

Introducing

Microsoft Virtual Server 2005:

Running Legacy Windows Environments on Dell PowerEdge Servers

Virtual server technology can enable organizations to consolidate their application workloads by porting legacy software environments to Dell™ PowerEdge™ servers that support the latest industry-standard hardware technologies. This article introduces Microsoft® Virtual Server 2005 server virtualization software, discussing its feature set and support for legacy Microsoft Windows® operating systems that run on Dell hardware platforms.

BY RANJITH PURUSH, JEFF WOOLSEY, AND EDWARD YARDUMIAN

Enterprise systems are often partitioned to create isolated operational environments for running multiple applications simultaneously. Initially, this technique was made possible through hardware and software partitioning implemented by proprietary technologies from enterprise server vendors. However, the disadvantages of proprietary technologies can include the associated cost of supporting legacy systems and limited scalability because of compatibility issues with other hardware and software technologies. With the future of computing looking toward standardized technologies such as those implemented in the Dell PowerEdge server line, proprietary technologies may no longer be a viable solution for the enterprise. In many cases, industry-standard systems such as Dell PowerEdge servers can provide better performance more cost-effectively than proprietary servers.¹

By allowing multiple operating systems to run on the same physical platform in complete isolation from one another, industry-standard server virtualization software is designed to allow organizations to consolidate² their legacy operating systems and applications from older, proprietary platforms to industry-standard Dell PowerEdge servers.

Microsoft has introduced Microsoft Virtual Server 2005,³ a virtual machine (VM) solution running on the Microsoft Windows Server™ 2003 operating system (OS), to help enable the efficient management of multiple OS environments on an Intel® x86-based hardware platform. Virtual Server 2005 is designed to provide a cost-effective approach for application migration and server consolidation. It is an integral part of the Microsoft Dynamic Systems Initiative (DSI), which aims to help organizations

¹For an example of this price/performance advantage, see "Migrating Databases from Sun Systems to Dell Servers Running Microsoft Windows Server 2003" by Todd Muirhead; Dave Jaffe, Ph.D.; and Kerstin Ackerman in *Dell Power Solutions*, October 2004.

²For more information about different server consolidation approaches, see "Approaches to Server Consolidation" by Todd Muirhead in *Dell Power Solutions*, August 2002, and "Using Windows Server 2003 on Dell PowerEdge Servers for Server Consolidation" by Paul Rad, Tim Cornett, and Ranjith Purush in *Dell Power Solutions*, August 2003.

³For more information about Microsoft Virtual Server 2005, visit <http://www.microsoft.com/windowsserversystem/virtualserver/default.aspx>.

reduce the complexity of their IT infrastructures.⁴ Based on virtual server technology from Connectix,⁵ Virtual Server 2005 leverages core Microsoft Windows application programming interfaces (APIs) for efficiently managing the virtual server environment. On September 13, 2004, Microsoft announced the release of Virtual Server 2005 (Enterprise and Standard editions), with general availability scheduled for October 1, 2004.⁶

Using Virtual Server for legacy application consolidation and improved hardware utilization

Continued support of legacy applications can require IT organizations to maintain obsolete operating systems and legacy hardware that are no longer supported. Porting legacy applications to a supported operating environment is usually not an option because of the significant development effort and investment required. One solution is to isolate the legacy operating environments from the hardware—which can help organizations overcome the burden of supporting legacy hardware and rewriting legacy applications. Virtual Server 2005 allows the physical server to be logically partitioned into several VMs, which can then host legacy operating systems and applications.

