

At a glance:

What is virtualisation?

Three virtualisation architectures

Microkernelised vs. monolithic hypervisor

What Hyper-V does

An introduction to Hyper-V in Windows Server 2008

RAJIV ARUNKUNDRAM

There has been quite a lot of talk about virtualisation recently, and most of the discussion is specifically about server virtualisation. This is one of the most exciting trends in the industry

and one that has the potential, over the next few years, to change the paradigm of how IT systems are deployed. But server virtualisation will not only change how IT administrators and architects think about servers and system utilisation, it is also going to affect the processes and tools used to manage what will certainly become an increasingly dynamic environment.

Virtualisation has actually been around for some time now, but the technology is still evolving. In fact, the word itself still means different things to different people. In broad

terms, however, virtualisation is about abstracting one layer of the technology stack from the next layer, like storage from servers or the OS from the applications. Abstracting the different layers, in turn, enables consolidation and better manageability.

As a concept, virtualisation applies to storage, networks, servers, applications and access. When you look at storage and networks, the goal of virtualisation is to aggregate a set of different devices so the total pool of resources looks and acts like a single entity. For example, you can configure a 40TB storage

solution instead of a set of 20 2TB storage devices. But with other components, virtualisation acts in the opposite direction, helping you to make a single system appear as though there are multiple systems. The most common example of this is server virtualisation, where you host multiple OS instances and environments on a single server.

Microsoft® has approached virtualisation at several different levels, extending from the desktop to the data centre with solutions for server virtualisation, application virtualisation, presentation virtualisation and desktop virtualisation. The common thread across all of these is the management piece with Microsoft System Center. For this article, I am focusing on the server virtualisation component and specifically on how Hyper-V, a key feature of Windows Server 2008, fits into the equation for a dynamic data centre.

The server virtualisation market

First, I think it would be worthwhile to look at what exists in today's environment and where the overall market is heading. Depending on what research you read, some analysts estimate that 5–9 per cent of all physical servers currently sold are being used as virtualisation hosts. You might consider this to be a big chunk of systems in a market where more than nine million physical servers are shipped every year. But one thing is certain: there is still a huge market opportunity as more customers become comfortable with virtualisation and want to employ it.

It is important to note where virtualisation is being adopted. Enterprise customers have certainly led the charge with testing and being early adopters. However, there are small and medium-sized businesses also deploying virtualisation. The adoption of virtualisation reaches across different types of workloads, from business applications and management to the web and e-mail.

So why is virtualisation now all the buzz? There are a few factors, not the least of which is timing. A few key industry factors have come together at the same time, helping to push for increased adoption of virtualisation. These industry factors include the move to 64-bit computing, multicore processors, and even the drive of sustainable computing to improve system utilisation.

Systems are becoming much bigger and they require a technology like virtualisation to make full use of system power. But while it is true that core technology (and Moore's law) has had a steady lead in terms of producing more processing capacity than systems can use, we are now also more conscious of the environmental impact, power requirements and cooling costs.

All of these factors, plus the easy justification on the return on investment (ROI) of adopting virtualisation, should together accelerate the adoption of virtualisation across both large and small businesses. And we, the IT professionals, can expect all the major players to continue to invest in this technology over the next few years and improve features and functionality.

How server virtualisation works

Server virtualisation in general terms lets you take a single physical device and install (and run simultaneously) two or more OS environments that are potentially different and have different identities, application stacks and so on. Hyper-V is a next-generation, 64-bit hypervisor-based virtualisation technology that offers reliable and scalable platform capabilities. Together with System Center, it offers a single set of integrated management tools for both physical and virtual resources.

All of this works to reduce costs, improve utilisation, optimise infrastructure and allow businesses to rapidly provision new servers. In order to help you better understand how Hyper-V is architected, I want to first take a look at the different types of virtualisation solutions.

