

## Intel VT vs. AMD Pacifica

The two chipmakers are building virtualization support right into the CPU. Will it virtualize Microsoft's monopoly? by Andy Dornan

**+** Comparing CPUs used to be relatively simple. Sophisticated buyers always knew to look beyond a chip's megahertz rating, but ultimately it was still about speed. The math coprocessors, multimedia extensions, and second-level caches were all in the service of crunching through code as fast as possible.

Not anymore—at least, not if Intel has any say in the matter. Instead of just trying to make its processors faster, it's adding functionality that can't be quantified in gigaflops. Intel hopes customers will do the same, looking beyond number-crunching performance to focus on features such as security, manageability, and power consumption.

Intel's stance could be seen as an attempt at distraction. Most independent tests put AMD at the front of the x86 speed race, so Intel's only hope of retaining market share is to make people look at something else. However, AMD is also going beyond pure performance. It has an equivalent to most of the new capabilities that Intel is promoting, and in some cases AMD's versions are more advanced.

Intel calls its new features "star technologies" (\*Ts, see table on page 52). Of the five announced so far, one is really just a rebranding of the 64-bit extensions it licensed from AMD. Three more are dependent on a fifth, Virtualization Technology (VT). Previously known

under the code names Vanderpool and Silverdale, VT is set to ship by the end of 2005. AMD's equivalent is the Pacifica Secure Virtual Machine (SVM), slated for early 2006. Both build virtualization support into hardware.

From the vendors' marketing slides, VT and Pacifica look quite different. Intel is promoting VT as a security and management architecture for laptops, while AMD is selling Pacifica as a way to consolidate servers in the data center. However, this is just spin, representing the

companies' strengths in other areas: The Pentium M has helped Intel consolidate its hold on the mobile market, while servers are increasingly turning to AMD's Opteron. The underlying technologies are almost identical and will be included across the full range of PCs within a year.

### FIVE-RING CIRCUS

Building virtualization into hardware sounds contradictory. The whole point of virtualization has traditionally been to avoid hardware, simulating it in software. Why crawl around in the data center every time a Unix server needs a memory upgrade when an IBM mainframe can provision virtual Linux instances automatically? Why keep that old Windows 95 box around when a modern XP workstation can virtualize legacy DOS applications in the idle time between key presses?

The difficult part is that true virtualization requires each Virtual Machine (VM) to simulate a real one exactly. This is a problem with the x86 architecture because OS kernels expect direct control of the CPU. In programming parlance, they run at "Ring 0," the deepest level of access, with the most functionality. A traditional x86 chip can't run a virtualized OS at Ring 0 because that's needed for the hypervisor, the master OS that hosts all the VMs.


The x86 architecture provides three more rings, each with progressively less functionality. For stability, modern OSs restrict applications to the least functional, Ring

### > Upshot

**Claim:** CPU extensions simplify the creation of VMs and other management operations, making server virtualization simpler and allowing an entire client OS to be run in a secure sandbox, separate from management tools.

**Context:** The AMD and Intel architectures were originally driven by Microsoft's Palladium initiative. Although Palladium was held up, the chipmakers pressed ahead, and Intel has been working closely with security vendors and the open-source community.

**Credibility:** Intel and AMD have a history of delivering on their promises. But hardware isn't much use unless there's software to run on it. Watch Xen and VMware closely.



Intel and AMD are moving ever more PC features onto the CPU.

# product roadmap

3. (This is why Windows XP is so much more reliable than its DOS-based predecessors, which let applications access Ring 0.) So the obvious approach to virtualization is to run the guest OS in one of the two vacant rings.

Unfortunately, some x86 machine code instructions only work at Ring 0. To run properly in higher rings, the OS must be rewritten (or at least recompiled) to avoid those instructions, an approach known as paravirtualization. This is popular in the Linux world—IBM uses a similar technique to run Linux clusters on a mainframe—but it takes work on the part of programmers, and it requires that the OS's source code be available.

## DRILLING DOWN

To run an unmodified OS outside Ring 0, the hypervisor must intercept the forbidden instructions and emulate them. This is the approach taken by VMware, as well as by Windows XP's own emulation of DOS. The disadvantage is that emulation can use a lot of computing power—not a problem for the occasional application written to run on DOS-era hardware, but a significant one for an entire OS that takes full advantage of a modern PC.

