Taking Advantage of Modern Memory and Storage with Windows R2

SPONSORED BY

Windows IT Pro

In 2003, Microsoft released a brand-new version of Windows Server that built on the radical changes introduced in Windows 2000 Server, such as a true directory service in Active Directory (AD) and a new interface based on the Start menu. Windows Server 2003 introduced several features and improvements: a huge new version of IIS (version 6.0), integrated .NET, Volume Shadow Copy Service (VSS) for application-aware backups, big changes to Active Directory, and more. Organizations quickly adopted this "second version" of the Windows 2000 Server OS, and time has proven it to be a rock-solid OS environment. Also in 2003, the movie Chicago won the Academy Award for Best Picture, and Apple released the third-generation iPod with a new backlight for its monochrome screen.

More than a decade later, Windows Server 2003 is still heavily deployed in many organizations, often on the same hardware it was initially installed on. In fact, the actual number of deployments is startling. As of July 2014, there are more than 20 million Windows Server 2003 deployments, and those deployments are evenly split between physical and virtual environments. Although Windows Server 2003 was available as a 64-bit OS, it wasn't a widely adopted option, which means most organizations are running the 32-bit version. Consider the requirements of Windows Server 2003:

- A minimum of 128MB of RAM (256MB recommended)
- 1.5GB of disk space
- Up to 8 processors supported

After July 14, 2015, Microsoft will no longer issue critical updates for the OS, and any company still running it will not only be unsupported but will be unable to pass any kind of audit that requires the OS to be supported

It's hard to find a mobile phone OS that can run within those specifications today, yet that is the OS many organizations rely on. It's an OS that was architected before the advent of today's large-scale servers, before the cloud, and really before virtualization—and the reality is that on July 14, 2015, all companies will need to have migrated from it because it will become an unsupported OS.

Even today, Windows Server 2003 is in an extended support phase, which means it receives limited free updates, restricted to those that impact security, and only paid support. Even this support ends on July 14, 2015. Microsoft has a countdown clock on its Windows Server 2003 end-of-life website, along with some helpful migration resources. After July 14, 2015, Microsoft will no longer issue critical updates for the OS, and any company still running it will not only be unsupported but will be unable to pass any kind of audit that requires the OS to be supported, which might also mean that company is vulnerable to fines and litigation.

My goal with this whitepaper is to entice you to migrate sooner rather than later to Windows Server 2012 R2 so that you can take advantage of exciting new capabilities and also maximize your investment in modern hardware, specifically memory and storage. Beyond understanding the ramifications and urgency of moving away from Windows Server 2003, you're in for a treat when you start looking at Windows Server 2012 R2.

Enter Windows Server 2012 R2

In 2012, Microsoft introduced Windows Server 2012, and the company released Windows Server 2012 R2 a year later, building on many of the huge investment areas of Windows Server 2012. Consider the requirements and some limits of Windows Server 2012 (which also apply to Windows Server 2012 R2):

- A minimum of 512MB of RAM (with 2GB being a more realistic minimum), with support for up to 4TB
- 32GB of disk space
- Up to 64 processors supported, including up to 640 logical processors
- Support for Unified Extensible Firmware Interface (UEFI), which enables features such as SecureBoot, providing a secure handoff from the physical hardware to the OS

The capability to fully leverage the hardware of modern datacenters—including the latest networking and storage fabrics—is vital, but it's really the capabilities of the OS that will drive its adoption. When migrating to Windows Server 2012 R2, it's critical to take the time to understand what's possible with the OS and not just emulate the way you work in Windows Server 2003. Only then will you see the real value of your OS investment.

Listing all the new capabilities introduced between Windows Server 2003 and Windows Server 2012 R2 would require multiple books. However, there are key areas that nearly every organization should evaluate and consider.

- ***Virtualization***—Nearly every organization has a virtual-first mentality these days, and Windows Server 2012 R2 offers Hyper-V, one of two x86 server virtualization solutions in Gartner's magic quadrant. Hyper-V provides industry-leading features such as software-defined networking, shared storage, 64 vCPU virtual machines (VMs) with NUMA support and 1TB of memory, full

VM mobility even outside of clusters, extensible switch architecture, dynamic storage resize, and more. Additionally, Hyper-V powers Microsoft Azure, giving full compatibility between VM workloads running on-premises on Hyper-V and workloads running in Microsoft Azure.

