



# WHITE PAPER

## The Microsoft Virtualisation Strategy

---

*What makes this enterprise class virtualisation solution different?*

*Author: Aidan Finn, BSc, MCSE, MVP (ConfigMgr 2007)*

*Systems and Infrastructure Manager, C Infinity.*

*<http://www.cinfinity.ie>*

## Table of Contents

Introduction .....	3
Virtualisation Technologies.....	4
Server Virtualisation .....	4
Session Virtualisation.....	9
Desktop Virtualisation.....	9
Application Virtualisation .....	12
Why Choose the Microsoft Strategy: Management.....	13
Health and Performance Monitoring.....	13
Backup and Recovery.....	14
Virtualisation Management .....	15
The Future of Microsoft Virtualisation .....	17
Hyper-V .....	17
Windows Server 2008 R2.....	18
Remote Desktop Services .....	20
Choosing a Virtualisation Platform .....	21
About Aidan Finn .....	23
About C Infinity .....	24

## Introduction

It used to be that virtualisation was locked away safe in the hands of geeks or software testers where it wouldn't ruin the day for the rest of us. However, the technology has matured to the point that virtualisation is now the primary architecture of choice for deploying servers, desktops and information worker environments. That's right; virtualisation is more than just for servers!

It seems that no matter what industry newspaper, conference or website we pick up or attend, we're bombarded with marketing materials and arguments about virtualisation solutions. If you're an architect or a decision maker you need to know more than the tiny details; you need to think about the big picture. It seems that everyone is getting into battles about tiny details and misquotes. Maybe it's time to think about the big picture.

How will these new technologies fit into my existing infrastructure, how will I manage them and will they add positively to my business?

Microsoft has been in the virtualisation business for years. When we first think of virtualisation we have a very limited point of view ... machines. As you'll soon read, there's much more to virtualisation than reducing the number of servers in the data centre.

It's true that Microsoft has been in the server virtualisation market for a while but they gained little acceptance outside of the test lab. They sold their solutions as being ready for the enterprise but we rightly greeted those comments with great scepticism. IT professionals viewed their 2 products for PC's and Servers as nothing more than point solutions with very limited functionality when compared with the competition.

However, Windows Server 2008 has changed the game somewhat. New technologies and a new strategy mean that we now need to re-evaluate the solutions we are being offered. It's important that we view the direction of the entire infrastructure and how we want to manage it when deciding what virtualisation technology we will adopt.

This document will discuss the Microsoft virtualisation strategy and why you might want to seriously consider adopting it. There will be a comparison with competitor solutions. These competitive solutions are excellent products and I'll try to discuss them a little. However, there is a good understanding in the market of these technologies so I'll focus a bit more on the less understood Microsoft strategy.

I will be discussing technology but I don't intend this to be a "level 400" technical analysis. This document is intended to be a briefing to help you with gathering information for a fair analysis of alternative solutions in your decision making process.

## Virtualisation Technologies

### Server Virtualisation

We've become familiar with the concept of server virtualisation and the market has now accepted it as a method for deploying the majority of production servers.

*If you're not familiar then here's a quick overview. Studies of servers found that the majority of them used only a small percentage of processing power, usually around 8%. We have powerful servers so why don't we try to do more with their CPU capacity? Using a very thin piece of software referred to as a hypervisor that runs directly on the hardware, we can simulate the hardware of many computers on a single physical server. The physical machine is referred to as a host and the simulated machines are known as virtual machines (VM's) or guests. Each VM executes independently from the others, just like physical servers. They can execute hand-in-hand on the network with other virtual machines and with physical servers. Normal networking strategies such as VLAN's, VLAN tagging and firewalls apply. Each virtual machine is represented on the storage of the host as a collection of files, usually a configuration file(s) and a virtual disk. With virtualisation we can make more of processing power but we often still need the same memory capacity (some solutions can minimise this by only supplying memory as required) and we actually need slightly more storage.*

It's still early days for the technology despite being stable, secure and long ready for production. Studies believe less than 10% of the world's servers are the virtualised. The vast majority of servers are still the traditional rack or tower machine with a single operating system and role on the network. There is a big move to change the face of the data centre. What's changing to make people radically redesign the server network?

If you talk to an IT administrator about server virtualisation he's going to talk about being able to quickly deploy machines. They're nothing but a bunch of files in the "real world". Combine these with traditional machine deployment solutions such as Microsoft's SYSPREP and you can deploy a virtual server in a matter of minutes. Compare that with a rack server which can take a day hour to get racked, cabled and installed – and we've not even talked about the logistics of ordering and delivering the hardware. Not everyone is a cabling whiz so virtual machines take care of that operational nightmare too. The administrator now is able to spend more time working on important projects for the business rather than dealing with time consuming operational tasks.

The IT architect will talk about disaster recovery and agile computing. Remember that a virtual machine is nothing more than a few files when you get down to it? That makes it relatively easy to replicate these files to a remote location. This makes it pretty easy to provide site to site replication of servers for business continuity or disaster recovery. There's plenty of out of the box and even appliance solutions for doing this on the market.

Virtualisation uses simulated hardware. The simulated drivers don't change, even if the underlying physical hardware changes. This means that we're no longer tied to physical hardware. If a machine ages beyond support we can migrate a VM to newer hardware with little effort and time. Even better, we can provide hardware fault tolerance to *all* virtual machines thanks to this abstraction of the hardware. Using enterprise ready virtualisation solutions, the physical hosts are connected together in a cluster with shared storage. All of the virtual machines that execute on those hosts are kept on the shared storage. During normal operations the cluster runs at N+1 or N+2 capacity, i.e. the virtual machines are load balanced across the cluster with enough host RAM, I/O and CPU capacity to allow for 1 or 2 host failures. If maintenance is planned on a host, its virtual machines are quickly moved to other hosts with little or even (virtually) zero downtime. If a physical host fails, its virtual machines will automatically fail over thanks to the cluster to the hosts with spare capacity. The architect now can provide a highly available infrastructure that is agile for future unforeseen circumstances such as migrations, failures and business strategy changes.