To assist virtualization, VT and Pacifica insert a new privilege level beneath Ring 0. Both add nine new machine code instructions that only work at “Ring -1,” intended to be used by the hypervisor. This way the OS doesn't have to be modified, and the performance penalty from emulation is reduced. However, it isn't eliminated completely: Each OS must be convinced that it alone has access to the machine's memory and I/O buses, while the hypervisor juggles access to the real devices to ensure that programs and data can't leak between OSs.

Memory has been partly virtualized since the 386 in the sense that the OS and a hardware memory controller allocate RAM (or disk space if the RAM runs out) between applications. AMD has a definite advantage here. Its CPUs include the memory controller, so Pacifica can simply re-use that. In contrast, Intel's CPUs off-load memory control to a separate chip that doesn't support VT, meaning the hypervisor must take on more of the memory management work. Intel's memory controller will eventually be able to use VT, but not until it's brought into the CPU, expected to happen in 2007.

At present, I/O virtualization requires that drivers run on the hypervisor, which then presents virtual drivers to the guest OSs. Future versions of Pacifica and VT will eliminate the drivers from the hypervisor, allowing guest OS drivers to communicate with the hardware directly. However, this will require support from all PCI devices and so needs to be built into the PCI specification. The PCI-SIG began work on this in June, but has no timetable for a final standard.

## IT INSIDE

Microsoft originally planned to support VT and Pacifica through Palladium, a new security architecture aimed mainly at consumer Digital Rights Management (DRM). The principle was that a new, more secure OS would run parallel to Windows and be invoked whenever extra security was wanted. For example, a media player on the secure OS would be able to play content that couldn't be captured by an application on regular Windows.

Microsoft demonstrated the technology in early alpha versions of Windows Vista, then called Longhorn.

## How Intel's "Star Technologies" Compare to AMD :::

Intel *T	AMD equivalent	Purpose	Winner
Virtualization Technology (VT-x, "Vanderpool")	Pacifica Secure Virtual Machine (SVM) architecture	Extra assembly-language instructions that assist virtualization, enabling a machine to run several OSs at once	AMD, at least in theory. Pacifica's built-in memory management should give it a clear edge. In practice, Intel's work on software could still pay off.
Active Management Technology (AMT)	None	Remote control of a PC through a hardware management agent embedded in the NIC	Intel. AMD has nothing like it even announced, whereas AMT is already shipping.
Extended Memory 64 Technology (EM64T)	x86-64 (also called AMD64)	64-bit extensions, fully backward-compatible with all 32-bit x86 programs	AMD. It invented the 64-bit x86 instruction set, and Intel is still trying to catch up. AMD powers the highest-performance 64-bit x86 servers and the only 64-bit laptops.
Embedded IT (EIT) architecture	Pacifica SVM combined with the TPM	Management software that runs in its own virtualized OS, secured by the TPM and, in Intel's case, integrated with AMT management hardware	Intel. Though AMD's virtualization hardware is better, Intel also has AMT and has announced tie-ups with management vendors such as Computer Associates and Tivoli.
LaGrande Technology (LT, may be renamed)	Presidio Security Technology (may be renamed)	Encrypted I/O, and integration between the TPM and hardware virtualization support	None. Neither exists yet, and encrypted I/O will be possible using the TPM alone—without the need for proprietary Intel or AMD technology.

# product roadmap

From the user's perspective, applications running on the second, secure OS appeared to run in Windows with highlighted borders. However, the extra OS wasn't included in later beta versions, and the plan has since been put on hold. Microsoft has announced a hypervisor for Windows Server 2007, but that will ship later in 2007 (or perhaps 2008), not with the OS itself, and may require an additional licensing fee.

Absent Microsoft, Intel is still promoting VT as a desktop (and laptop) security technology, but focused on enterprise management. The slogan is "Embedded IT Architecture"—a VM dedicated to anti-virus, anti-spyware, or backup software (see figure). In most cases, this software would be controlled remotely by the IT department, invisible to the user. Another VM can run Windows and all its applications normally—except that a malicious program or user wouldn't be able to disable the security software.