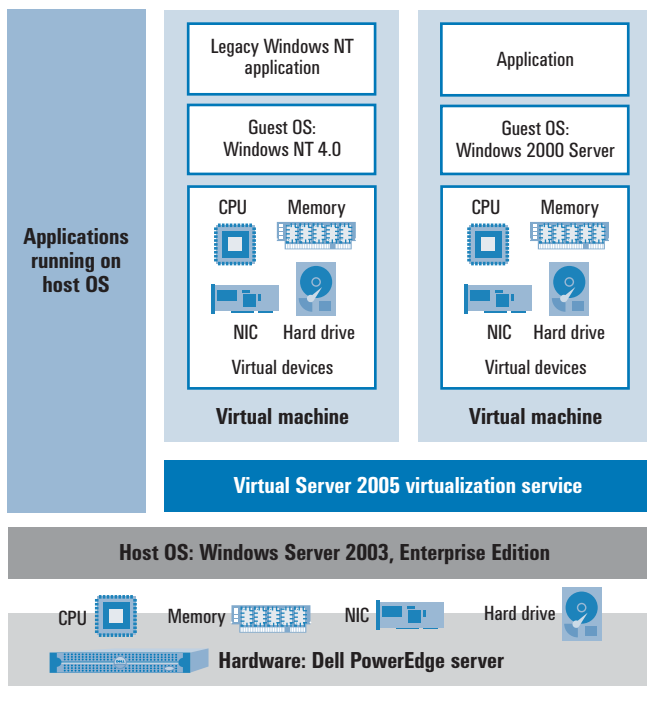


Figure 1. Application view of Virtual Server 2005

With a broad complement of hardware support for legacy environments such as the Microsoft Windows NT® Server 4.0 OS, Virtual Server 2005 on Windows Server 2003 is designed to be an excellent platform for running legacy operating systems. Most applications running on a legacy OS can be migrated to a Virtual Server 2005 VM environment. Exceptions include applications that have a specific hardware dependency on a unique peripheral. Figure 1 illustrates legacy operating systems and applications running simultaneously on top of Microsoft Windows Server 2003, Enterprise Edition, on the latest Dell PowerEdge servers.

Because many departmental line-of-business applications were written to run on dedicated servers, many application servers today are underutilized. Although innovations in hardware technologies can provide processing power that far exceeds application needs, some applications cannot coexist on the same server because of security, application compatibility, or data integrity concerns. Virtual Server 2005 is designed to provide the required application isolation. This isolation allows such applications to be consolidated from many unsupported legacy servers to fewer more powerful servers, thus helping to increase hardware utilization. Deploying fewer servers can enable enterprises to achieve economic benefits resulting from lower capital expenses and from savings on power, cooling, and space requirements.

In addition, Virtual Server 2005 is designed to provide easy-to-configure VM-to-VM clustering support for failover between VMs on the same physical server. This capability can increase productivity for software developers to create cluster-aware applications.

Although server virtualization has its benefits, the software virtualization layer between the application and the host OS introduces overhead because of the frequent context switching between the guest and host operating systems.⁷ This overhead is more obvious when processing workloads have high network or disk I/O, as compared to compute-intensive applications that do not require as much interaction with the network or disk subsystems. Although Virtual Server 2005 is well suited for applications with low I/O or CPU dependencies, such as point-of-sale (POS) software, high-I/O applications are recommended to be run natively on the host OS, not in a virtual environment.

Supporting multiple operating systems and hardware emulation

Figure 2 lists the host operating systems and guest operating systems that are supported by Virtual Server 2005. The Virtual Server Service (vssrvc) that runs on the host OS as a system service (with Network_Service logon credentials) implements the virtualization

⁴For more information about the Microsoft DSI, visit <http://www.microsoft.com/windowsserversystem/dsi/default.aspx>.

⁵Microsoft acquired the initial version of Virtual Server from Connectix, the original manufacturer of the Connectix Virtual PC software, in early 2003.

⁶For more information, visit <http://www.microsoft.com/presspass/press/2004/sep04/09-13AvailabilityVS2005PR.asp>.

⁷In software-based server virtualization terminology, the OS that runs on the physical server hardware is called the host OS and the OS that runs on the VMs is called the guest OS. Context switching refers to the transfer of control from the guest OS to the host OS; when a VM needs to access the physical hardware, control must be transferred from the guest OS to the host OS.

