




7 Reasons Why VDI with Hyper-V and RDS is Ready for Take Off

Microsoft RDS and Hyper-V is half the cost of implementing virtual desktops with Citrix or VMware technology. But the prevailing opinion has been that Microsoft VDI is behind the other solutions. This eBook explains the 7 technical reasons why Microsoft has caught up to, and in some cases, surpassed the competition with help from partners like Unidesk.



Introduction

Microsoft Hyper-V is rapidly gaining traction with IT organizations serving numerous use cases, geographies and customer segments. According to IDC, Hyper-V has grown to capture nearly one third of the hypervisor market in the last 5 years, as customers look to reduce spending on what has become a commodity.

However, Virtual Desktop Infrastructure (VDI) deployments on Hyper-V haven't grown at the same pace, even though leveraging native, in-the-box Windows Server technology such as Hyper-V and Remote Desktop Services (RDS) costs far less than refreshing PCs or implementing VDI with third-party Citrix and VMware solutions. For example, a [VDI price analysis](#) on 500 desktops shows that Microsoft VDI with Unidesk on Hyper-V is about half the cost of an equivalent Citrix or VMware solution.

The reason organizations continue to pay almost double for VDI is the perception that Microsoft VDI and Hyper-V are not up to par with Citrix and VMware. That may have been true at one time. But that's no longer the case. This eBook offers 7 reasons why VDI with Hyper-V and RDS is now equal to or better than the competition. And why, if you haven't looked at Hyper-V as your platform for VDI recently, it's time to look again.

Reason #1:

RDP with RemoteFX Offers a Rich User Experience

Remote Desktop Protocol (RDP) has come a long way, and now rivals the other leading VDI protocols. Back in the Terminal Services days before VDI, it was admittedly weak. But in typical Microsoft fashion, it has been steadily improved to the point where RDP with RemoteFX now offers a rich user experience. A quick look at RDP's evolution shows this improvement.

In RDP's first incarnations, you couldn't print to local printers; that feature was added in Windows Server 2000, albeit with limitations. The Windows XP RDP client added support for 24-bit color and sound.

RDP 6.0 that shipped with Windows Vista added the first support for multiple monitors. When version RDP 6.1 shipped as part of Windows Server 2008 and Vista SP1, full local printing capabilities were delivered, without the need to install printer drivers on the server.

RDP 7.0 in Windows 7 and Windows Server 2008 R2 really upped the ante on the user experience. Multi-monitor support was improved and you could even get the Aero Glass experience. Windows Media Player redirection was also added for enhanced video.

Then, RemoteFX was introduced as an enhancement to RDP 7.1, taking RDP to a new level.

RemoteFX Overview

RemoteFX enables a richer user experience by providing a 3D virtual adapter, intelligent codecs, and the ability to redirect USB devices in virtual machines. The RemoteFX codec is designed to provide high quality video, photos and text over limited bandwidth networks in real time by compressing the data while reducing latency. The RemoteFX codec can run in software on the CPU or on the client GPU, or on a custom chip on a card or embedded in thin client hardware.

With RemoteFX, graphics are rendered on the host computer (the RDP server) instead of the client device, and applications use the host's GPU and CPU to run at full speed. The data on the screen is compressed and then encoded by the RemoteFX encoder. Then the compressed bitmaps are sent to the client, where the RemoteFX decoder decodes it. The experience is almost the same as running the applications on a physical PC. RemoteFX also enables the single GPU on a Hyper-V server to be shared by multiple virtual desktops on multiple clients, including thin and ultrathin (zero) clients.

The first version of RDP with RemoteFX shipped with Windows 7 SP1 and Windows Server 2008 R2 SP1. It provided the virtual GPU that is the basis for all RemoteFX capabilities. The vGPU gives virtual machines access to the advanced functions of the hardware GPU, such as 3D rendering. Another important part of RemoteFX is the USB redirection component. This is what makes it possible for you to access USB devices that are connected to your local client device when you're working on your remote virtual desktop.

In October 2013, Microsoft released Remote Desktop apps for iOS and Android devices that use RemoteFX to provide a rich virtual desktop experience on those devices. The apps are now available in the [Apple App Store](#) and [Google play store](#).

The latest version – RDP 8.1 – ships with Windows 8.1 and Windows Server 2012 R2. This version supports as many as 4 monitors per virtual machine at 1920 x 1200 resolution. If you're willing to drop resolution back to 1280 x 1024, you can run up to 8 monitors per VM.