The financial controller will also be happy. Based purely on the upfront costs it appears that virtualisation is a strategy for the larger enterprises only. True enough, a large capital investment is required to install the required large capacity hosts and shared storage. The hardware chosen is typically blade servers (due to their dense rack capacity) and either an iSCSI or fibre channel SAN. Those aren't exactly cheap ticket items. Without digging a little further, the financial controller won't approve any budget requests. Virtualisation means we get the same computing power using less physical equipment. We're reducing the size of the computer room. This reduces rent, electrical power for servers and cooling and air cleaning/conditioning requirements. These operational costs far outweigh the cost of a server over its typical life of three years. The biggest of these costs is power and the cost of electricity is going only one direction. In certain markets, we're facing the imminent threat of carbon taxes and that's only going to compound the issue. The more we can minimise these costs, the better. Replacing server hardware is also a major project that typically involves lots of consultants. We can replace hundreds if not even thousands of servers with a single rack of blade servers. We've already seen how a virtual machine can be moved from one host to another with no downtime. I think I see a smile forming on the face of that financial controller as the IT budget requirements go down over a 3 year period.

There are many vendors in the server virtualisation market. In fact, 2008 has been an exciting year with new solutions and new entries into the hypervisor sphere.

VMware is the best known of the solutions and easily has the largest share of this market. They rightly deserve credit for driving this technology through the dark ages when many of us saw virtualisation as an RND technology only and finally burst our misconceptions. VMware has got two very basic offerings with ESXi and ESX. ESXi is a low end and efficient solution where the host only runs a hypervisor. There native operating system on the host

and ESXi does not allow for clustering, i.e. hardware fault tolerance of virtual machines. ESX is the enterprise solution. It features a Linux operating system for advanced management of the host and allows for virtual machine clustering. VMware's VMotion is the best known feature of ESX and it makes for a great live demonstration. A virtual machine is shown on the network performing some operation. The VM is migrated from one host to another. It does go offline but only for milliseconds. It's virtually zero downtime between running on one host and running on another. This technology has made VMware a very attractive solution.

Xen is a virtualisation technology that has tried to compete against VMware for some time. Citrix acquired this technology and displayed the seriousness of their intent to the market by renaming everything they have under the Xen banner. Citrix's strategy ties XenServer tightly with their other solutions.

In the summer of 2008, we saw the newest of the players to enter the hypervisor virtualisation market. Microsoft had spent the previous year demonstrating and publicly testing Hyper-V, a hypervisor based on Windows Server 2008. The strategy was simple. Hyper-V would be version 1.0 of their product, good enough to satisfy the needs of most installations.

Hyper-V is a hypervisor virtualisation product that is similar in concept to VMware ESX. Hyper-V runs at a lower layer in the CPU architecture thanks to virtualisation assistance technologies provided by AMD and Intel in the more recent generations of their products. If you remember back to your Computer Science classes, you might remember that the kernel of an operating system runs at Ring 0. This is a very privileged and secure location in the CPU. VMware ESX's hypervisor runs at Ring 0 and its virtual machines run at Ring 1. Hyper-V is a little clever; it runs at Ring -1. This allows the virtual machines to run at Ring 0. This means that an operating system running raw on a physical machine runs at the same location on the CPU as an operating system on a Hyper-V virtual machine or "child partition". In Hyper-V we've now got software running at a very important location in Ring -1. Microsoft protect that using a technology called Data Execution Prevention (DEP). Working with the hardware, Ring -1 is protected from buffer overflow attacks by preventing blocks of memory that are marked as data from being executed as instructions.

Installing Hyper-V is pretty easy. First you install Windows Server 2008 x64 onto a host where the CPU has virtualisation assistance enabled and the BIOS has DEP enabled. You download an update with the latest version of Hyper-V from Microsoft and install this update. In Server Manager (Windows Full Installation) or OCSSETUP (Windows Core Installation) you enable the role and reboot. The hypervisor is slipped underneath the original operating system at this point. The OS installation is now referred to as the "parent partition". It is used to manage the hypervisor and the server (operating system and hardware). You can install monitoring and management agents on this parent partition to

manage the complete machine. Hyper-V is nothing like the hosted solutions offered by Microsoft's Virtual Server 2005 R2 SP1.

The next advantage of Hyper-V is the huge amount of hardware that it can be run on. The first thing you need to do with VMware ESX is check their very tightly controlled hardware compatibility list (HCL). VMware builds their virtual machine drivers into the hypervisor. This means they must write drivers for every piece of hardware that they will support ESX on. There's a plus and a minus to this. On the positive side, VMware can tightly control the hardware in terms of testing and development. They can provide a near 100% predictable experience much like Apple Macintosh owners would expect with their computers. On the negative side, VMware is pretty much restricted to high end hardware and they will naturally be relatively slow to add support for new hardware innovations. You'll find the supported hardware on the HCL tends to be on the pricey side. Hyper-V works by directing guest machine driver traffic to the parent partition via a 1-1 secure channel called a VM Bus. The hardware drivers reside in the parent partition. This means that any hardware with a Windows Server 2008 logo on it can run Windows Server 2008 and Hyper-V. That's a pretty huge HCL.

On a tangent, this comparatively small ESX HCL means that ESX is an expensive solution for test environments and is beyond the budget of small scale implementations. It makes it a near impossibility for self paced training, even if you do get your hands on a legal trial copy. In my opinion that will make ESX more of a niche skill set compared to the more available Hyper-V solution which will run on nearly any hardware you buy now and is only a download and a tick box away.

