

Microsoft Highlights Changes to the Windows Kernel

May 17 2007

At WinHEC, Microsoft technical fellow Mark Russinovich discussed, among other things, how uniprocessor kernel variants were now gone from Windows Server 2008, which reduces the need for downtime.

LAS VEGAS - Attendees at Microsoft's Windows Hardware Engineering Conference here got an in-depth and deeply technical keynote outlining the changes to the Windows kernel and other key areas, and how partners can take advantage of these.

In his presentation on May 16, Mark Russinovich, a technical fellow in Microsoft's platform and services division and the third keynoter of the day, talked about how uniprocessor kernel variants were now gone from Windows Server 2008, which reduces the need for downtime by supporting hardware configuration changes without the need to reboot the system.

The new server product, which is the basis for Microsoft's new virtualization offering, also introduced a new common infrastructure called WHEA (Windows Hardware Error Architecture).

On the time accounting front, Russinovich said that Windows previously accounted for CPU time based on the interval clock timer and that thread quantum expiration was not always fair.

"Windows Server 2008 now reads the Time Stamp Counter at the context switch, which allows for more accurate quantum reporting," he

said.

Other infrastructure changes include an enhanced thread pool mechanism, new synchronization APIs, private namespaces and hard resource quotas.

"SMB is the original Windows remote file system protocol, but it could not adapt to the new NTFS features and was not designed for today's large data sizes. As such, SMB2 has been introduced in Windows Server 2008," Russinovitch said.

The server product also defers I/O completion until the thread pulls the I/O off the completion port, which avoids the context switch and improves the performance with no change to the applications, he said.

Windows Server 2008 also introduces I/O prioritization, based on the priority of the issuing thread or the explicitly set I/O priority.

"Also, in 32-bit Windows Server 2008, virtual memory is assigned as needed, with kernel page tables allocated on demand instead of at boot time. Kernel stack usage is also reduced through stack jumping. This all supports more users on terminal services," he said.

The product also brings memory manager performance improvements, with fewer and larger disk reads for page faults and system cache readahead. NUMA enhancements include more memory allocations being NUMA aware and now the I/O system directs interrupt completion to the node that initiated I/O.

The ideal node is used more effectively for process memory allocations and the new NUMA APIs allow applications to specify the preferred node number for memory applications and file mapping, Russinovitch said.

With regard to startup and shutdown enhancements, there is now parallel session creation in Windows Server 2008 rather than serial session creation, which was how it was done in previous versions of the product, he said.

Windows Server 2008 also has a clean service shutdown, which allows the system more time to shut down. Services can now request a pre-shutdown notification and, after these services stop, the system performs Windows XP-like shutdown, he said.

"This is also the first server system that is completely bug free," Russinovitch joked, adding that it was the third-party applications that could cause it to crash. But the new product is better at handling process crashes, with unhandled exceptions sending a message to the Windows Error Reporting service and all process crashes now getting recorded, he said.

"In summary, there are lots of exciting kernel changes for performance, scalability, reliability and security," he said.

Copyright 2007 by Ziff Davis Media, Distributed by United Press International

Citation: Microsoft Highlights Changes to the Windows Kernel (2007, May 17) retrieved 10 April 2024 from <https://phys.org/news/2007-05-microsoft-highlights-windows-kernel.html>

<p>This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.</p>
--