- ***Desktop Virtualization***—Remote Desktop Services (RDS) provides a rich set of capabilities related to providing virtualized desktops, which could be session-based environments running on a shared OS, a Virtual Desktop Infrastructure (VDI) scenario with each user connecting to a client OS running in a VM, or simple application publishing, which is a great solution for mobile devices. RDS has full gateway functionality, enabling remote sessions to operate over HTTPS, driverless printing, and even GPU virtualization via RemoteFX, enabling rich graphical applications to run in virtual desktop environments.
- ***File Services***—People often think of Windows file services as simple shares on which to store users' Office documents, but the reality is that Windows has a powerful set of capabilities related to file services in File Server Resource Manager (FSRM). FSRM enables real-time quotas to control how much space is used, real-time screening to control which type of data is saved, detailed reports, and file classification that can scan file content and classify it based on the type of content, with classifications being used to trigger actions such as application of rights usage policies and encryption. BitLocker is also a core feature, providing volume-level encryption that ensures data protection even in locations where servers can't be physically secured.
- ***Active Directory***—The new OS brings not only huge improvements to the core Directory Services component, including fine-grained password policies and the ability to undelete objects, but also new components such as Active Directory Federation Services (ADFS), enabling single identities to be used in a secure fashion between trusted organizations.
- ***Web Application Proxy***—This feature not only provides a reverse proxy to securely offer services to the Internet but also integrates with ADFS to perform pre-authentication and enable single-sign on (SSO) across multiple services for end users.
- ***Management***—For administrators of Windows Server 2003, Windows Server 2012 (and Windows Server 2012 R2) is a completely different experience. In the new OS, the graphical shell and management tools can be removed and added to a server at any point in the server's lifecycle, enabling reduced patching and

reboots for maintenance. Management is now performed through a multi-server management tool called Server Manager and through PowerShell. You can use PowerShell as a command-line interface and for scripting to automate any task not just in Windows but in nearly all Microsoft and partner solutions.

The above list enumerates just some of the core functionalities, but certainly Windows Server is a very different animal than it was a decade ago. I doubt there are many organizations that wouldn't benefit from the improvements Microsoft has made to the OS over the past 10 years.

Maximizing Memory Utilization

There are very few applications that could come close to utilizing Windows Server 2012's 4TB memory maximum—even SQL Server, unless you have a database that needs 4TB. In that case, the best approach would be to break the memory into smaller Non-Uniform Memory Access (NUMA)-sized VMs and have multiple SQL Server instances. With virtualization, 4TB of memory support is very important because it equates to great VM density on servers as processor cores, disks, and network bandwidth can be shared. A single Hyper-V VM can have up to 64 vCPUs and 1TB of memory. When you consider VMs of this magnitude, NUMA becomes essential.

What is NUMA? Put simply, a server has multiple processor sockets, and typically each of these processor sockets is directly connected to a certain number of memory banks. The memory banks that are directly connected to processors are considered local and give the fastest response. The memory and processor that are directly connected are known as the NUMA node, and where possible workloads should be NUMA-aware so that when performing processing, the memory used is within the same NUMA node as the processor cores perform the processing. Hyper-V passes through the NUMA topology of the physical system to VMs, enabling NUMA-aware applications such as SQL Server and IIS to intelligently leverage resources in the most optimal way.

Hyper-V also features Dynamic Memory, which enables a minimum and maximum amount of memory to be specified for each VM, as Figure 1 shows. The amount of memory actually used by processes within the VM is then monitored, and if the amount of available memory drops below a defined threshold, and if memory is

physically available within the host, then additional memory is added to the VM. Conversely, if a VM no longer needs the amount of memory it has, and other VMs could benefit from it, then memory can be removed from a VM and reallocated through a process called "ballooning." What this means to an environment is increased density as memory is removed and added as required in an intelligent manner.