Category	Requirement
Host-server hardware requirements	<ul style="list-style-type: none"> Processor: Virtual Server 2005 Enterprise Edition supports up to 32 processors on the host server; Standard Edition supports up to 4 processors on the host server Memory: Virtual Server 2005 supports up to 64 GB of RAM on the host server
Host-server software requirements	<ul style="list-style-type: none"> Microsoft IIS 6.0 and later for supporting the Virtual Server 2005 Web application Microsoft Internet Explorer 5.5 and later for supporting the VMRC interface
Supported host operating systems	<ul style="list-style-type: none"> Windows Server 2003, Standard Edition Windows Server 2003, Enterprise Edition Windows Server 2003, Datacenter Edition Windows Small Business Server 2003, Standard Edition Windows Small Business Server 2003, Premium Edition Windows XP Professional (in nonproduction roles only)
Supported guest operating systems	<ul style="list-style-type: none"> Windows Server 2003, Standard Edition Windows Server 2003, Enterprise Edition Windows Server 2003, Web Edition Windows Small Business Server 2003, Standard Edition Windows Small Business Server 2003, Premium Edition Windows 2000 Server Windows 2000 Advanced Server Windows NT Server 4.0 with Service Pack 6a*

*Although functional in Virtual Server, Windows NT Server 4.0 will reach end-of-life status on December 31, 2004; after that date, Microsoft will no longer support Windows NT Server 4.0.

Figure 2. Hardware and software support for Virtual Server 2005

that allows multiple VM instances to run on the host OS (see the “Understanding the architecture and components of Virtual Server 2005” section in this article for more information). The Virtual Server Service is the only process that is exposed to the host OS, regardless of how many VMs have been created.

Virtual Server 2005 is designed to support up to 64 GB of host server memory and up to 64 VMs running concurrently on a single physical server.⁸ Each Virtual Server 2005 VM exposes a consistent set of virtual hardware devices to its OS regardless of the underlying physical hardware. This hardware independence enables the movement of VMs from one physical server to another—even when source and target physical servers have entirely different hardware configurations.

Figure 3 lists the hardware that Virtual Server 2005 is designed to emulate. A Virtual Server 2005 VM is designed to provide a completely isolated environment. Hardware on the host server is not presented in the VM. For example, support for USB devices such as smart card readers and scanners may be added in the future.

Understanding the architecture and components of Virtual Server 2005

Figure 4 shows the high-level architecture of Virtual Server 2005. Virtual Server provides two user interfaces to access the Virtual Server VMs: the Virtual Server Administration Web site and the Virtual Machine Remote Control (VMRC) interface. The Virtual Server

Hardware type	Emulation support
Processor	<ul style="list-style-type: none"> VMs use actual processor on host server. VMs support only 1 processor; symmetric multiprocessing (SMP) is not supported.^a
Memory	<ul style="list-style-type: none"> Up to 3.6 GB per VM can be supported.
IDE devices	<ul style="list-style-type: none"> Up to 4 virtual IDE devices can be connected to the virtual IDE bus. IDE devices can be virtual hard disks or CD-ROM drives; CD-ROM ISO (International Organization for Standardization) images^b are also supported.
SCSI disks	<ul style="list-style-type: none"> Emulated Adaptec[®] AIC-7870 SCSI controller chip set is supported. Up to 4 virtual SCSI adapters are supported. Up to 7 virtual SCSI disks are supported by each virtual adapter (for a total of 28 virtual SCSI disks). Each virtual SCSI disk can be as large as 2 TB.
Ethernet adapter	<ul style="list-style-type: none"> Emulated 100 Mbps DEC 21140 Ethernet card is supported. Up to 4 virtual NICs can be supported in a VM; Virtual Server emulates a multi-port NIC with 1 to 4 network connections on a single virtual Peripheral Component Interconnect (PCI) slot. An unlimited number of virtual networks can be supported.
Graphics adapter	<ul style="list-style-type: none"> Emulated S3 Trio 32/64 VGA/SVGA adapter with 4 MB VRAM is supported.
Floppy drives	<ul style="list-style-type: none"> A virtual 1.44 MB floppy drive (or floppy image^c) is supported.
Serial (COM) ports	<ul style="list-style-type: none"> Two virtual serial ports are supported; the virtual serial ports can be mapped to server serial port(s).
Printer (LPT) ports	<ul style="list-style-type: none"> One virtual printer port is supported; the virtual printer port can be mapped to a server parallel port.
Mouse	<ul style="list-style-type: none"> One PS/2-based Microsoft IntelliMouse[®] pointing device is supported.
BIOS	<ul style="list-style-type: none"> AMI AMIBIOS chip is supported.
Motherboard and chip set	<ul style="list-style-type: none"> Intel 440BX is supported.
Audio adapter	<ul style="list-style-type: none"> No emulated audio adapters are supported.

