

New Approach to Virtualization Is a Lightweight

Steven J. Vaughan-Nichols


Virtualization is a hot topic in the technology world. The technology enables a single computer to run multiple operating systems simultaneously. It lets companies use a single server for multiple tasks that would normally have to run on multiple servers, each running a different OS.

For example, with virtualization, a single server could run several instances of the Apache Web server. Each could support a different Web site with its own network card while sharing the host server's RAM, hard drive, and processor.

The technology thus reduces the number of servers needed, thereby saving money and using resources more efficiently, according to Dan Chu, senior director of developer and independent-software-vendor products at leading virtualization vendor VMware.

Because of this, the virtualization market is growing rapidly, led by vendors such as SWsoft, VMware, and XenSource. Accelerating this trend are improvements in processor performance that have made it possible to run virtualization software even on small servers.

Now, vendors are releasing products based on two *lightweight virtualization* approaches that also let a single operating system run several instances of the same OS or different OSs. Each virtualized OS can then run its own set of applications, as Figure 1 shows.



This, too, lets users more fully utilize their computer's resources.

However, today's new virtualization approaches do not try to emulate an entire hardware environment, as traditional virtualization does. They thus require fewer CPU and memory resources, which is why the technology is called "lightweight" virtualization.

Lightweight approaches should experience a higher growth rate in the marketplace than traditional virtualization during the next five years, predicted senior analyst Andi Mann with Enterprise Management Associates (EMA), a market research firm.

However, lightweight virtualization still faces several barriers to widespread adoption.

STANDARD VIRTUALIZATION

Virtualization research dates back to the mid-1960s and the joint work of IBM and MIT on the M44/44X Project. Soon thereafter, IBM created and began selling a series of virtual-machine-enabled OSs, including the VM/370, to optimize usage of the company's mainframe systems.

Standard virtualization

The key to any kind of virtualization is the *hypervisor*, also called a virtual machine monitor. The hypervisor is the application that allocates the system resources for each virtualized OS or program. It can be installed as an application, or it can be part of the OS.

Multiple operating systems can run atop the hypervisor, which handles each OS's communications with the CPU, stored data, and the network.

Typically, the hypervisor provides a complete emulation of hardware devices to fool each virtualized OS into thinking it has exclusive access to an entire computer, said Jeff Jaffe, Novell's chief technology officer.

VMware's VMware, which runs on Linux or Windows, and Microsoft Virtual Server, which runs on Windows, are well-known examples of commercial traditional virtualization software.

There are also virtualization programs for the Mac OS, such as Parallels Inc.'s Parallels Desktop, and for Unix, like Hewlett-Packard's Virtual Server Environments for HP-UX.

Increasing popularity

According to EMA's Mann, "Almost 75 percent of surveyed enterprises in the US have already deployed virtualization in one form or another, and the virtualization market is increasing by approximately 26 percent [annually] on average."

Al Gillen, vice president of system software at market research firm IDC, predicted the number of servers worldwide using virtualization will grow from just under 175,000 in 2004 to just over 1 million in 2009.

LIGHTWEIGHT VIRTUALIZATION

Vendors such as Sun Microsystems, SWsoft, and XenSource are now selling software products—as well as the accompanying tools and services—that enable lightweight virtualization.

Leading Linux vendors Novell and Red Hat have begun building Xen virtualization software into their OSs.

SWsoft makes Virtuozzo products for Linux and Windows and recently released its core code to the open source OpenVZ virtualization project for Linux.

In developing its Viridian hypervisor, Microsoft is working with Xen.

Also, Intel with its Virtualization Technology and Advanced Micro Devices with its Secure Virtual Machine provide APIs that enable virtualization software to work optimally with their microprocessors.

There are two primary approaches to lightweight virtualization.

Containers

A *container* is software that virtualizes an OS's operating environment, according to Daniel Price, a Sun kernel engineer.

There is only one underlying operating system kernel, Price explained, which the containers enhance by providing walls that offer increased isolation between groups of processes. "The result is the ability to run multiple instances in isolation from each other within a single operating system," he said. Each instance would run within its own container.

Containers do not emulate any of the underlying hardware. Instead, the virtualized OS or application talks to the host OS, which then makes the appropriate calls to the real hardware.

In each case, a hypervisor apportioned system resources to the containers from both the local server and servers networked to it.

Container-based virtualization software includes SWsoft's Virtuozzo for Linux and Windows, the Zeosoft Technology Group's XR Enterprise for mobile environments, and FreeBSD Jail and Linux-VServer, both open source. Sun has also developed container software for its Solaris 10 OS.

With the container approach, applications must run on the host OS,

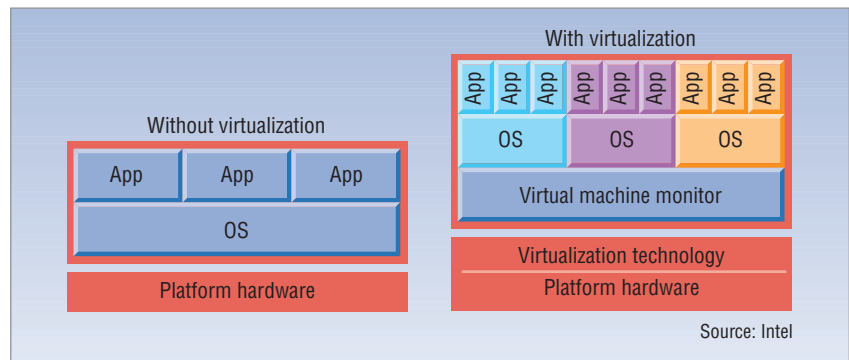


Figure 1. Lightweight virtualization lets a single operating system run several instances of the same OS or different OSs. Each virtualized OS can run its own set of applications. The virtual machine monitor, also called a hypervisor, allocates the system resources for each virtualized OS or program.

according to Sun kernel engineer Russ Blaine. Users thus generally can't run, for example, Linux applications on a Virtuozzo-enabled Windows system. Therefore, containers work best for running multiple native applications on their home OS.