RDP 8.1 is also touch-enabled to align with the current trend toward mobile computing, the increased use of tablets, the availability of touch-enabled monitors, and the adoption of Windows 8.1 (or Windows 10) as a touch-centric operating system. It has full support for multi-touch, so that users can use a touch-enabled device such as a tablet to interact with their virtual desktops. That's an important factor in making the client-side experience as seamless as possible and makes the modern UI much more user-friendly with Windows 8/8.1/10 desktops.

Lastly, performance has been greatly improved. The first versions of RDP were a bit laggy, even over a fast LAN. Even when RemoteFX first shipped, it worked only on a LAN. However, RDP 8.1 significantly improves the performance of virtual desktops, even across a WAN. It can detect the connection speed automatically and adjust to provide the best VDI performance.

The adaptive encoding feature introduced with RDP 8 also improves performance. It chooses the best codec for the type of content being transmitted, and can use hardware accelerators if they are available on the client system. In RDP 8.1, this feature gets even better. The H.264 video compression support now includes images as well as video, so if there is an H.264 hardware accelerator module detected, for example, RemoteFX can offload both video and image processing to it. The CPU will continue to decode text. This results in much higher frame rates. The virtual desktop experience through Microsoft Surface RT and similar devices will be greatly enhanced as a result.

As you can see, RemoteFX and the many other enhancements described above have made RDP a viable protocol for VDI that's worthy of consideration.

Reason #2: Microsoft RDS Scales with GUI-Based Management

So RDP with RemoteFX has become significantly better, and now rivals the other leading protocols used for VDI. But the remoting protocol isn't everything. What about manageability of Remote Desktop Services (RDS) and its RD Connection Broker role that ships with Windows Server 2012? How easy is it to manage desktops and apps when you're deploying Microsoft VDI on Hyper-V at any kind of scale?



Until recently, the answer was "not easy." As you scaled RDS toward 500 users, it became harder to manage. Missing was a real user-centric view of RDS broker connections. Adding desktops to RDS collections and updating all of the VMs in a pooled virtual desktop collection were manual processes. Extensive use of PowerShell scripting was required, creating a serious implementation hurdle for small and mid-sized businesses with lean IT teams.

This is no longer the case. With the availability of [Unidesk® on Hyper-V](#), managing RDS and updating pooled virtual desktop collections can now be driven entirely through the graphical user interface of the leading virtual desktop management platform. The need for PowerShell scripts has been eliminated, enabling lean IT organizations to scale RDS to 500 users and beyond.



Unidesk enables you to create and edit RDS collections from an easy-to-use web-based management console, as shown here.



You can then use Unidesk to provision Windows 7/8/10 and Windows 2008/2012 Server desktops – persistent or non-persistent – into any RDS collection, as shown here.



Updating all of the VMs in each of the collections is fully automated by Unidesk's Composite Virtualization® layering technology, which enables the Windows OS and any app to be delivered to many VMs – without installation – as virtual disks (VHDs). Patch the one Windows OS layer that is used as the basis for all desktops, or update any of the application layers – all desktops that are using those layers will be automatically rebuilt based on the maintenance schedule you define.

If the desktops are provisioned as persistent, all user personalization – local settings, user-installed applications, plug-ins, and data – will remain intact, no matter how often the underlying OS and application layers are changed.

If the desktops are provisioned as non-persistent, the writable Personalization layer that is unique to each VM will be erased, removing all customizations so that the next user starts with a fresh, clean experience.

With Unidesk, simple, scalable management of RDS collections is no longer an obstacle to deploying Microsoft VDI on Hyper-V.

Unidesk also offers similar capabilities when deploying Citrix XenDesktop as your broker on Hyper-V, integrating with the XenDesktop SDK to provision layered desktops into XenDesktop catalogs and groups.

Reason #3:

Hyper-Converged Systems for Hyper-V Offer Infrastructure Scale and Simplicity

IT infrastructure has traditionally been silos of compute, network, and storage resources. In these legacy environments, the storage group handles the purchasing, provisioning, and support of the storage infrastructure and maintains the relationship with the storage hardware vendor. The same situation exists for the server and network groups. The burden of integration falls on IT, with the three groups being called upon to coordinate deployments, changes and updates. And there is no workload better at exposing flaws in any one of the three components than VDI, which relies heavily on all three.

The concept of converged systems combines two or more of these infrastructure components as a pre-engineered solution that works “out-of-the-box,” relieving IT of the integration burden and mitigating VDI project risk.