^aVMs are multithreaded, allowing threads to take advantage of multiprocessor servers and to be scaled accordingly. However, a VM can “see” only a single processor.

^bUse of CD-ROM ISO images is recommended over physical CD-ROMs because the CD-ROM ISO images have faster read speeds (times).

^cUse of virtual floppy disks (.vfd files) is recommended over physical floppy disks because they are software-based virtual devices that incur no mechanical overhead; read/write operations on a virtual floppy device are much faster than on a physical floppy. Virtual Server provides the capability to create virtual floppy disks.

Figure 3. Emulated hardware available to a Virtual Server 2005 VM

Administration Web site allows administrators to create, configure, and control VMs while the VMRC interface allows administrators only to connect to VMs and view the VM screen. The Virtual Server Administration Web site can also be used to connect to Virtual Server instances on remote machines. The Virtual Server Administration Web site communicates with the Virtual Server Service through the Virtual Server Web application (VSWebapp.exe), which is built into Virtual Server 2005 and requires the Microsoft Internet Information Services (IIS) 6.0 Web engine.

Figure 5 shows the Virtual Server Administration Web site of a Virtual Server instance running on an Intel Xeon™ processor-based Dell PowerEdge 6650 server. The customizable Virtual Server

⁸Support for more than 4 GB of memory by Virtual Server depends on support for Physical Address Extension (PAE) by the host OS. For more information about PAE and enabling PAE on the host OS, refer to http://www.microsoft.com/resources/documentation/WindowsServ/2003/datacenter/proddocs/en-us/Default.asp?url=/resources/documentation/WindowsServ/2003/datacenter/proddocs/en-us/paex86_2.asp.

Administration Web site allows administrators to monitor the status of all VMs running on a Virtual Server instance with thumbnail images of VM desktops and graphical views of the CPU utilization by the VMs. From the Master Status page, administrators can use the Turn On, Pause/Resume, Shutdown Guest OS, Turn Off, and Reset commands for managing the virtual machines.

The Virtual Server Administration Web site is the only user interface provided by Virtual Server 2005 that allows administrators to configure a VM with the available subset of virtual devices (as listed in Figure 5) before installing an OS on the VM. Even after the VM has been created, some virtual devices on the VM are by default not mapped to the actual physical devices on the host server. For example, the virtual serial ports (COM ports) are not automatically mapped to the physical serial ports of the host server. The mapping between the virtual serial ports and physical ports can be made using the Virtual Server Administration Web site.

The VMRC interface, which is accessed through the VMRC client, allows users to connect to and view the VMs. However, the VMRC client interface does not allow for the creation, configuration, or management of VMs. The VMRC client console connects to the VMRC server that is built into Virtual Server. As shown in Figure 4, the VMRC server communicates directly with the Virtual Server Service. To connect to VMs remotely using VMRC, administrators must install the VMRC client on the remote system.