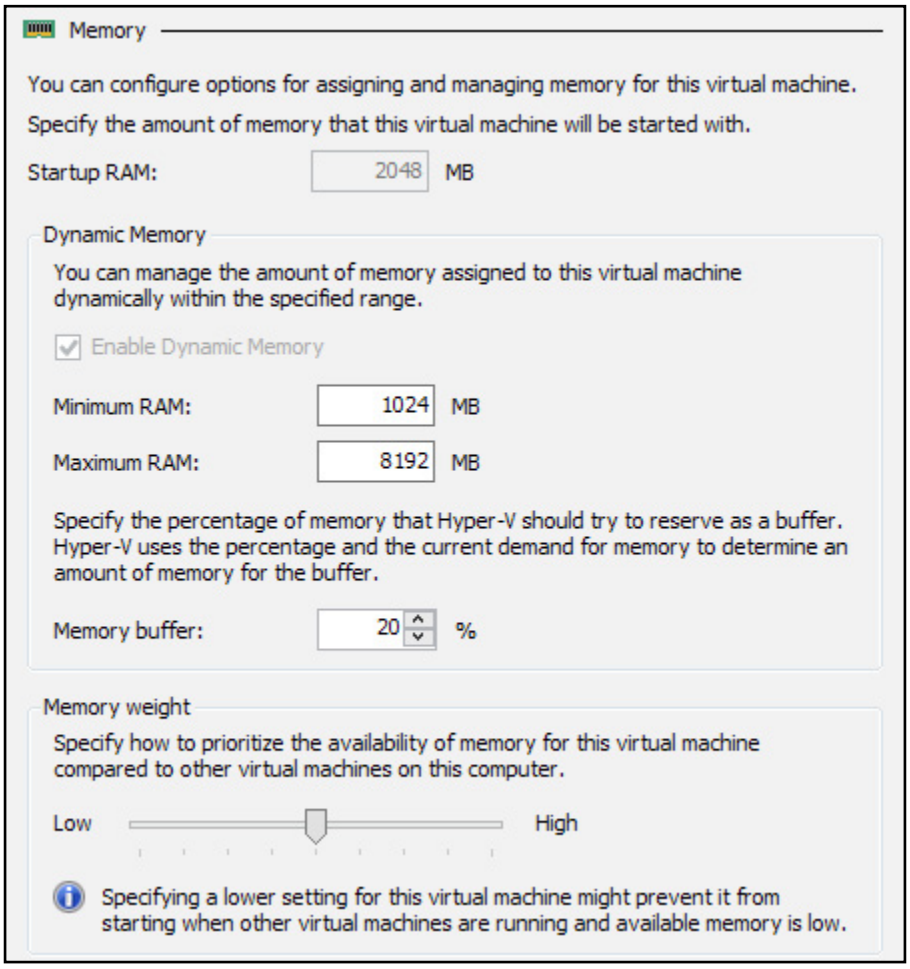

Figure 1: Configuring memory for a Hyper-V VM

Memory has become denser and more advanced, with enterprise memory featuring error-correcting code (ECC) capabilities, which enables the memory itself to detect internal data corruption and remediate the corruption. It's common to see ECC memory in environments where no type of data corruption can be tolerated.

Windows Server takes these capabilities to the next level by reading the telemetry of the ECC memory, which includes information about the corrected errors and the frequency of the errors, to predict impending failures of memory pages and take the pages offline before a failure occurs. This functionality is called Predictive Failure Analysis (PFA) and is implemented as part of the Windows Hardware Error Architecture (WHEA). The net effect is a reduction of the probability of a memory fault impacting the availability of your services.

If you ask the average person to name the greatest advancement in storage over the past decade, the most likely answer is solid-state drives (SSDs)

Leveraging Storage Improvements

If you ask the average person to name the greatest advancement in storage over the past decade, the most likely answer is solid-state drives (SSDs). SSDs are a radical re-architecture of persistent storage, moving away from mechanical spinning platters to a solid-state solution that uses flash memory with no moving parts. In a traditional disk, most of the performance penalty is the movement of the data-reading head to the correct location above the platter. This is known as the seek operation, and when reading non-sequential data, there are continual seek operations that significantly impact performance. With an SSD, there is no seek time, there is no head to move, all the data is available instantly, and the only limitations are the controller, the bus type, and so on. These are the reasons SSDs are so much faster than traditional hard drives—potentially hundreds of times faster for random I/O reads. Additionally, because they have no moving parts, SSDs use significantly less power and run much cooler, which are attractive qualities in desktops, tablets, and servers as datacenters strive to be more energy-efficient.