The same thing will be possible with Pacifica, though Intel's Active Management Technology (AMT) gives Intel an edge in embedded IT. AMT places a hardware management agent inside the NIC that can perform basic management tasks even when the CPU is switched off. For example, it could reboot a crashed PC or install a new hypervisor.

## HYPE VISION

The big issue for both VT and Pacifica is software support. The management VM will probably run a stripped-

down version of Linux, simply because it costs nothing and is easy for vendors to customize. However, there's no reason in principle that it couldn't run a hardened version of Windows or any other x86 OS. And the possibilities aren't mutually exclusive.

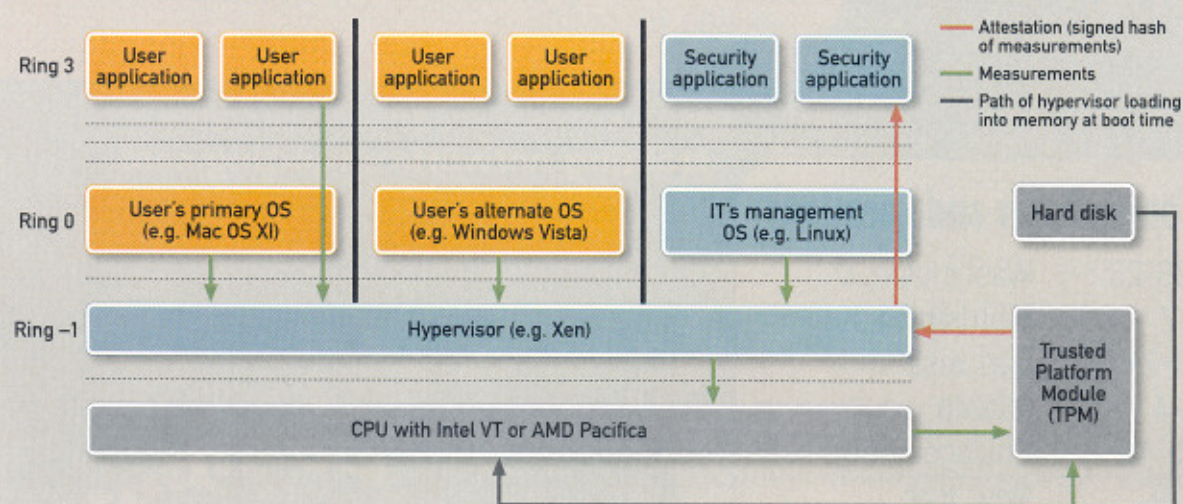
Similarly, users can have access to more than one OS. The concept is similar to current dual-boot systems, except that several partitions can run at once. For example, Intel says it's giving its software developers a Linux VM for their programming work, an empty x86 VM to test the compiled code, and a Windows VM to run Office applications. Even users who don't want to leave Windows could see benefits: They can use one VM to surf the Web and another to hold sensitive documents that shouldn't be exposed to the Internet.

Competition for the hypervisor has higher stakes. While VMs allow several OSs to share a system, there can only be one hypervisor. Windows servers will probably end up using Microsoft's. Clients and other servers will have a harder choice.

So far, there are two main contenders: VMware and Xen, an open-source hypervisor. The current versions still run at Ring 0—Xen uses paravirtualization, VMware emulation—but Intel and AMD are helping them move down to Ring -1. Both plan to support VT and Pacifica by the time the hardware is available.

Xen is the early favorite for embedded client management. It's used in all of Intel's embedded IT demos and has attracted code contributions from IBM as well as the

## Hypervisor Software Authenticated By PKI Hardware



VT and Pacifica both provide low-level access to the CPU for a hypervisor, the software that manages VMs. Each VM emulates the full capabilities of a legacy x86 CPU, requiring its own OS and applications. To ensure that the hypervisor isn't replaced by a rootkit, a physically separate TPM chip intercepts the hypervisor as it loads from the hard disk, digitally signing it for verification by security applications. During system operation, the TPM also provides security applications with signed measurements of data, hardware, and the virtualized software.

## product roadmap

chip vendors. For customers who don't feel comfortable downloading free software, some of its developers have formed a start-up, XenSource, to provide support and custom development work.