Hyper-V has most of the functionality required of a virtualisation platform. The ability to create virtual switches, build virtual machines on secured VLAN tagged networks, and create snapshots and roll them back.

It might sound like that I'm doing a hatchet job on VMware. I'm not. VMware still has some advantages.

VMware ESX has the ability to bind multiple network cards to a virtual switch. This gives A+B networking to every virtual machine, thus offering divergent networking paths. This fault tolerance is important if a physical networking path should fail outside of the virtual machine, e.g. a host NIC or physical switch. Microsoft has long held onto the mantra of “Windows does not support NIC teaming”. So Microsoft did not build in the ability to bind multiple host machine NIC’s to a virtual switch into Hyper-V. They originally said that OEM’s would allow for it with their NIC teaming solutions. Unfortunately, when Hyper-V is installed on a host, it changes networking enough that the NIC teaming solutions provided by the likes of Intel, Dell, HP or Broadcom no longer work. We then had a little finger pointing between Microsoft and the OEM’s. I’ve asked around and it appears that the OEM’s and Microsoft are working together from executive level down to work out an industry solution. At the time of writing, I was told to expect an announcement in the coming weeks.

The next problem is sort of complicated and is really two intertwined issues. VMware ESX with Virtual Center allows virtual machines to move between hosts with no perceivable downtime. This is called VMotion. It relies on two things: a shared file system between the hosts where file locks can be respected and a clustering solution. Hyper-V provides virtual machine fault tolerance by using Windows Server 2008 failover clustering. It has to be said that failover clustering in Windows 2008 is excellent. It’s a *huge* improvement over Windows 2003 and all the previous versions. 3 day projects can now be completed in a matter of minutes. However, Windows has lacked a cluster file system like VMware’s VMFS. This means that each virtual machine that is independently able to failover from one host to another requires its own SAN storage or LUN. We only have 26 letters in the English language and at least 4 of those are taken in a 2 node Windows 2008 cluster. We can’t get around the problem of requiring one LUN per VM but we can use letterless drives that are identified by GUID’s. This takes a little getting used to but clear naming of storage in the SAN management and in the failover clustering administrative console makes things a little more manageable.

Finally, VMware has the ability to oversubscribe to RAM on a host machine. The theory is that each virtual machine is assigned a certain amount of RAM, e.g. 2GB. However, in reality, that VM may only need 1GB of that RAM 90% of the time. That means we have assigned 100% of RAM that is used 50% of the time. By only allocating RAM as required, ESX can allow many more VM’s than there is RAM for. Be careful with this; you need to be sure you can accommodate RAM loads across your host cluster during peak periods. Hyper-V does not have this feature. Any VM that is allocated RAM will immediately take this RAM as soon as it starts up.

But like I said earlier, Hyper-V is a version 1.0 product. It's extremely stable and performs excellently. Microsoft is aware of some of the competition ... but I'll come back to what we can expect in 2010 in a little while.

### Session Virtualisation

This is probably one of the more familiar forms of virtualisation on the market but we've never really thought of it as a virtualisation technology until recently. Terminal Services based solutions have been around for years. Even before that, Citrix pioneered the technology and their brand name is often synonymous with the technology, such has been their success. They are not the only brand name out there; there are many more and they've all reached certain levels of success by offering features not found in the solutions offered by Citrix or Microsoft.

The current release of Windows Terminal Services has shaken things up somewhat. By now offering a secure access, web-based interface (TS Web), application publishing (RemoteApp), a way to access those over SSL via the Internet (TS Gateway) and single-driver thin printing (Easy Print) Windows Server 2008 Terminal Services might just be good enough by itself for many small to medium business implementations. This isn't to say that Citrix and the other providers have no place in the market. Quite the opposite in fact! They still offer cleaner solutions, better architectures and more features that will be required in larger or more complex deployments.

Microsoft may have hit a winner with Easy Print. For once, the marketing material can be 100% believed, solving a highly complex and painful issue for users and administrators alike. It is "log in and play" technology requiring absolutely no configuration at all and offering hard to believe results.

### Desktop Virtualisation

This is one of the more interesting solutions available in the virtualisation world. It's been known to us as Virtual Desktop Infrastructure (VDI) for some time but Microsoft's "naming department" would like us to refer to it as Virtual Enterprise Centralised Desktop (VECD); I'll just refer to it as VDI from now on.

VDI leverages the concept of machine virtualisation, a now well known and accepted technology. A desktop operating system is used instead of installing a server operating system in the central virtual machine. Each user is given access to a virtual machine via a protocol such as the Remote Desktop Protocol. The result is that each user has a familiar working environment (a virtual PC with its own execution environment) and they don't share an operating system with other users.

A very crude solution has the user directly connecting to the virtual machine. This might be fine for lab work or for the smallest of implementations. However, a broker solution can offer a very smooth and efficient experience. Microsoft doesn't have a broker service yet but numerous partners are already building upon Hyper-V:

- The user signs into a single central point, e.g. an SSL secured website, to log into their virtual machine.
- The broker connects the user to a virtual machine.
- If a virtual machine is not available then a new one is deployed on the underlying machine virtualisation infrastructure.
- Policies define how virtual machines are deployed from one of potentially many templates *and destroyed* – we can't have "VM creep" endlessly consuming resources. Some users might be assigned virtual machines that only exist as long as the user remains logged in. Some users might get virtual machines that only contain the differences between the template virtual machine and that user's session. That's using operating system streaming technologies. And some users will have virtual machines that permanently reside on the virtualisation platform.
- A secure website can be presented to the Internet allowing roaming access to the internal infrastructure via a secure SSL session. The user can work from home, an airport or a hotel room without having any data on their laptop or thin terminal thus minimising risks to the organisation, its clients, suppliers and shareholders.