Why do you need Windows Server 2012 to take advantage of SSDs? Surely you could just put an SSD in a Windows Server 2003 system and you would get faster I/O and use less power. The answer is yes, at least initially, but Windows Server

2008 R2 and later are SSD-aware and actually perform a number of optimizations when the OS detects an SSD. Optimizations include the following:

- Defragmentation is automatically disabled on SSDs. Fragmentation is actually a good thing on SSDs because it levels out wear on the SSD and helps prolong its life. Because there is no seek time on SSDs (which is typically the problem of fragmented data), there is no negative effect.
- Superfetch, which normally defragments the files used for startup, is disabled. Prefetch is also disabled.
- Readyboost, which enables fast USB drives to supplement memory, is disabled; however, this is not required with SSDs.
- Partitions are automatically aligned to the SSD structure.
- BitLocker in Windows Server 2012 enables encryption only of used space, which avoids operations on unused areas of the SSD.

TRIM support is critical to the ongoing of an SSD

A more important feature that's critical to the ongoing optimal performance of an SSD is TRIM support, which is also part of Windows Server 2008 R2, Windows 7, and later. To grasp the value of TRIM, some understanding of the structure of an SSD is necessary. Essentially, an SSD is made up of cells that each store a bit of data, and sometimes more than one bit of data in multi-level cell drives. These cells are organized into pages that are typically 4KB in size, and the page is the smallest unit that can be written to and read from. The pages are organized into blocks that are typically 512KB in size, and a block is the smallest unit that can be erased. You can't overwrite data in a page; its content must be removed first, which means erasing the entire page. What this means is that when data is deleted before fresh data can be written, the entire content of the block must be copied to a memory cache, the new data written to the cache, the block erased, and then the data copied back, which slows down performance. Figure 2 shows this process of copying the content to memory, wiping the block, then writing back. SSDs have powerful algorithms to distribute the writes throughout the SSD to level wear and reduce the impact of the overwriting.