However, there are efforts to change this, such as Sun's BrandZ project, which can run Linux binaries on Solaris 10.

Paravirtualization

Products such as XenSource's Xen and Virtual Iron Software's Virtual Iron work via *paravirtualization*.

Paravirtualization virtualizes part of an OS's operating environment but also selectively emulates the hardware devices that a virtualized OS requires, according to Novell's Jaffe. The process does not necessarily require complete hardware emulation—for example, an application might not require an emulated graphics card—and thus demands less management overhead.

Paravirtualization provides both a virtual machine and access to the native hardware and thereby lets users run non-native applications on an OS.

According to XenSource chief technology officer Simon Crosby, paravirtualization enables the operating system, the virtualized hardware, and the real hardware to collaborate to achieve optimal performance. In other words, the paravirtualized hypervisor lets the virtual OS make calls

to either the real hardware or the virtualized hardware, depending on what's better for performance.

Management and other tools

Tools made by both virtualization vendors and their allied OS vendors help users install and manage a virtualized system's multiple operating systems and applications.

Vendors enable management of lightweight virtualization systems in different ways. For example, Sun uses command-line-based scripts, which the GUI-based Solaris Container Manager controls.

SWsoft gives users a choice of a stand-alone management program with a graphical interface, a Web-based administration portal, or command line instructions.

Many industry observers see a push toward tighter integration between virtualization and management software to make the management process easier and more effective.

SWsoft is making container-management tools for multiple virtualization systems, not just its own. For example, the company's Virtuozzo management tools also support VMware, said SWsoft CEO Serguei Belousov.

Advantages

Both lightweight virtualization approaches minimize the use of processor and memory resources by sharing system calls with the host

operating system. The virtual OS thus doesn't need to duplicate the host system's functionality.

Moreover, systems load OS components such as the kernel and application libraries into RAM only once even though there may be multiple instances running. This also uses fewer system resources, which are then available for other purposes.

In addition, lightweight virtualization requires fewer computing resources than traditional virtualization, which emulates all system hardware. Containers virtualize only the operating environment. Paravirtualization virtualizes only the hardware that the virtualized OS requires.

Lightweight virtualization is thus faster than traditional virtualization because it doesn't spend time emulating all of the system hardware and looking for items it will never use.

In some cases, lightweight virtualization can help reduce system downtime. For example SWsoft can migrate virtualized OS instances from one server to another, allowing availability if the first machine crashes.

ON THE OTHER HAND

Lightweight virtualization faces several important hurdles to widespread adoption. For example, because traditional virtualization emulates all system hardware, it is more likely that any OS or application could run on such a system, making the technique more useful under some circumstances.

With containers, each implementation is tied to just one OS, which reduces user flexibility, said VMware president Diane Greene. Thus, containers might not work in systems using many legacy applications that require diverse OSs.

Moreover, Greene added, lightweight virtualization's functionality has outrun the ability to manage it. Thus, she contended, today's management tools are often still primitive.

Another problem with containers or paravirtualization is that a set of users in a virtualized system is basi-

cally working off a single OS kernel, noted Gabriel Consulting Group analyst Dan Olds. Thus, he said, a kernel crash would affect all users.

Lightweight virtualization uses fewer resources than traditional virtualization.

Multiple approaches

Because there are not only two types of lightweight virtualization but also different approaches to each type, companies might have trouble both selecting the best product for their purposes and making multiple systems interoperate.

Oracle, which has expressed interest in lightweight virtualization, is pushing for a simple, universal way to integrate a variety of virtualization solutions, said Bob Shimp, vice president of the company's technology business unit.

Greg Kroah-Hartman, a key Linux kernel maintainer, said the Linux community is trying to get virtualization vendors to work together on this.

As a result, companies such as IBM, Red Hat, VMware, and XenSource are collaborating on a paravirtualization API with which any Linux-based hypervisor could work.

Pricing

The way software vendors charge customers is a key marketplace-related problem for lightweight virtualization.

The standard pricing model assumes a single application or a single OS will run on a single computer. Now, a computer can run multiple operating systems or multiple instances of the same OS or application. But having to pay for every instance of an OS or application would make the process expensive.

Many vendors, including Microsoft, have somewhat modified pricing models in which they charge users per OS or application instance.

However, even the newer models could become expensive for customers that extensively use virtualization. And the general lack of agreement on how to bill for virtualization muddies the business case for using the technology.

Users are beginning to work with lightweight virtualization products, such as Virtuozzo and Xen, on PCs as well as servers. On PCs, the technology could create virtual machines for trying out possibly troublesome new software in a safe environment or isolating potentially vulnerable applications such as Web browsers.

The technology could also prove useful in *virtual relocation tools*, which let users move a virtual machine from one physical server to another, said Tom Bittman, a vice president with market research firm Gartner. This technique would be useful in disaster recovery and maintaining system availability.

With virtualization, multiple OSs and applications compete for the same hardware resources, so managing and troubleshooting problems could be difficult. Solving these problems will be a major factor in the competition among lightweight-virtualization vendors, said IDC's Gillen.

Regardless, lightweight virtualization is still relatively new, so it will take time for potential users to determine what they want from the technology and whether it will prove useful to them, particularly in mission-critical settings. ■

Steven J. Vaughan-Nichols is a freelance technology writer based in Arden, North Carolina. Contact him at sjun@vna1.com.

Editor: Lee Garber, [Computer](mailto:I.garber@computer.org),
I.garber@computer.org