What problems does VDI resolve and how can it make a difference?

- Session Virtualisation (or Terminal Services) seems to be a fantastic solution when there is a stable set of end user applications. However, what do you do when you add those dodgy line-of-business applications that we're painfully aware of? You know what I mean; they have a service running in the operating system that crashes and requires a reboot to fix, maybe they share user data between sessions that they shouldn't, or they cause random issues between user sessions. You will probably have experienced these issues with Terminal Services at some point in the past. When they occur, they usually require a reboot or worse, a regularly scheduled set of rolling reboots across the server farm. VDI solves this quite simply; each user has their own (virtual) machine and they don't share a session on that (virtual) machine. If an application misbehaves on that machine, it does not affect anyone else. In the end, users are getting a near PC-like experience ... just remotely.
- The branch office network is a problem that just doesn't seem to go away, especially now when we have virtual teams (there's that "V" word again!) trying to collaborate with each other across the Internet or corporate WAN. The laws of physics restrict the performance of the technology in the form of network latency. We've seen network hardware solutions from the likes of Riverbed and Citrix that surely help but the end problem is still there. We have machinery out on the corporate network that is difficult and expensive to deploy, update, secure, fix and replace. What if every user logged into machines that resided in central locations and allowed them to centrally store and share data? That's what VDI does! The virtual machine farms are in central data centres, just a few milliseconds away from the

application and storage servers. Users generate, modify and save their data in a central location. A person in Krakow can easily share a file with a person in Dublin without the file ever leaving the data centre. The need to have frequently replaced PC's and servers in the branch office is gone as well is the need to maintain, power and secure large computer rooms in every corporate location.

- One of the things I dislike about Terminal Services is that it diminishes the usefulness of the helpdesk. I happen to believe that the helpdesk is a powerful asset in any company. They know the end users, their roles and their challenges. When used properly, helpdesk not only fixes problems (allowing senior engineers to focus on projects and servers) but they also act like an intelligence gathering or early warning system. When listened to, the helpdesk can forewarn IT about upcoming issues and enable IT as a whole to prevent them. Consider this: what happens when a user calls up saying that something is wrong with Office on their PC? A helpdesk engineer can quickly log in (hopefully using Remote Assistance or Remote Control), diagnose and fix the problem by making the required changes. What happens when the same user calls up saying that Office is misbehaving on a Terminal Server? The helpdesk engineer will check the problem out and diagnose it. Now it has to be escalated to a Terminal Services engineer because the helpdesk administrator doesn't have the required privileges or skills (yet). The engineer investigates the problem. Once a solution is found it has to be tested in a lab and documented. Change control takes place and a signatory signs off on the fix. A maintenance window is planned and the fix is deployed across the entire Terminal Server farm and monitored. About 2 weeks after the call was opened, the user got their fix and they could continue doing whatever task was interrupted.
- The migration from a PC to centralised computing solution for the end user isn't a cheap one and there's a certain amount of re-education of users and particularly administrators that needs to take place. The last things any organisation will want to do is rip and replace absolutely every management component and procedure that has been developed over time with significant amounts of effort and personal investment. The beauty of VDI is that you can reuse existing practices and procedures because we've still got individual desktop operating systems per user. We can recycle PC's as terminals as dumb terminals: locking them down using Active Directory Group Policy. System Center Configuration Manager can still be used to manage the end user infrastructure just like it did before. Users need the smallest amount of re-education to familiarise them with how the VDI broker works; they still have a desktop and applications like they always did on their monitor. Now it just runs in a centralised data centre.
- Security. I'm not going to sit here and type ghost stories or list out registry entries. With a VDI deployment you no longer have data on the desktop. With roaming Internet access laptops no longer need to store data. Using Group Policy you can even prevent storage of data on those physical machines! Users log into physically secure central data centres where there is auditing and complete control of all services and data. The access point is a simple SSL gateway. You don't have to worry about domain controllers being stolen from insecure branch office computer rooms, laptops going missing without disk encryption, etc.

Isn't it starting to sound like things are getting easier? What remains to be seen is how these VDI technologies will mature. There are still questions regarding thin printing

(Terminal Services Easy Print is king), issues regarding document scanning from the Terminal (large data streams being transferred up to the central virtual machine), Internet access for roaming users and peripheral device compatibility.

## Application Virtualisation

This is a newer technology and is really only familiar to a smaller set of organisations, consultants and administrators. The concept is driven by client application complexity:

- Consider a large corporation, finance company or university with hundreds if not thousands of desktop applications. Every time there is an operating system or single application upgrade a huge series of compatibility tests must be performed, thus delaying the deployment and costing the organisation in terms of time, effort and money.
- Software deployment is becoming more and more complicated with infinitely more software installation mechanisms: scripts, setup.exe, INI files, MSI installations, transform files, network installs, CD only installs, administrative rights VS installing as the user as a requirement, etc. The desktop administrator doesn't need all this complexity and the end user doesn't care about it; they just want their client application now so they can do their job in time for a deadline.
- Multiple PC's per user is becoming more and more common due to incompatible applications that are a business requirement. I've heard stories of users having 3 PC's. I've seen core line-of-business applications having conflicting requirements, e.g. Office 2000 VS Office 2003, where they need 2 machines. And I've heard of much worse!
- Administration is becoming more complex and time consuming. Demands of the business whether they be project or financial mean that IT departments are trying to do more than ever with fewer staff than ever. This impacts the business user badly. Again they don't care why getting access to Visio is taking so long they just need it now to do a job. Self service installation is possible with application virtualisation. That sounds like a licensing and budget nightmare but combined with a workflow, the owner of the budget can approve or reject requests for applications. IT is removed from the deployment-to-the-PC process and the end user can do their job with only the minimal and required financial delays.