New hyper-converged systems that support Hyper-V take the concept of convergence to the next level. While converged systems are separate components engineered to work well together, hyper-converged systems are modular systems designed to scale out by adding additional modules. These systems are most often designed around storage and compute on a single x86 server chassis interconnected by 10 GB Ethernet.



The differences between a hyper-converged system and servers with a bunch of disks are engineering and software. Hyper-converged solutions leverage improvements at the storage controller software layer to enable the systems to scale out. The more appliances you add, the greater the performance and capacity. Instead of scaling up by adding more drives, memory, or CPUs, you scale out by adding more appliance modules.

In addition to the simplified architecture, hyper-converged systems offer simplified administration. Instead of having a set of applications and a team to manage your storage

array, a team to manage the network, and a team to manage the server hardware, one team (or in some environments one person) can manage the complete hyper-converged stack.

Here are some of the capabilities and benefits that hyper-converged systems such as [Gridstore](#) and [Nutanix](#) now bring to VDI on Hyper-V:

- **Elimination of Storage Complexity:** Network storage is presented as one large container of virtual disk storage, eliminating the need to manage volumes, LUNs, RAID groups and a dedicated storage area network (SAN).
- **Easy Scalability:** As compute and storage needs grow, new hyper-converged systems can simply be added, resulting in linear scaling without needing to tweak the host cluster or SAN. No rip and replace of old storage systems. No waiting for the old systems to be migrated. No need to provision a LUN within the appropriate group of disks. No configurations for performance and fault tolerance to meet storage SLAs.
- **Efficient Copy and Zeroing Operations:** The Hyper-V ODX (Offloaded Data Transfers) feature allows Windows to place the responsibility of data copy and zeroing operations on converged infrastructure, resulting in as much as 10x better performance than equivalent client level operations.
- **Thin Provisioning:** The Hyper-V TRIM capability enables converged infrastructure solutions to reclaim capacity the moment it is available, with minimal impact to system resources, making thin provisioning more efficient.
- **Failover Clustering:** Failover is integral to VDI availability and requires storage to be shared across all hosts. Hyper-converged infrastructure makes this easier than legacy, siloed IT infrastructure. Hyper-converged solutions are typically based on shared storage pooled from SSDs and HDDs that are local to the nodes in the cluster. All hosts are presented with a single, scale out, distributed file system, enabling virtual machines to access storage on any host within the cluster. Any host within the cluster can contribute cluster capacity and performance. When capacity or performance is expended, IT simply hot-adds new appliances for linear scale out of capacity and performance with no downtime. If you need to balance virtual desktop workloads across cluster hosts, you can leverage the Hyper-V Performance Resource Optimization (PRO) capability in System Center Virtual Machine Manager (SCVMM) to configure load balancing rules. These rules determine when and where virtual desktops must move within the cluster and govern physical resource distribution across VMs. VM data is appropriately moved within the cluster following the VM it is tied to, without any manual intervention.
- **Quick Installation:** Installing Windows Server 2012 R2 with Hyper-V on hyper-converged systems across multiple hosts should complete in less than an hour. Setup is fully automated including joining hosts to Active directory, creating a Failover cluster and optionally registering hosts with SCVMM.

Reason #4:

Application Layering Solves App Delivery “Bottleneck”

As Gartner explains in [“The Secret Bottleneck of VDI,”](#) application delivery is one of the biggest challenges facing any VDI project. The traditional application delivery methods that have been available on the Hyper-V platform all have issues:

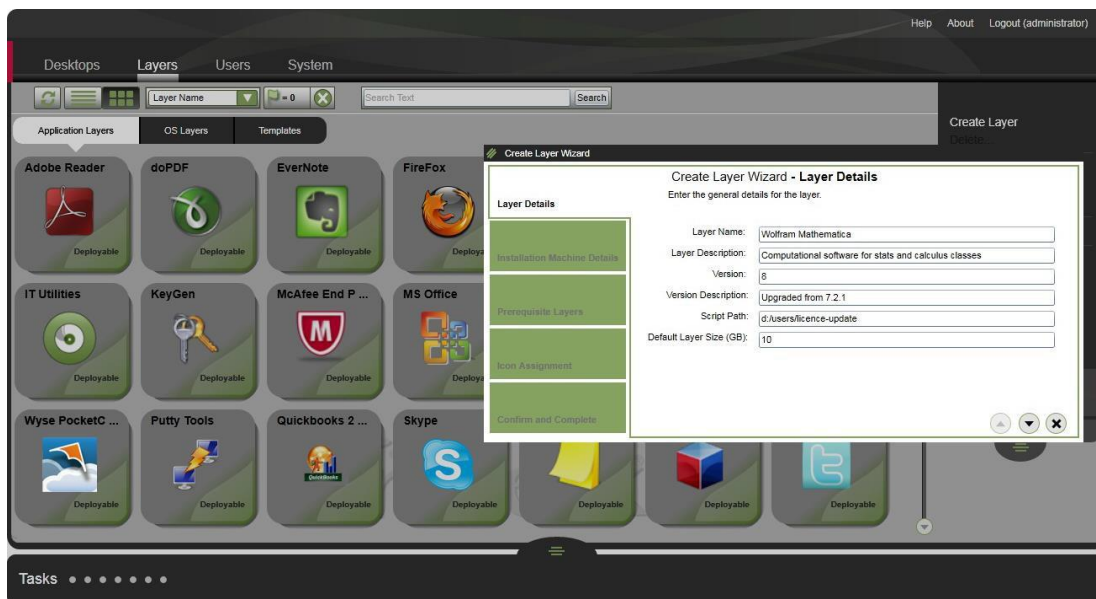
- **PC-focused Application Management** tools were designed for physical desktops with local CPU and local disk. The impact of using these agent-based tools to repeatedly install the same patch or update on many VMs has a severe impact on shared resources like VDI servers and centralized storage. Furthermore, agent-based updates don’t always “take,” resulting in open service tickets and escalations to Tier 2 and 3 IT administrators that increase OpEx.
- **Server-Based Computing (App Publishing)** essentially duplicates the VDI environment, with one set of virtualized servers to host the desktops and another to host the applications. Most organizations that are deploying VDI are trying to move away from legacy server-based computing and don’t want a duplicate environment. There are also many applications that don’t play nicely with app publishing.
- **Application Virtualization** isolates applications in their own protective “bubbles,” effectively hiding them from Windows and other apps. This approach was intended to let applications that would normally conflict with each other coexist on the same desktop. It was never intended to deliver every application in your environment. So it’s no surprise that trying to virtualize every app usually fails due to limited application compatibility, insufficient IT staff time and expertise to master the complex packaging process, and the inability for sandboxed apps to cross-communicate.

Application Layering is the newest method of application delivery in VDI, and it is [now available on Hyper-V](#). As Gartner points out – and the [many customer comments confirm](#) – application layering has clear advantages over the other methods of application delivery:

1. **Layering is easy.** Where other methods require hours, even days, of packaging, scripting, and pre/post-processing time and expertise, layering is as easy as installing an app on your home PC. You are given a prepped and ready VM with just a clean version of Windows. You run the application’s setup procedure, click Finalize, name the layer, and start assigning it to desktops. Layering enables Tier 1 IT staff to become app packaging experts.
2. **Layering is compatible with all applications.** Layering technology captures every file and registry key that is different from the base Windows OS, and stores the application

“layer” as a read-only virtual disk (VHD) that can be mounted and shared by many VMs. Apps are never installed again. After app layers are assigned to desktops in various combinations, the layering engine merges the files and registry keys in each layer into a unified file system to form a perfect virtual C: drive. This approach – which operates above the hypervisor but just below Windows and occurs before a desktop boots – enables any application to be layered. System services, apps with boot-time drivers, ring 0 applications, and apps with complex installation procedures are all compatible with application layer.

3. **Layered apps are not isolated.** Layered apps behave as if they were locally installed. Windows and other apps aren’t even aware that the app is actually a read-only, shared VHD, layered into a C: hard drive. All files, data, and registry settings are stored in the usual places, and can be seen by native Windows tools such as regedit and Add and Remove Programs. Apps and plug-ins cross-communicate and share data, just like they’re supposed to.
4. **Layered app delivery is reliable and uses fewer resources.** Unlike traditional physical desktop management solutions that use an agent to install an application 100 times to deploy it to 100 VMs, layering requires only one install. Once an app is packaged as a layer, the same shared VHD is mounted by all 100 VMs. This “package once, deploy to any number of desktops” innovation leverages virtualization to eliminate the install failures and help desk escalations common with legacy agent-based tools.