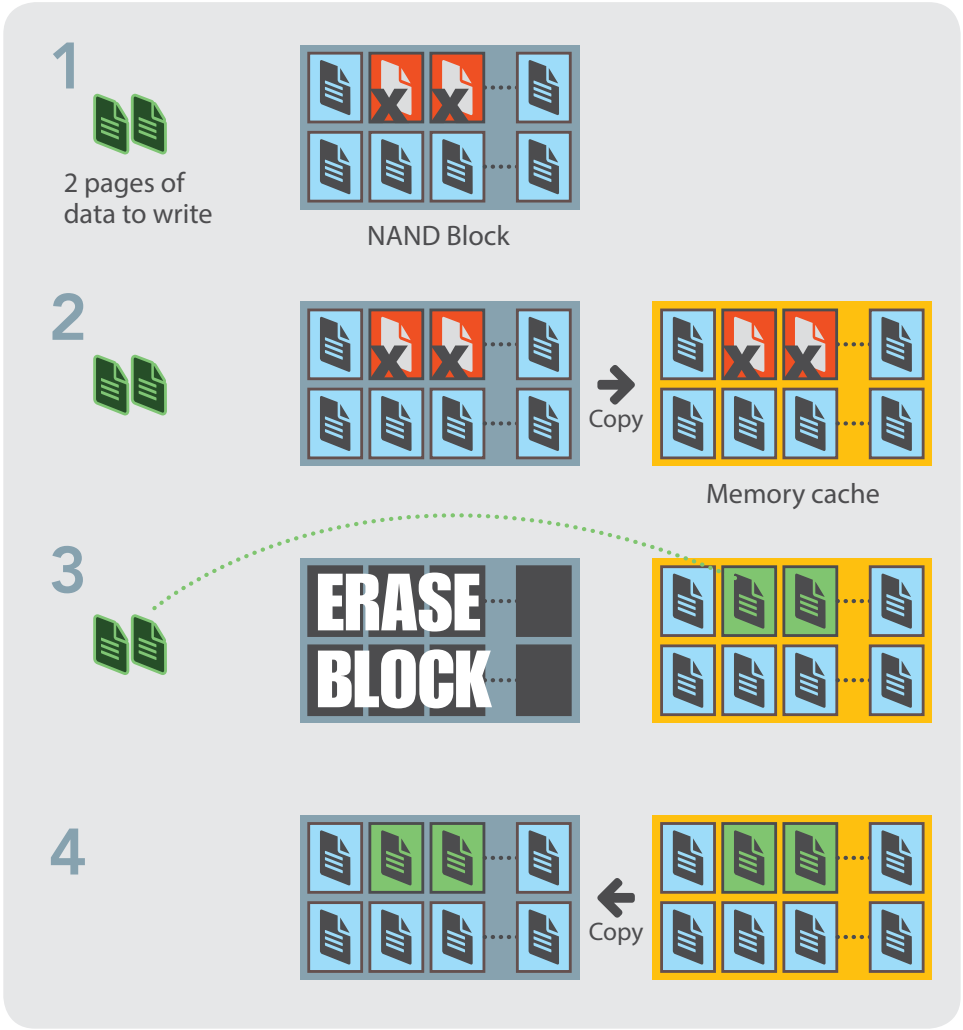

Figure 2: The process of overwriting data

This is where TRIM is leveraged. When data is deleted on a normal file system, no action actually takes place on the physical disk; the area on the disk is just marked as available. However, in the case of an SSD, it's advantageous for the OS to tell the SSD when areas of the disk are deleted, because it allows the SSD to perform the delete ahead of future writes, thus removing performance impacts. Remember, because writes can occur at a page level, it's only an overwrite that requires the complete block to be erased. This is exactly what TRIM enables: The OS notifies the

SSD when areas on disk are deleted, and the SSD then deletes the pages, as Figure 3 shows. Windows Server 2003 has no concept of TRIM, so over time, as all blocks are filled, the performance of the SSD would degrade, reducing the benefit of the SSD investment.