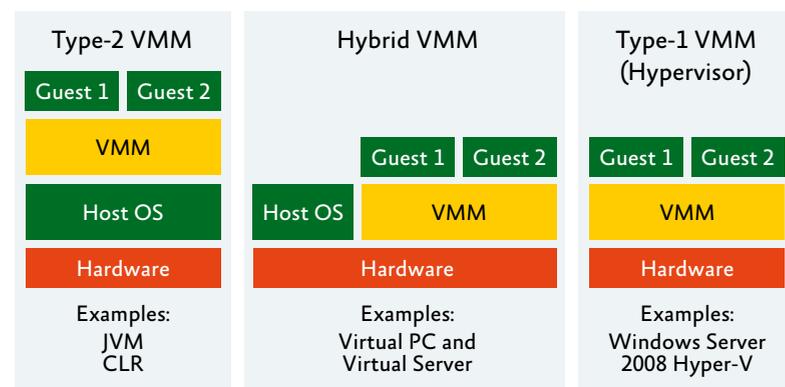


Figure 1 The three architectures of virtualisation

The move to 64-bit processors and the need for sustainable computing have helped to enable virtualisation

Types of virtualisation solutions

There are essentially three general architectures used for server virtualisation, as shown in Figure 1. The fundamental differences have to do with the relationship between the virtualisation layer and the physical hardware. By virtualisation layer, I mean the layer of software called the virtual machine monitor (VMM, not to be confused with Virtual Machine Manager). It is this layer that provides the ability to create multiple isolated instances that share the same underlying hardware resources.

The Type-2 VMM architecture is exemplified by Java Virtual Machines. Here, the goal of virtualisation is to create a runtime environment within which the process can execute a set of instructions without relying on the host system. In this case, the isolation is for the different processes, and it allows a single application to run on different OSs without having to worry about OS dependencies. Server virtualisation does not fall into this category.

Type-1 VMM and Hybrid VMMs are the two approaches you are most likely to find in wide use today. The Hybrid VMM is a stage where the VMM runs alongside the host OS and helps to create virtual machines on top. Examples of the Hybrid VMM are Microsoft Virtual Server, Microsoft Virtual PC, VMware Workstation, and VMware Player. You should note that while these types of solutions are excellent for a client scenario where you are only running virtual machines part of the time, the VMMs add considerable overhead and therefore are not suitable for resource-intensive workloads.

In a Type-1 VMM architecture, the VMM layer runs directly on top of the hardware.

This is often called the hypervisor layer. This architecture was originally designed in the 1960s by IBM for mainframe systems and has recently been made available on the x86/x64 platforms with a variety of solutions, including Windows Server 2008 Hyper-V.

There are solutions available on which hypervisor is an embedded part of the firmware. This, however, is simply a packaging option and does not really change the underlying technology.

As you look at Type-1 VMMs, there are essentially two main ways to architect the hypervisor solutions: microkernelised and monolithic. Both of these approaches, as shown in Figure 2, are true Type-1 VMMs that have the hypervisor installed directly on the physical hardware.

The monolithic hypervisor approach hosts the hypervisor/VMM in a single layer that also includes most of the required components, such as the kernel, device drivers and the I/O stack. This is the approach used by such solutions as VMware ESX and traditional mainframe systems.

The microkernelised approach uses a very thin, specialised hypervisor that only performs the core tasks of ensuring partition isolation and memory management. This layer does not include the I/O stack or device drivers. This is the approach used by Hyper-V. In this architecture, the virtualisation stack and hardware-specific device drivers are located in a specialised partition called the parent partition.

Windows hypervisor

Ensuring that there is strong separation between multiple OSs is done by creating virtual processors, memory, timers and in-

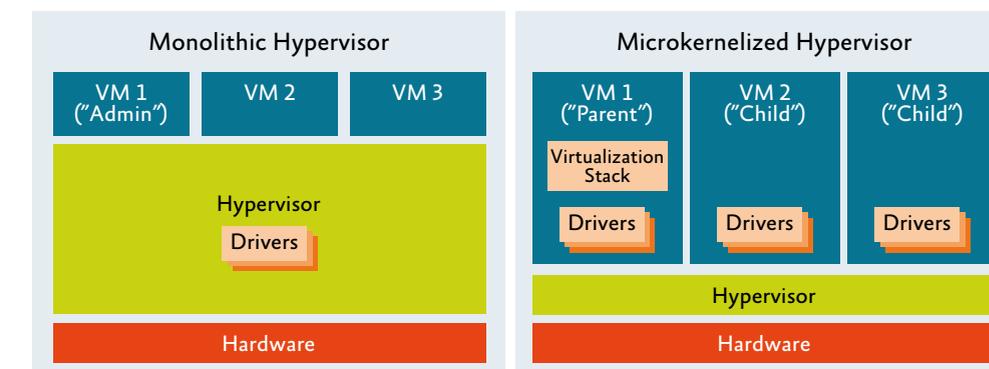


Figure 2 The two ways to architect hypervisor solutions

terrupt controllers. OSs use these virtual resources just as they would use their physical counterparts.