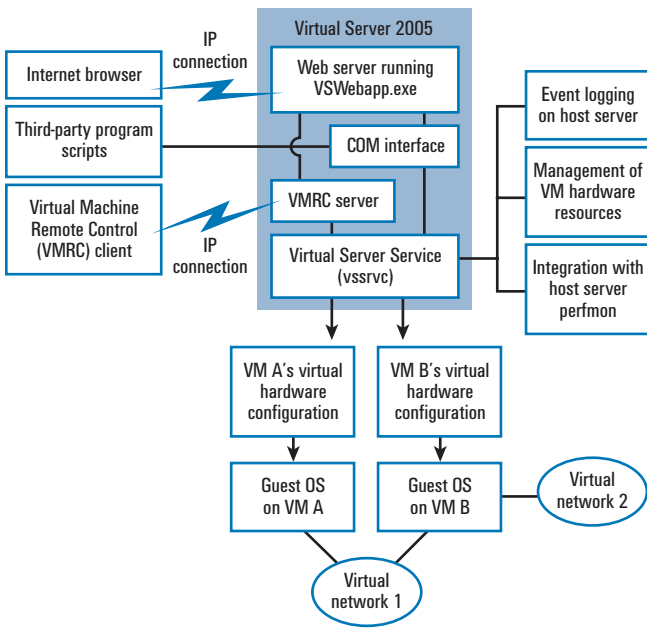


Figure 4. High-level architecture of Virtual Server 2005

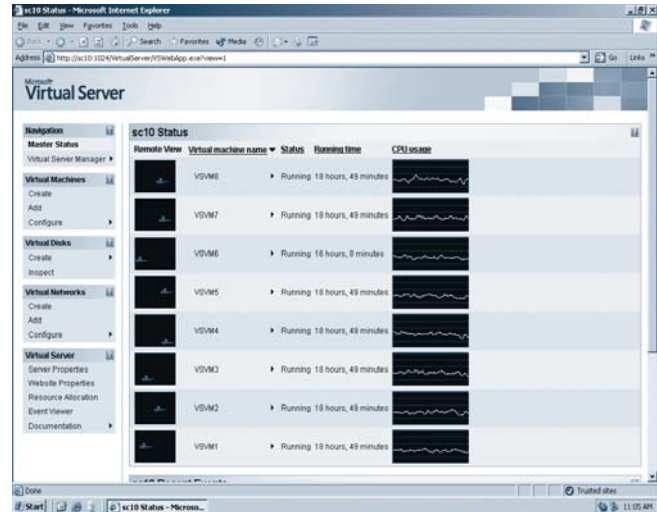


Figure 5. Virtual Server 2005 Web console

The Virtual Server Administration Web site also uses the VMRC server to allow administrators to view individual VMs.⁹ Once a VMRC session is started, either through the Virtual Server Administration Web site or through the VMRC client, the VM desktop can be remotely controlled in the same way a physical server's desktop is controlled from a direct attach console. Figure 6 shows a screen display of the VMRC client connected to an instance of Virtual Server.

The Virtual Server event-logging capability logs all Virtual Server-generated events in the host server's Windows Event Viewer under a category called Virtual Server. The supported event sources include Virtual Server, Virtual Machine, Remote Control, and Setting Change. These events are also accessible from the Virtual Server Administration Web site, allowing administrators to easily retrieve details about Virtual Server-specific events from the host physical server. The event log keeps track of changes made by VM users as well as those that are triggered by the Virtual Server Service.

Virtual hard disks

Virtual Server represents virtual hard drives as files called virtual disks, which are stored on the host OS file system as .vhd files. The Virtual Server Administration Web site serves as a single point for managing all the virtual disks in any Virtual Server instance, and it allows administrators to create different types of virtual hard drives. From within a VM, these different types of virtual disks appear the same; however, the difference is transparent only on the medium that actually hosts these virtual disk files. The medium may be the host physical server or a remote storage area network (SAN). Virtual Server supports the following types of virtual disks:

⁹This feature—viewing individual VMs using VMRC through the Virtual Server Administration Web site—is implemented by embedded Microsoft ActiveX® controls in the Virtual Server Administration Web site and hence is supported only by Microsoft Internet Explorer 5.5 and later.

- **Dynamically expanding virtual hard disk (VHD):** This VHD is designed to start with a sparse file and grow as data is added. Virtual Server 2005 can send low-disk warnings if a VHD grows too large on the host file system.
- **Fixed-size VHD:** This VHD consists of a fixed-extent file that resides on a host system's hard drive. Fixed-size disks cannot be resized.
- **Linked VHD:** This VHD points to and uses an entire physical disk for the purpose of converting a physical disk to a VHD.
- **Differencing VHD:** This VHD enables a disk hierarchy, with one or more child disks (the differencing disks) and a parent disk (the associated VHD). The differencing disk stores a record of all changes made to the parent disk and provides a way to save changes without altering the parent disk. Administrators should write protect the parent disk before using the differencing disk. Otherwise, if the parent disk is modified by some other process, all differencing disks related to it become invalid. On the host file system, differencing disks appear as two files, but from within the VM they appear as a single VHD.