Application virtualisation can be used on the PC but where it's really making an impact is when the technology is bound with VDI or Terminal Services. You're getting the best of those technologies while minimising application complexities and ownership costs.

Microsoft acquired SoftGrid several years ago to add this technology to their arsenal. Unfortunately they bundled it into the Microsoft Desktop Optimisation Pack, a license that is only made available to Software Assurance customers. This very clever solution has been further developed and rebranded as App-V.

## Why Choose the Microsoft Strategy? Management

When I look at bringing a new solution onto the network, I evaluate how I will be able to manage it. Microsoft has anticipated this and they are aiming to differentiate their virtualisation product suite with their management solution. Microsoft provides enterprise management using their System Center family of products.

## Health and Performance Monitoring

Most mid to large networks have some form of systems monitoring for network, hardware, operating systems and applications. There's a few ways to do this:

1. Use niche or point solutions: This is often considered as the budget approach but the budget may be misleading over the long term. Generic products or point solutions need to be deployed or heavily customised every time something is added to the network. This requires understanding what can go wrong, specific scripting languages (often many of them at once), SNMP and MIB's, etc. And each solution can require hardware and operating system licensing to operate. You're also expected to understand everything that can go wrong so you know how to monitor for that circumstance. Unfortunately, these monitors often fail to go beyond the most basic of ping-style tests and offer little intelligence about the health and performance of the services that IT provides to the business. However, this may be the right solution for some networks.
2. Use a single solution: There are tried and tested enterprise management products such as Tivoli, OpenView, etc, that aim to be all things to all products. The advantages are that expertise is only required in one brand of products and that operators only need to use a few or maybe even only one console to manage a huge network. These products tend to be highly scalable. However, these products tend to be highly expensive too! Licensing is huge and a purchase of these solutions often requires a team of expensive consultants camping out in your office for what may seem like an eternity. There's also another thing to consider; one vendor is trying to provide you with expertise in every product on your network. Is that really feasible?
3. Use vendor expertise: This is an approach where the original manufacturer of a product provides the management expertise for this product. Each of these *management packs* is added to a single management infrastructure.

It just so happens that the latter is the approach that Microsoft took with System Center Operations Manager. OpsMgr 2007 allows management packs to be added as modules to grow the management expertise of products that are deployed within the organisation.

Imagine if Microsoft's Active Directory team wrote the management product for their directory services solution, HP produced expertise for their storage products, Dell produced management packs for their servers or Citrix produced maintenance knowledge for their solutions? You know what? They do. Each of these vendors or teams within Microsoft writes a management pack for their product. A management pack is made up of expertise

for the product in question. That expertise can vary from simple event log monitors, scripts to check functionality or agents that monitor hardware. The management pack can detect anything in terms of health or performance that the original vendor thought was important. This is what makes this solution very different from the competition. You're not just limited to this out-of-the box experience either. Management packs are easily customised to your unique requirements. Specialist firms develop expertise for products that don't have manufacturer authored management packs. And there are utilities available for writing your own management packs.

## Backup and Recovery

The basis of Microsoft's backup solution is a service called Volume Shadow Copy Service (VSS) that is built into Windows. Additional products such as Exchange, SQL Server and Hyper-V allow their files to be backed up in a consistent state.

There's two ways we can backup a virtual machine. The first is to backup the few files that make up the virtual machine while it is running. This "snapshot" is a photograph of the VM as it was at that time:

- The file system.
- The RAM.
- The processor.

In theory, if we recover that snapshot and start it up the machine should execute exactly as it was when it was backed up.

You can't just backup files and hope to get this effect. Hyper-V allows this to function correctly by using VSS. VSS will slow the virtual machine down and stop it in an unperceivable time frame to allow a consistent snapshot to be taken. Microsoft's System Center Data Protection Manager is capable of using this Hyper-V VSS writer to get this functionally. A number of Microsoft partners have either already or are soon to release updates of their backup solutions to allow snapshot backups.

Be wary that snapshot backups should never be recovered into production for certain types of server, e.g. Active Directory domain controllers. Do a search for "USN Rollback" and the scare you get there will ensure you never make that mistake.

The other type of backup is to do what you would have done with a physical machine. Install an agent into the virtual machine's operating system and back it up at a system state and file/service level. This is probably the best solution anyway for normal operational backup/recovery.

Microsoft would contend that DPM is the best solution to go with because it integrates tightly with Active Directory and can be monitored by OpsMgr 2007. Personally, I'd say that

you should stick with what you're familiar and comfortable with as long as it has the support of your virtualisation platform. Just be sure the backups are tested, secure, are securely stored offsite and can be rapidly and predictably recovered from.

You also might want to have a look at some of the DR solutions provided by third parties that will replicate virtual machines to a powered off DR site. This allows for the simplest form of DR site replication and invocation – yet another reason to consider machine virtualisation for your production environment *where appropriate*.

## Virtualisation Management

Of the management products I've mentioned so far, this is probably the one you place highest on your list when you are planning a virtualisation environment. True, it goes hand in hand with the host cluster(s) but I'd argue that being able to monitor the health & performance as well as securely backing up the applications and data of your servers are just as critical.

VMware's ESX can be managed with the excellent Virtual Center. I say "excellent" because it really is a product you can pick up and work on with almost no training. If you're familiar with the concepts of virtualisation then it's no challenge at all. It's Active Directory integrated and allows you to set up delegated administration based on your existing administrative model. You can also provide easy remote access to the console and to virtual machines via a SSL secured web portal.

Virtual Center has easily been the poster boy for virtualisation platform management up until now. If you're a VMware customer you're probably extremely happy with it and I can't blame you.