The Windows hypervisor, part of Hyper-V, performs the following tasks:

- Creates logical partitions.
- Manages memory and processor scheduling for guest OSs.
- Provides mechanisms in order to virtualise input/output and communicate among partitions.
- Enforces memory access rules.
- Enforces policy for GPU usage.
- Exposes a simple programmatic interface known as hypercalls.

Since it uses the microkernelised approach, the Windows hypervisor is fairly small – less than 1MB in size. This minimal footprint helps improve the overall security of the system.

One of the requirements for running Hyper-V is that you have an x64 system that has Intel VT or AMD-V technologies. x64 technology enables access to a larger address space and support for systems with more memory, and thus allows more virtual machines on a single host system. Intel VT and AMD-V are hardware-assisted virtualisation solutions that provide an ultra-privileged layer in the ring architecture that helps to keep the execution environment of the hypervisor separate from the rest of the system. They also allow Hyper-V to run an unmodified guest OS without incurring significant emulation performance penalties.

The parent partition

Hyper-V consists of one parent partition, which is essentially a virtual machine that has special or privileged access. This is the only virtual machine with direct access to hardware resources. All of the other virtual machines, which are known as guest partitions, go through the parent partition for their device access.

The existence of the parent partition is fairly transparent. When you begin to install Hyper-V, the first thing you must do is install Windows Server 2008 x64 Edition on the physical system. You then need to go to Server Manager, enable the Hyper-V role and restart the system. Once the system has re-

started, Windows hypervisor is loaded first, and then the rest of the stack is converted to be the parent partition.

The parent partition has ownership of the keyboard, mouse, video display and other devices attached to the host server. It does not have direct control over the timers and interrupt controllers that the hypervisor uses.

The parent partition contains a Windows Management Instrumentation (WMI) provider to facilitate management of all aspects of the virtualised environment, as well as a virtualisation stack that performs hardware-related tasks on behalf of the child partitions. In addition, any independent hardware vendor (IHV) drivers needed for host system hardware are contained in the parent partition, and any drivers created for Windows Server 2008 x64 editions will also work in the parent partition.

Device-sharing architecture

One of the innovative architectural components in Hyper-V is the new device-sharing architecture that supports emulated and synthetic devices in each guest OS. Device emulation is quite useful for supporting older OSs with device drivers designed for older generations of hardware. For example, Hyper-V includes an emulation of the Intel 21140 network adapter, which was called the DEC 21140 network adapter at the time many older OSs were being shipped.

Generally, device emulation is slow, not easily extendable, and doesn't scale well. But emulation is still important because it allows you to run most x86 OSs on Hyper-V. Since virtualisation is now moving from a niche technology primarily for testing and development to an essential technology for production environments, users require better performance in order to run larger workloads. Emulated devices no longer meet these growing demands.

An alternative solution to this is to use Hyper-V synthetic devices. Synthetic devices are virtual devices that are mapped directly to physical devices. Unlike emulated devices, synthetic devices do not emulate legacy hardware. With the Hyper-V hardware sharing model, supported guest OSs can interact directly with synthetic devices that may have no physical counterparts. These OSs use vir-

One of the most innovative components of Hyper-V is the new device-sharing architecture

tual service clients (VSCs), which act as device drivers within the guest OS.

Instead of accessing physical hardware directly, VSCs use the VMBus, which is a high-speed, in-memory bus, in order to access virtual service providers (VSPs) in the parent partition. The parent partition VSPs then manage access to the underlying physical hardware, as illustrated in **Figure 3**. A key benefit of synthetic devices is that performance of synthetic devices over the VMBus is closer to performance of non-virtualised hardware devices.

Integration components

Hyper-V was built to provide strong boundaries between various instances running on one computer. To enable interaction between the guest OS and the host OS and to supply some additional functionality for supported guest OSs, Hyper-V provides integration components.

The Hyper-V integration components support the following features:

- Time synchronisation
- Volume Shadow Copy Service (VSS)
- Heartbeat functionality
- Guest shutdown
- Key value pair exchange (used to access the registry of a guest OS)
- OS identification

The Hyper-V feature set

It goes without saying that the closer the virtualisation platform comes to acting like the physical server, the easier it becomes for organisations to deploy and rely on virtual workloads. In my view, there are four key areas under which you can view the different features of the virtualisation platform.