The server virtualization market still belongs to VMware. And to protect its position, it has formed a consortium including hardware vendors IBM and Dell, Linux leaders Red Hat and Novell, and Intel and AMD (see "Linux Virtually Ready For the Data Center," April 2005, page 26, or search for Doc ID# 2004f2 at [www.ITArchitect.com](http://www.ITArchitect.com)). The consortium aims to develop an open hypervisor standard, though it isn't clear yet whether Xen, Microsoft, and other competitors will be able to implement that standard.

VMware is also targeting home users with an intuitive user interface, offering features such as tabbed desktops (similar to tabbed browsing, but with VMs instead of Web pages). And it promotes virtualization as a security technology for the family PC. If you believe its demos, you'll be able to let your kids play with your computer, safe in the knowledge that even if they corrupt the OS, the damage will be limited to their own partition.

### -1 RING TO RULE THEM ALL

Virtualization can help protect a system against OS bugs or vulnerabilities, but it really just pushes security and stability problems down a level. The whole system is only as good as the hypervisor.

Fortunately, hypervisors tend to be robust. Most VMware products have never suffered a security advisory, a refreshing change to anyone accustomed to the frequent patches required by other software. And that's not just because of the programming skills of VMware employees. A hypervisor can be much smaller than a full-scale OS—Microsoft calls its own a "microkernel"—so auditing one for security is easier.

But VT and Pacifica can still introduce new vulnerabilities, especially for users who don't want the new VM capability. An attack on a system running

a single, non-virtualized OS wouldn't even require hacking the hypervisor, as the attacker could just slip a virus or Trojan into the unused Ring -1.

A Ring -1 virus is the ultimate rootkit. Because it operates beneath the OS and simulates the legacy x86 chip exactly, it can attack even perfectly secure software. What's more, it's OS-independent: The same virus can compromise every x86 OS, from CP/M to Solaris. Worst of all, it's mathematically impossible for software alone to detect.

To protect against such a virus, the system needs a hardware component that can't be virtualized. This is provided by the Trusted Platform Module (TPM), the controversial PKI chip already included in many PCs. The TPM watches the hypervisor and other programs as they load into memory, checking that they match precomputed hash values. Once it's sure that the hypervisor hasn't been tampered with, it signs a digital certificate that can be verified by the virtualized OS or security software.

This process, known as attestation, isn't limited to software. It can also prove whether or not particular components are present. In the original Palladium DRM architecture, it would be used to reassure a media player or video-streaming site that movies aren't being saved to a TiVo.

Intel and AMD both plan to do something similar in 2007, with technologies known respectively as La Grande and Presidio. Supposedly intended for enterprise security, these will encrypt the link to local USB and video devices, protecting against hardware keyboard sniffers.

In the meantime, VT and Pacifica both provide a compelling application for the TPM—even for enterprises that don't yet need VMs and hypervisors on desktops or laptops. While the chip has other uses such as disk encryption, virtualization-aware hardware could be what persuades users to activate it.

→ Read Andy Dornan's posts at <http://blog.ITArchitect.com> and write to him at [adornan@cmp.com](mailto:adornan@cmp.com).



MOBILE  
TECHNOLOGY

**AN OFFER  
THAT'S RIGHT  
ON TARGET.**



**Dell recommends  
Windows® XP Professional**

#### LATITUDE™ D610 NOTEBOOK

**Business-Class Performance, Compact Design – Starting at 5.17 lbs.\***

- Featuring Intel® Centrino™ Mobile Technology
- Intel® Pentium® M Processor 740 (1.73GHz)
- Intel® PRO/Wireless 2200 802.11b/g Wireless
- Microsoft® Windows® XP Professional
- 14.1" XGA Display
- 512MB Shared\* DDR2 SDRAM
- 40GB\* Hard Drive
- Modular 24x CD Burner/DVD Combo Drive

**\$1559** (as low as \$42/mo., 48 mos.†)  
E-VALUE Code: 06200-S71115n



**Call (toll free) 1.877.365.5415**

**Click [www.dell.com/smb/itarchitect](http://www.dell.com/smb/itarchitect)**

Dell, the Dell logo and Latitude are trademarks of Dell Inc. Intel Inside, the Intel Inside logo, Intel Centrino, and the Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. Microsoft and Windows are trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. © 2005 Dell Inc.