Microsoft released Virtual Machine Manager 2007 as part of the System Center suite to manage their free hosted virtualisation product, Virtual Server 2005 R2 SP1. It was a decent first attempt but working only with a limited virtualisation product, it was never going to compare well with Virtual Center.

With the release of Hyper-V, Microsoft developed and recently released VMM 2008. This added support to Hyper-V. It features Active Directory integrated administrative delegation, a web portal for managing and accessing VM's, the ability to convert physical machines into virtual (P2V) and it greatly simplifies managing Hyper-V clusters, especially when it comes to dealing with GUID (letterless) drives.

What is truly daring is the ability to manage ESX and Virtual Center using VMM 2008. Imagine this: you can manage many Hyper-V clusters alongside many ESX clusters. Microsoft actually released a "manager of managers" before VMware did! The "manager of managers" function will really only be useful to the huge deployments. Microsoft was probably thinking of those smaller/mid-sized deployments that already has ESX and were

interested in Hyper-V. You have the ability to manage both platforms will complete functionality from one “pane of glass” into you management suite. And of course, you could always migrate those VM’s on ESX onto Hyper-V but I doubt anyone in Redmond considered that when developing this functionality.

I said earlier that the differentiator for Microsoft’s virtualisation strategy was management. All we’ve seen so far is a trading of punches between two software giants. VMM 2008 offers something really unique called PRO (Performance and Resource Optimisation). Virtual Machine Manager 2008 and Operations Manager 2007 are integrated using a connector. A management pack referred to as “PRO Tips” is installed on Operations Manager 2007 to make it fully aware of the hypervisor functionality of Hyper-V and the management ability of VMM 2008. OpsMgr 2007 is now able to monitor the health and performance of everything from the hardware (thanks to HP, Dell, IBM, etc) and all the way through the virtualisation, operating system, service and application layers of the server platform. VMM 2008 manages the machine virtualisation infrastructure. Now, when something goes wrong, OpsMgr can detect it automatically and raise an alert. VMM 2008 will, by default, suggest a fix to the problem that OpsMgr has detected. You can configure this response to be completely automated. This configuration can be filtered by computer group and alert severity level. The magic of Pro Tips doesn’t end there. Microsoft’s hardware partners have announced that they are developing PRO Tips for their products. These include NetApp, Emulex and HP. Imagine if a SAN fibre channel host bus adapter (HBA) reports that it’s fully utilised and causing storage latency issues for the VM’s hosted on its physical server. The VM causing the issue is identified and can be moved to another host. If you enable automatic responses in VMM 2008 then this happens silently and automatically minimises the impact on the business.

You’ve now seen aspects of dynamic IT (rapidly responding to change and issues with expertise and automation built into the network) and optimised IT (minimising impact on human effort and the business by using automation and good design) in action. The business was affected by its own workloads and the issue was quickly resolved to allow the business to generate revenue.

This strategy is what makes the Microsoft virtualisation solution different. Being able to integrate the configuration of the management platform and integrate that seamlessly with the health and performance monitoring of it offers huge benefits and simplifies the ownership of the architecture. Unfortunately for VMware ESX, to get the same results we rely 100% on third party solutions which cannot offer the same levels of expertise that VMware could if their management platform went beyond the hypervisor.

## The Future of Microsoft Virtualisation

### Hyper-V

Due sometime in 2010, Windows Server 2008 R2 will feature version 2 of Hyper-V. We can expect some new features including support for more logical processors per host and hot-swap “hardware” (virtual devices). The most anticipated changes are Live Migration and the Cluster File System.

The Cluster File System (CFS) will be Microsoft’s answer to VMware VMFS. This will allow a cluster of Hyper-V servers to share a single or a small number of very large disks. Many virtual machines can be saved per disk. This will simplify VM deployment and make the self-service portal of Virtual Machine Manager a more realistic option – right now we can control how much virtual resources a user can consume but not which pre-configured physical disks they take. The main purpose of the CFS is to allow many Hyper-V servers lock individual files or VM’s for their own use without having an exclusive lock on the disk. This means that the first step of Live Migration is possible.

Live Migration is the near instant failover of a VM from one host to another. VMware has branded their solution as VMotion. With Windows Server 2008 Hyper-V, this isn’t possible; we have the slower quick migration which still can serve most people’s needs but not everyone’s. Part of the relatively slow failover of Quick Migration is the failover of the individual LUN for each VM. By putting all VM’s on a single shared LUN, this is no longer required. Instead there is simply a transfer of locks on the VM’s files by the source and target host servers.

Live Migration requires a second piece of functionality where the state of the virtual machine (the VM configuration, the RAM and the state of the CPU) can be consistently transferred between the source and target servers without any visible stoppage of the VM. Hyper-V V2 has a clever decrement system where pages of RAM are transferred between the hosts during a migration. Any pages that change during the process are sent again. This process repeats up to 10 times. If all RAM is transferred successfully before the 10<sup>th</sup> iteration then:

- The VM is paused on the host server
- The configuration of the VM and the state of the CPU are copied over to the target server.
- The VM files are unlocked by the source and locked by the target servers.
- The VM is un-paused on the target server.

If the RAM is not successfully transferred after the 10 RAM copy iterations then the remaining “dirty” pages are copied over during that state transfer after the pause. This entire process should be unperceivable to humans and take nothing more than a few

milliseconds. If you search around, you can find Microsoft demonstrating this technology now.

Microsoft is also searching for new ways to improve performance. Some of these solutions involve working with hardware vendors to develop new architectures to improve virtualisation performance. Microsoft is working with CPU manufacturers Intel and AMD to improve how memory transfer happens. Microsoft is also working with network card manufacturers to move virtual machine networking functionality from the hypervisor and into the hardware. These may seem like tiny savings per transfer but they soon stack up into big processing savings.