Today most hypervisor-based virtualisation solutions are pretty close to each other in terms of features and functionality. As we move forward, things like total cost of ownership (TCO) and ease of use will be key differentiators. And the management solutions will see continued investments and development to bring us closer to the vision of a dynamic IT environment, where the infrastructure is flexible enough to adapt to the needs of the business, and models and policies help drive increased automation and management.

Scalability

Using the microkernelised hypervisor architecture, Hyper-V has very low CPU overhead, leaving plenty of room to virtualise workloads. By enabling virtual machines to take advantage of powerful features and hardware, such as multicore technology, improved disk access and greater memory, Hyper-V improves scalability and performance of the virtualisation platform.

Combined with the rest of the Windows Server 2008 capabilities, Hyper-V allows you to consolidate most workloads – including 32-bit and 64-bit workloads – on a single system. And it can help you balance 64-bit technology adoption with continued support for 32-bit workloads already used throughout your environment.

The fact that Hyper-V requires a 64-bit host system with hardware-assisted virtualisation helps ensure that the host system can access a large pool of memory resources. Hyper-V can support up to 1TB of memory on the host, with up to 64GB of memory per virtual machine. This is key for those who plan to virtualise memory-intensive workloads such as Exchange Server and SQL Server.

Hyper-V also supports up to 16 logical processors on the host system, making Hyper-V scalable to most commodity two-socket and four-socket systems with multiple cores. You

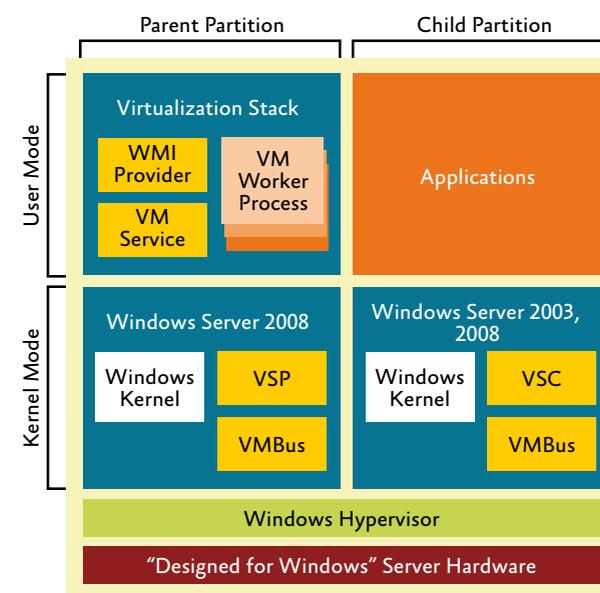


Figure 3 VSCs use the VMBus to access VSPs, which then manage access to underlying physical hardware

can also create a virtual machine with up to four virtual processors in order to support workloads that require or take advantage of multi-processor capabilities.

Consolidating servers through Hyper-V also enables those servers to make use of robust networking support, including VLAN, Network Address Translation (NAT), and Network Access Protection (NAP) policies (quarantine). And as a Windows Server 2008 feature, Hyper-V works well with other Windows Server features, such as BitLocker™ and Windows PowerShell™.

High availability

High availability is a scenario where Hyper-V and host clustering capabilities work together to help address business continuity and disaster recovery needs. Business continuity is the ability to minimise both scheduled and unscheduled downtime. That includes time lost to routine functions, such as maintenance and backup, as well as unanticipated outages.

Disaster recovery is an important component of business continuity. Natural disasters, malicious attacks and even simple configuration problems such as software conflicts can cripple services and applications until administrators resolve the problems and restore data. A reliable business and disaster recovery strategy must offer minimal data loss and powerful remote management capabilities.

When looking at high availability, you should consider three different categories – planned downtime, unplanned downtime and backups. Protection for planned downtime is typically needed to help move the virtual machines off the host system so you can either perform hardware maintenance or apply patches to the host system or the virtualisation platform (which may potentially require a reboot).

Most organisations have planned maintenance windows, and what you are really looking to do here is to minimise or eliminate the period of time in which the virtual machines will not be available while the host system is down for maintenance. With the Quick Migration feature, you can rapidly migrate a running virtual machine from one physical node to another in a mat-

ter of seconds. By doing that, you can keep your virtual machines available for production use while you perform maintenance on the original host. Once the maintenance is done, you can then use Quick Migration to return the virtual machines back to the original host system.