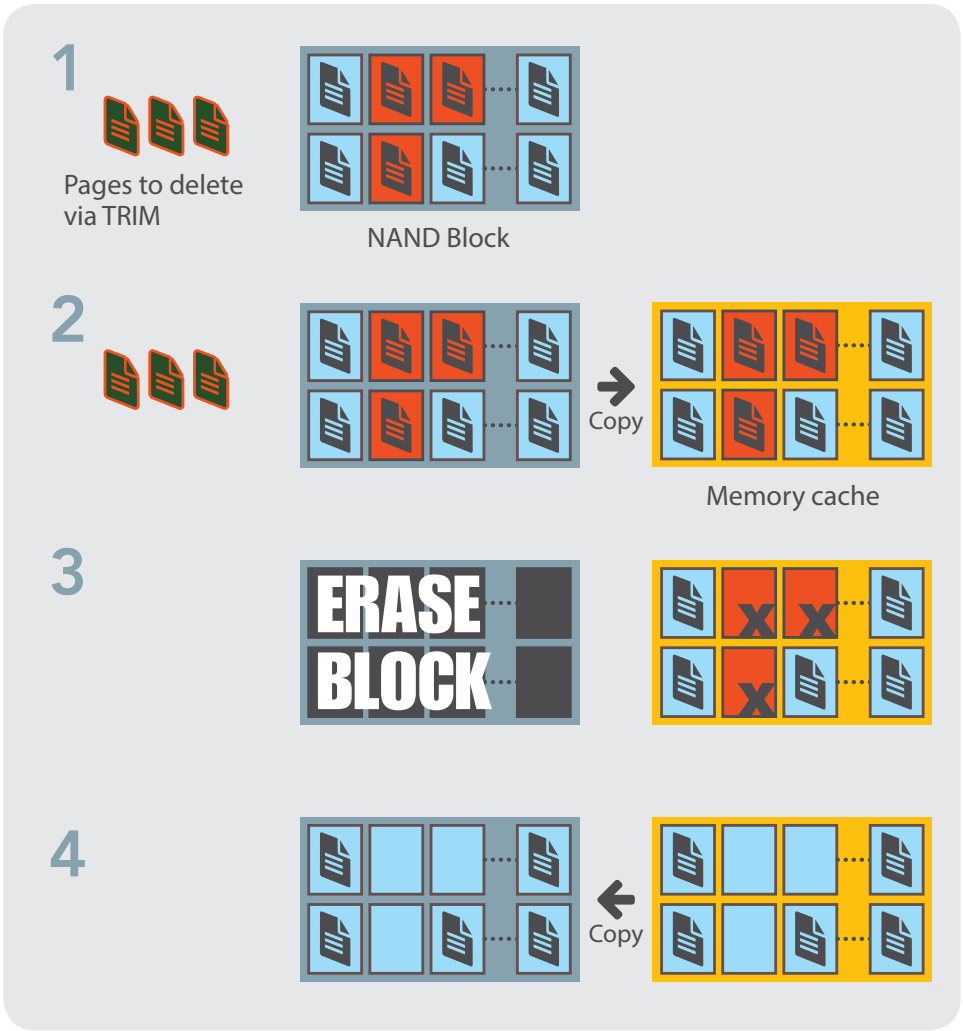

Figure 3: Pages cleared, ready for future data writes when TRIM is used

One more key area where Windows Server 2012 R2 offers huge benefits for SSD drives is tiered Storage Spaces. Storage Spaces was first introduced in Windows Server 2012 and provided the ability to group locally attached disks into a Storage Pool. From that pool, virtual disks could be created that supported features such as thin provisioning and various types of resiliency such as parity and mirroring. In Windows Server 2012, there was no concept of different types of disks, but this changed with Windows Server 2012 R2, which added tiering and enabled a distinction between HDDs and SSDs in the pool. This evolution enabled a solution with very large amounts of storage using HDDs and automatic movement of the most used blocks to the SSDs, giving the highest performance. Each virtual disk can be allocated a specific amount of space from the HDD tier and a certain amount from the SSD tier, as Figure 4 shows.

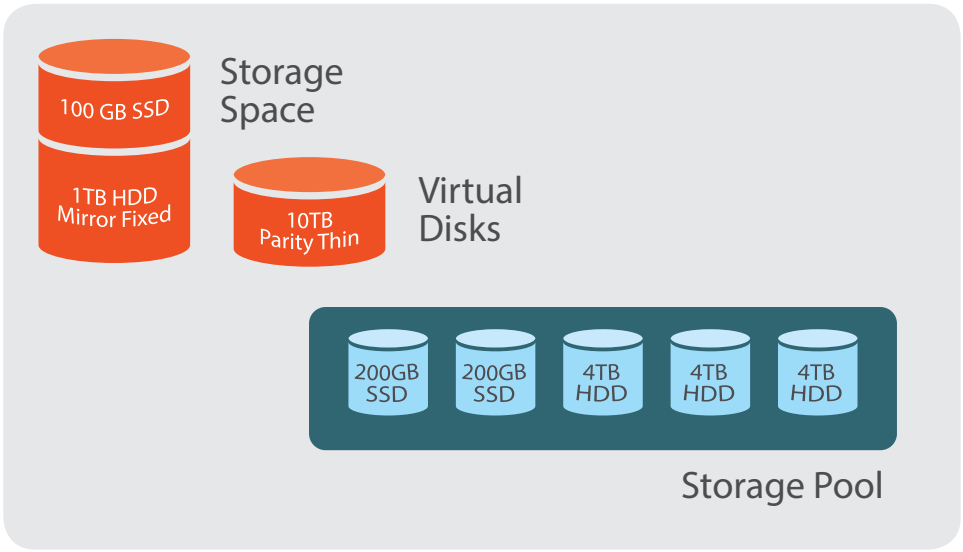

Figure 4: Example of Storage spaces in action

By default, an optimization is performed each night, detecting the most used blocks and moving them to the SSD tier. Additionally, the SSD tier is used as a write-back cache. When writes are performed, the write is initially written to the SSD tier, providing the fastest write times, and then background writes are "lazily" written to the HDDs without impacting the performance of the storage

solution. When using tiered Storage Spaces, it's possible for Windows to act as an enterprise storage solution.

For larger deployments, it's possible to use an external enclosure that connects to multiple Windows servers in a cluster, and a clustered storage space is created. This clustered storage space can have disks created that are visible to all the Windows server boxes, and that shared storage—for example, Cluster Shared Volumes—can be used by workloads such as Hyper-V and SQL Server.

Closing Thoughts

Windows Server 2003 was a phenomenal OS that provided great capabilities, but when you compare it with the capabilities of modern-day Windows Server, it doesn't come close in terms of scalability, capability, or security. Windows Server 2012 R2 enables a wide range of solutions for organizations and forms a foundation from which companies can embark on a journey to the cloud. Windows Server 2012 R2 also enables the maximum benefit from your hardware investments, especially when considering memory and SSD storage.

MKP-484