calculation speed on native IA-64 code was unparalleled, but its integer performance didn’t dramatically exceed that of similarly clocked x86 counterparts, and it ran x86 instructions at roughly one-eighth native speed. The Itanium 2 family improved on the initial Merced design in a number of areas.

Development delays and competitive offerings (specifically, the 64-bit extensions and multicore extrapolations of AMD’s Athlon) have hindered the Itanium family. In response to AMD’s threat, Intel added 64-bit, Hyper-Threading and multicore features to its own Xeon family, thereby competing with itself and pigeonholing Itanium into scientific-workstation, large-database-server, and supercomputer applications. However, company executives recently trumpeted that “top server vendors,” support the CPU, along with “six of the eight top” mainframe vendors.—by Brian Dipert