Unplanned downtime is downtime that is not foreseen. It can be catastrophic in nature or as simple as someone accidentally unplugging a power cord and bringing a server down. Although that may sound unlikely, over the years I have met quite a few administrators at Tech•Ed, VMworld and other conferences who have stories to tell about how some server was accidentally powered off by a colleague.

With Hyper-V, you can set up a host cluster for the different host systems and configure all the virtual machines as cluster resources that can then failover to a different system in case one of the hosts fails. Meanwhile, the multi-site clustering capability of Windows Server 2008 will enable you to set up a geographically dispersed cluster so that if your primary data centre fails, you have the ability to recover the different virtual machines to a remote data centre.

This is also handy for protecting all your branch offices. One of the advantages of the unplanned downtime support with Hyper-V is that it is guest OS agnostic, which means you can extend its high availability benefits to Linux virtual machines and older versions of Windows Server to protect and recover those systems similarly.

As you look at unplanned downtime, it is important to note that the recovery is equivalent to powering off the system and restarting, which means you will have lost all state information. This might or might not be a problem, depending on the workload you are running in the virtual machine. That's why it is important to look at backup in the context of high availability.

Hyper-V lets you take backups of each virtual machine or use VSS to take consistent backups of all the VSS-aware virtual machines while they are still running. With VSS, you can set up backups to occur at set intervals without having an impact on the production workload availability while ensuring that you have a continuous backup

With high availability, Hyper-V and host clustering capabilities work together to address continuity and disaster recovery

plan that can help you easily recover state in the event of an unplanned downtime. For more information on high-availability solutions with Hyper-V, see the article

The microkernelised hypervisor architecture is designed to minimise the attack surface and to enhance security

“Achieving high availability for Hyper-V” by Steven Ekren (<http://technet.microsoft.com/magazine/cc837977>).

Security

The microkernelised hypervisor architecture is designed to minimise the attack surface and to enhance security, especially when Hyper-V is implemented as a Server Core role. Server Core is an installation option of Windows Server 2008. The hypervisor contains no device drivers or third-party code, promoting a more stable, thin and secure foundation for running virtual machines. Hyper-V also delivers strong role-based security with Active Directory integration. And Hyper-V allows virtual machines to benefit from hardware-level security features, such as the execute disable (NX) bit, further helping to increase the security of virtual machines.

Hyper-V has gone through the Secure Development Lifecycle (SDL) like the rest of the Windows Server components, and extensive threat modelling and analysis has been carried out to ensure that Hyper-V is a highly secure virtualisation platform. As you deploy Hyper-V, be sure to follow the best practices for deploying Windows Server 2008 and also

RAJIV ARUNKUNDRAM is a senior product manager at Microsoft focused on server virtualisation in the Windows Server Marketing division. Rajiv's primary responsibility is to work with customers and partners to help them understand the virtualisation strategy and solutions of Microsoft from a technical and business perspective.

the best practices for Hyper-V. Include Active Directory® as well as antivirus and anti-malware solutions as part of your plan. And use delegated administration capabilities to ensure that you use admin access privileges appropriately for Hyper-V hosts.

Manageability

It is easy to go from a slight server sprawl problem to a massive virtual machine sprawl. This risk arises from the ease with which you can deploy virtual machines. And with the increased mobility of virtual machines, you also have the added responsibility of knowing exactly where the different virtual machines are running, keeping track of their security contexts, and so on.

Fortunately with Hyper-V, you don't have to create a separate management infrastructure for your virtual environment. It integrates with Microsoft management tools, System Center Virtual Machine Manager and Microsoft System Center Operations Manager, as well as with third-party management tools, so you can manage your physical and virtual resources from one console. For details of System Center Virtual Machine Manager 2008, see Edwin Yuen's article, “Manage your virtual environments with VMM 2008,” in this issue of *TechNet Magazine* (<http://technet.microsoft.com/magazine/cc836456>). Meanwhile, support for Windows PowerShell makes it easy to automate tasks.

Hyper-V also provides virtual machines with an unprecedented ability to use available hardware. Because all Windows Hardware Quality Lab (WHQL)-certified drivers are able to run in the parent partition, Hyper-V delivers broad compatibility for drivers and devices, making it easier to manage the different drivers that are running in your environment.

Wrapping up

As I mentioned earlier, management is going to be a key area of development and differentiation. You will certainly see a lot of activity in that area in the years to come. With virtualisation becoming more mainstream, these are exciting times. ■

For more information on Hyper-V, visit: www.microsoft.com/Hyper-V