Virtual Server 2005 supports the conversion of dynamic VHDs to fixed-size VHDs, fixed-size VHDs to dynamic VHDs, and physical disks to either dynamic or fixed-size VHDs. However, the conversion of disk types is not recommended in a production environment because of the large amount of processor resources required for the conversion process. Although the dynamically expanding VHD is the default disk type and is recommended because it does not take up disk space unless required, the

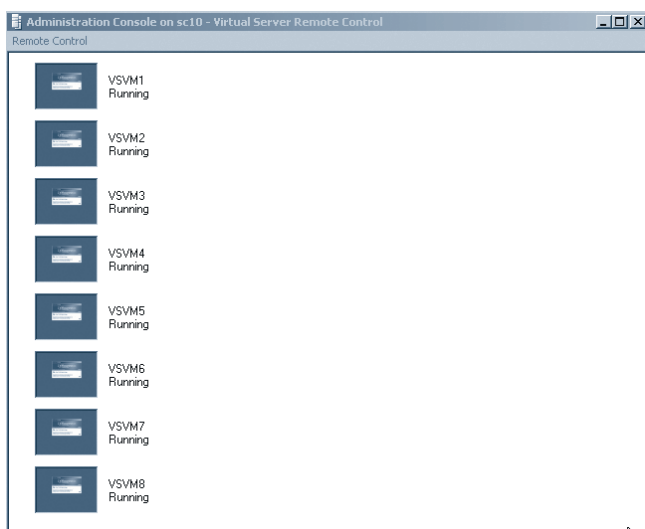


Figure 6. VMRC client console

additional overhead of dynamically increasing the file size of the virtual disk may marginally degrade performance.

Virtual networking

The Virtual Server Administration Web site allows administrators to create a virtually unlimited number of virtual networks on a physical server, and a virtually unlimited number of VMs can be connected to any given virtual network. Functionally, connecting the virtual network interface card (NIC) of a VM to a virtual network is analogous to connecting the network cable from the NIC of a physical server to a physical switch. Although the number of available ports on a physical switch limits the number of NICs that can be connected to that switch, a virtual network can be configured with as many virtual NICs as the implementation requires, depending on the workload and the network throughput. A virtual NIC can be connected to only one virtual network at a time, but virtual NICs may be dynamically added and removed from a virtual network using the Virtual Networking Manager.

For a VM to connect to the external network (the network outside the physical server), the VM should be connected to a virtual network that is associated with a physical NIC on the host server. The VM can then communicate with the external network to which the physical NIC is connected. In the absence of a physical adapter in a virtual network, only VM-to-VM communication (within the same physical server) is possible through that virtual network. The Microsoft Loopback adapter enables communication between the guest OS and the host OS on the same physical server without using the external physical network. For more information about how to configure the Microsoft Loopback adapter, refer to Virtual Server 2005 documentation from Microsoft, which is available as part of the Virtual Server 2005 installation.

CPU resource allocation

The Virtual Server Administration Web site allows administrators to allocate CPU resources to individual VMs based on application requirements or business priorities, helping to prevent resource contention among VMs. The advanced resource-allocation mechanisms in Virtual Server 2005 offer the flexibility of managing resources proportionally, absolutely, or through a combination of both methods.

Administrators can specify resource allocation settings for all VMs from a single page—the global CPU Resource Allocation page (see Figure 7). Using the CPU Resource Allocation page, administrators can designate the relative priority of VMs, set the maximum limit for CPU resources available to each VM, and reserve¹⁰ a specified amount of CPU resources for each VM.