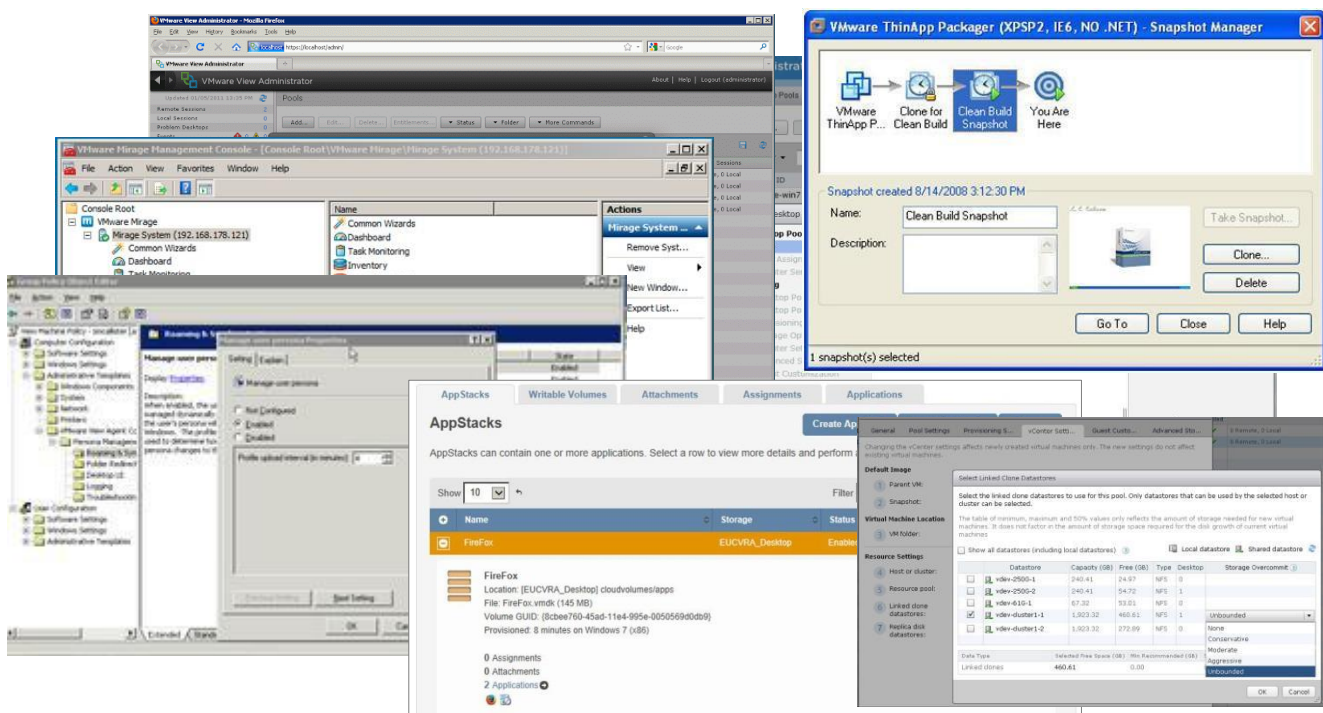

Application layering solutions such as Unidesk enable all applications to be captured as read-only VHDs that can be mounted and shared by many VMs on Hyper-V to greatly simplify and streamline VDI application delivery

Reason #5: Helpdesk & Tier 1 IT Staff Can Manage VDI Day-to-Day

Managing desktops of any kind has always been a challenge for IT. Managing virtual desktops has been an even bigger challenge. Administrators have had to cobble together a collection of independent tools and technologies for day-to-day desktop provisioning, image management, storage optimization, application delivery, Windows patching, break/fix, and personalization.

The complexity of this “silosed” management approach and the sheer number of management consoles that must be mastered are why VDI deployments have been difficult to scale and transition to less experienced IT staff.

The graphic of the VMware Horizon View VDI management stack on the right illustrates this complexity. When you look at it from the IT administrator’s perspective, as shown below, VDI management looks even more daunting.



The availability of the Unidesk virtual desktop management platform on Hyper-V, its innovative layering technology, and its tight integration with RDS paints a much different picture.

[Microsoft has chosen Unidesk as its VDI management partner](#) because Unidesk offers VM provisioning, application layering, full personalization (including user-installed apps), patch management, disk space reduction, and scalability enhancements to RDS for collection management – all in one easy-to-use, single pane of glass management solution.

VMs can be provisioned simply by selecting which application layers should be combined with the clean Windows OS layer. Persistent and non-persistent desktops can both be provisioned using the same set of shared layers for maximum storage efficiency and streamlined patch management. Desktops can be automatically added to RDS collections without PowerShell scripting.