## Windows Server 2008 R2

At the recent TechEd EMEA 2008 conference, there was a big emphasis by Microsoft on power savings and greener IT. Saving on power costs is one of the aims of virtualisation. We're certainly saving on electrical and cooling costs by consolidating the hardware into fewer physical machines. However, they are bigger machines. Server power consumption seems to be split between two big consumers: the CPU which generates heat and the rest of the chassis which is designed to remove that heat from the CPU.

Microsoft is developing "Core Parking" for Windows Server 2008. By monitoring the workload of the cores within a CPU, the operating system can see when there is an opportunity to consolidate workloads to fewer cores. When this happens then the idle cores can be paused or "parked". This analysis happens across all cores of all CPU's in the server. Workloads can be consolidated to fewer CPU's and allows idle CPU's to be parked. When the workload starts to grow the idle cores or CPU's can be awakened as required. This solution will run probably every 100 milliseconds. Again, it seems like little in the way of savings but across a year there probably will be big savings. The CPU will be consuming less power and generating less heat. The rest of the server chassis will consume less power to dissipate that heat. And the computer room or data centre will require less power to cool the server racks.

Running many servers on a single device is not always appropriate. Sometimes we need to allow an operating system and its applications access all of the processing, memory and throughput resources a physical machine can offer. However, do we want to always have a split between the physical machine and the virtual machine?

Microsoft has recently unveiled their long term storage and deployment strategy for Windows. It will first emerge in the light of day in Window Server 2008 R2. A VHD file is a single file that simulates a physical disk in a Microsoft machine virtualisation product. The VHD format was first used in Virtual Server 2005 and Virtual PC. Microsoft published this format with the intention of allowing partners to make use of it. Hyper-V continues to use this format as its virtual disk.

Windows Server 2008 R2 will introduce Native VHD. The idea here is very clever and will move the data centre further down the road of dynamic IT. A server will have a small boot partition that includes enough to start up an operating system and the paging file. A second partition will contain a VHD. The operating system, the applications and the data that will execute on that server will reside within the VHD file. You could even spread the applications and data across a couple of VHD files spread across more physical disks on the server.

Because the substance of this server is in a self contained bubble and it is effectively a virtual machine, it's now very portable and very easy to deploy from a central location. Here are a few scenarios:

- Deploying a new server image will require nothing more than copying a VHD file from a central point to a new piece of dedicated hardware. Operating System deployment will become a one file operation. Microsoft hopes to replace the WIM file we now use for OS deployment with the VHD format. This is a published format so there's nothing to stop others from doing the same.
- Server hardware generally has component support and availability from the vendors for a maximum of 3 years and 3 months. After that point you rely on local resellers to provide hardware maintenance and finding spare parts involves trawling through auction sites and gray markets. In an ideal world the accounting department should depreciate the hardware over three years and after this the hardware should be replaced. In the real world a business wants to get the very most out of a hardware investment. Swapping out hardware involves complications such as migrating applications and data, not to mention dealing with consultants and application vendors. If everything resides in one or a few VHD files then a hardware migration can possibly be a simple process of moving those files from the old hardware to the new. Drivers might not be an issue; the VHD files can be mounted by the Hyper-V hypervisor which will abstract the underlying hardware.
- There are businesses and applications that are near idle most of the year and have peak periods. These peak periods will drive resource consumption to a point where virtualisation will not be an option. With VHD we have a portable bundle of operating system, applications and data. This file format will be able to run on dedicated hardware using Native VHD and is the same format used by Hyper-V for virtual machine storage. In the future the VHD for these peaking applications can reside on a Hyper-V cluster for most of the year. This minimises hardware and power consumption. When peak periods are approaching or when your monitoring solution detects a peak (e.g. an Operations Manager/Virtual Machine Manager PRO Tips management pack) then an action can be fired to move the VHD to a dedicated server. The action would configure Native VHD on that machine to boot from the new VHD. Now the application can serve the peak that's expected. Once the peak subsides the application can be migrated back to a Hyper-V cluster, freeing up the hardware for other purposes, e.g. testing, research and development or maybe other peaking applications.

## Remote Desktop Services

We've already discussed how Microsoft has Terminal Services for session virtualisation and how this solution can be accessed via a web interface using TS Web on the LAN or securely over the Internet via TS Gateway. Microsoft intends to continue to develop this technology.

Microsoft is also aware of the need to supply a Virtual Desktop Infrastructure solution. Microsoft acquired Kidaro a VDI vendor with the intention of including a new VDI solution in Windows Server 2008. They already have TS Web and TS Gateway for providing easy, secure and remote access to RDP protocol sessions on Terminal Services. And they will have a virtualisation platform with Hyper-V V2 and Cluster File System that will allow easy and rapid deployment of virtual machines with no human supervision.

The new Microsoft VDI solution will be combined with Terminal Services and the package will be called Remote Desktop Services. Users and administrators of RDS will only need to be able to use or manage one access platform: TS Web bundled with TS Gateway.

Combined with Hyper-V, VMM and OpsMgr there will be a single pane of management for the server, session and desktop virtualisation platforms on the network.

## Choosing a Virtualisation Platform

So which virtualisation product do you choose? I hope that it's clear now that you are not just picking a standalone product; you're picking a strategy.

VMware is the established brand name and is an excellent product. If you need VMotion to get near 100% uptime then it's currently the solution to adopt. It offers genuine support for a large amount of guest operating systems. RAM oversubscription is an enticing feature that lowers the cost of the host capital investment. ESXi may be available as a free solution but the features of interest to people only come with the expensive ESX and Virtual Center implementations. VMware produces more than just the ESX family. They are leveraging their product to produce complimentary solutions such as VDI and cloud computing.