The Figure 7 example shows the effect of changing the resource allocation on total and available capacities of the system. As seen in this Virtual Server configuration with eight VMs, the maximum

¹⁰The CPU resources reserved for a VM are available to other VMs or even the host server when the owner VM does not need the resources.

capacity for any VM is limited to 25 percent of the host server's processing power. This is because the absolute capacity available to any VM is measured relative to the total available capacity of the host server. Even though Virtual Server 2005 provides advanced mechanisms to manage resource allocation among VMs, virtualization technology is not a recommended application consolidation solution if heavy resource contention is likely to occur between the residing applications.

Automating the deployment and operation of VMs through scripting

Virtual Server implements support for programmatic control of Virtual Server VMs through the Microsoft Component Object Model (COM) technology. This capability is designed to expose COM objects to control most aspects of Virtual Server functionality such as turning on, restarting, and shutting down the VMs. The Virtual Server Administration Web site uses this rich COM interface—through Microsoft JScript® development software—for monitoring and controlling the Virtual Server VMs.

Virtual Server 2005 is designed to provide easy-to-configure VM-to-VM clustering support for failover between VMs on the same physical server. This capability can increase productivity for software developers to create cluster-aware applications.

To write Virtual Server automation scripts, administrators can use any scripting language that the Windows Scripting Host can execute and that can also connect to COM objects. This includes Microsoft Visual Basic®, C#, Perl, Java, and most other contemporary development languages. The COM components can also be used by other third-party applications to control the VMs. Third-party applications can communicate directly with the COM interface to control the VMs without the Virtual Server Administration Web site. For more information about the automated deployment support in Virtual Server 2005, refer to the documenta-

tion available as part of the Virtual Server 2005 installation.

As an add-on product to Windows Server 2003, Virtual Server 2005 requires its own license for installation. For more information about Virtual Server licensing, visit *Dell Power Solutions* online at http://www.dell.com/magazines_extras. More information is also

Virtual Machine	Relative Weight	Reserved Capacity (% of one CPU)	Maximum Capacity (% of one CPU)	Reserved Capacity (% of system)	Maximum Capacity (% of system)	CPU Usage
Virtual Machines Currently Running						
VSM1	100	0%	50%	0%	12.5%	
VSM2	100	50%	100%	12.5%	25%	
VSM4	100	0%	100%	0%	25%	
VSM5	100	0%	100%	0%	25%	
VSM6	100	0%	100%	0%	25%	
VSM7	100	0%	100%	0%	25%	
VSM8	100	0%	100%	0%	25%	
VSM3	10000	0%	100%	0%	25%	
Total Capacity Reserved				12.5%		
Available Capacity Remaining				87.5%		

Figure 7. CPU Resource Allocation page

available at <http://www.microsoft.com/windowserversystem/virtualserver/evaluation/virtualizationfaq.msp#EBAAA>.

Providing comprehensive support for legacy systems

By offering comprehensive support¹¹ for Windows Server operating systems that run on its VMs, Microsoft Virtual Server 2005 is designed to enable organizations to leverage the latest hardware technologies while maintaining the legacy Windows environments that are required to support business-critical applications. Cost-effective Virtual Server 2005 software running on Dell PowerEdge servers can help provide an industry-standard approach to server virtualization and application consolidation.

Ranjith Purush is a systems engineer in the Server Operating Systems Engineering Department at Dell. His current areas of focus include operating systems, virtualization software, and performance benchmarking. Ranjith has an M.S. in Electrical and Computer Engineering from The University of Texas at Austin.

Jeff Woolsey is a program manager for the Virtual Machine team in the Microsoft Windows Server Division at Microsoft Corporation. Prior to joining Microsoft, Jeff was the director of quality assurance at Connectix Corporation and was integral in the development and release of numerous award-winning products such as Virtual PC for Windows, Virtual PC for Macintosh, and the Connectix QuickCam.

Edward Yardumian manages the Operating Systems Engineering and Certification teams in the Dell Enterprise Product Group. Previously, Edward led engineering projects for clustering and next-generation PowerEdge servers. He has published numerous articles and has patents pending on cluster computing and scalable solutions.

¹¹Microsoft tests Windows operating systems running on a Virtual Server VM and provides support for its products running on a VM. For more information about supported host and guest operating systems, see <http://www.microsoft.com/windowserversystem/virtualserver/evaluation/virtualizationfaq.msp>.