The simplicity offered by Unidesk on Hyper-V enables VDI to be deployed and managed by IT generalists and helpdesk staff. VDI becomes more agile and more responsive to business needs. Project risk is mitigated. Best of all, from the IT administrator's perspective, VDI now looks like this:



Reason #6:

Hyper-V Data Deduplication Reduces Storage Costs

Hyper-V now offers built-in data deduplication that radically reduces space requirements for both non-persistent and persistent virtual desktop use cases – a feature that was previously available only on expensive storage arrays and is not available on VMware vSphere.

Data deduplication is a method of reducing storage needs by eliminating redundant data. Only one unique instance of the data is actually retained on storage media, such as disk or tape, and redundant data is replaced with a pointer to the unique data instance.



Microsoft first introduced data deduplication for NTFS file systems in Windows Server 2012. With the release of Windows Server 2012 R2, this feature was greatly improved to offer significant benefits for virtual environments, especially virtual desktop workloads.

The use of dedupe with virtual hard disks (VHDs) in VDI environments can result in substantial capacity savings as well as more consistent, higher-performance I/O response. Fixed VHD files typically have a common Windows operating system and other bits in common. As VDI and the number of VHDs (including VM template libraries and ISO file shares) scale, this common data gets stored over and over.

Microsoft's deduplication feature saves significant disk space by storing the common bits only once. An added benefit is that when a volume with live VHDs is deduplicated, forcing all I/O to hit the common blocks, these blocks become "hot." Windows will move these frequently used "hot" blocks to Tier 1 storage - typically solid-state or cache memory. The result is significantly greater performance, since the majority of I/O will now be served from your highest-performing storage.

Reason #7:

Hyper-V Dynamic Memory Increases VM Density

The Dynamic Memory improvements for Hyper-V in Windows Server 2012 enable organizations to attain higher VM consolidation numbers and improved reliability for restart operations compared to VMware's Memory Overcommitment feature. This can lead to lower costs, especially in pooled VDI environments that have idle or low-load virtual machines.

VMware uses a balloon driver and paging to allow memory to be overcommitted. The maximum amount of memory assigned to a VM is described in the BIOS Memory Map, and is available as soon as the VM is started. If there is not enough physical memory on the vSphere server to satisfy demand, the hypervisor will use paging.

Once the OS is booted, the balloon driver will communicate with the hypervisor to control how much physical memory is assigned to the VM. By default, Windows Cache Manager will consume all of the memory in the system. But the balloon driver prevents this from happening by consuming memory that would otherwise be used by the Cache Manager, giving it back to the hypervisor instead. When a VM needs more memory, it can get it back from the hypervisor.

With VMware Memory Overcommitment, the OS has to track the maximum amount of memory, whether it's used or not. So, if you set your max memory really high, you're going to pay a price in resource utilization.

By comparison, Microsoft Dynamic Memory enables administrators to set their own memory thresholds and allows VMs to be prioritized in terms of memory usage. Consequently, when memory contention occurs, high-priority VMs receive memory first. The BIOS Memory Map contains the memory in the "Startup RAM" field. Hyper-V has minimal need for paging, which occurs only during a reboot of a VM, and only if memory can't be freed up any other way.

If a VM needs more memory than the Startup RAM, the Dynamic Memory Hot Add mechanism is used. The hypervisor gives a block of memory to the VM, and signals to the OS that memory has been added. The OS then starts tracking and using the added memory, using a balloon driver (like vSphere) to communicate with Hyper-V.

With Hyper-V's approach, the OS has to track only the Startup RAM and any additional memory that's actually used. Hyper-V Dynamic Memory minimizes the amount of memory the OS has to worry about by using Hot Add, making it much easier on the balloon driver, and freeing up server resources to host additional VMs.

The ability to make run-time configuration changes using Hyper-V Dynamic Memory can reduce downtime, provide increased agility, and increase VM density - all key considerations when implementing VDI.

Summary

Deploying virtual desktops and offering Windows app mobility no longer requires third party software from Citrix or VMware. Day-to-day VDI management no longer requires high-level IT staff resources and costly IT consultants. As this EBook shows, enterprises of all sizes can now leverage native Windows Server technologies such as Hyper-V and Remote Desktop Services, together with partner solutions such as Unidesk, to achieve the holy grail of VDI:

1. Simplicity to mitigate VDI project risk
2. Higher quality of service for end users
3. Greater agility for businesses
4. Lower cost than buying new physical PCs.