Xen has had a bit of a cult following up to now but Citrix wants it to become mainstream. It shares a low level architecture with Hyper-V making the two solutions very interoperable; Citrix and Microsoft have always had an intriguing competitive/cooperative partnership that offers lots of options to their shared customer base. Xen is more mature than Hyper-V so it is capable of competing on a feature basis with VMware ESX. Citrix can be credited as being one of the leading players in the history of virtualisation thanks to their key role in the development of session virtualisation. Citrix has developed a family of products for server based or virtualised computing that includes session based, machine based and VDI virtualisation. Their differentiator is the ICA protocol, widely thought of as the best remote computing protocol on the market.

Hyper-V is a version 1 product but thanks to a hugely open testing program, the product is stable and performs very well. Hyper-V is the basis of Microsoft's strategy for machine virtualisation. It is the basis of their Remote Desktop Services platform for session and desktop virtualisation. The Azure cloud services platform, currently in development by Microsoft, is based on Hyper-V. Application virtualisation in the form of Microsoft App-V is the often overlooked technology but it may start playing a bigger role in the future of server consolidation. It's already made an impact in desktop management but we may see a day where server applications are virtualised to make greater use of server licenses and installations.

Hyper-V lacks some of the headline or salesman-friendly features of its competition. Clustering and the provision of disk require a little more administration but Windows 2008 clustering is pretty easy now. Hyper-V also has a relatively small set of potential guest operating systems compared to the competition. Beyond the expected Windows platforms, there's official support from Microsoft for SUSE Enterprise Linux. Other than that we're told that any Xen enabled Linux kernel should run on Hyper-V. Testing has shown that RedHat and CentOS run fine but not as efficiently as SUSE because of the lack of Microsoft provided integration components (to enable better performing drivers).

There are a few things that really draw your interest to Hyper-V. The price is a major factor. A SKU of Windows Server 2008 x64 with Hyper-V is only \$28 more expensive than a SKU without Hyper-V. Although clustering requires either Enterprise or Datacenter editions these costs are easily offset by the licensing features. You can run 4 free licenses of Windows on a host licensed with 1 Enterprise edition license. You can even oversubscribe the licenses with 2 licenses for the host which would allow 8 VM Windows licenses. With Datacenter you get unlimited free licensing for your guest virtual machines. That will be very attractive to people who plan to run large host servers.

The big feature of Hyper-V is management. Being able to manage your virtualisation platform seamlessly with the same suite of tools that you use to manage the rest of your servers and network is very appealing. Adding in hardware partner support for knowledge and automated management further increases this appeal. As of now, it appears the competition cannot match what Microsoft has done.

I'm not going to say you should go one way or another. VMware ESX, Xen and other solutions may be the right way for you to go. What might draw you towards the Microsoft virtualisation platform is their long term strategy with the VHD virtual disk file format. Maybe it's the integrated systems management that features cradle-to-grave, top to bottom control and visibility through a single pane of glass.

A year ago I was a very happy VMware Virtual Infrastructure customer. More recently I had responsibility as part of a decision making group for another virtualisation project. The complete management solution of System Centre was the deciding factor for us. The more automation and expertise we could build into the network, the more responsive we could be to our business.

My advice for you is to download trial editions of any product as well as its management components. Thoroughly test the solution to ensure that it meets your actual business requirements – don't get caught up with the bright lights of features that appeal to the geek like me. This is a business infrastructure platform and it should match your business's technology demands. Research the future strategies for the products to see how it will match your businesses intentions. Consider how you will manage this new virtualisation solution; will you be able to integrate its management with your existing solution? And remember that virtualisation is more than just servers. Consider any future or current needs for other virtualisation technologies such as desktop, session and application virtualisation.

## About Aidan Finn

Aidan Finn has been working with large and complex IT infrastructures since 1996. Working as a consultant, administrator, contractor and service provider, Aidan has seen all aspects of IT working for Internet services companies, financial firms and government in Ireland and abroad. He works mainly with Microsoft technologies including Windows Server, desktops, System Center and virtualisation. Integrating these technologies into the larger network has given Aidan an appreciation for other IT skills such as security, networking, Linux, UNIX and application development/deployment. Aidan is a Microsoft MCSE and an MVP for System Center Configuration Manager. He currently works as the Systems and Infrastructure Manager for C Infinity in Dublin, Ireland.

## About C Infinity

**C Infinity** is an innovative company specialising in managed hosting services, data protection and security. **C Infinity's** success is based on customer focused values, its continuous investment in people and their skills, its adoption of best practice methodologies and processes and its ability to select and partner with leading technology suppliers.

C Infinity's offerings include:

- Data protection: Based on the Iron Mountain solutions, LiveVault and Connected, **C Infinity** offers secure online backup services for servers and desktops. The service is hosted by **C Infinity** in its data centre infrastructure in Dublin, Ireland. The solution offers network efficient and encrypted backups with unmatched performance, reliability and data recoverability.
- Data Security: **C Infinity** provides data encryption security for laptops and desktops using Iron Mountain DataDefense. This centrally and Internet based security solution ensures the owner of their data that it is secure in the event of loss of physical control over the computer.
- Managed Hosting Services: This managed infrastructure allows clients of **C Infinity** to host servers and server applications within the secure **C Infinity** data centre infrastructure. Using a Dublin based data centre with a 100% service level agreement backed up by its history, **C Infinity's** clients can have an enterprise class Internet based server/application deployment with superior hardware, storage and management. The solution offers physical and virtual servers packaged with a service that is based on partnership and a consultative approach.

Our commitment to high quality consultancy, our ability to design, implement, manage and support complex systems, and our reputation for people who deliver business solutions, are both innovative and cost-effective. We are not merely a company that resells other vendors solutions and ideals regardless of quality. We create business solutions using world leading software, hardware and IT infrastructure through carefully chosen partnerships. Using these solutions our clients can avail of enterprise class services that are consumed as required.