Lunch Break SERIES

The 24/7 Availability Challenge: Host Level Backups Gain Momentum

Today, enterprises need to operate seven days a week, 24 hours a day, and 365 days a year. With no exceptions. Downtime is simply not tolerated because the business is so tightly connected to the performance of its Information Technology (IT) systems. This is a challenge for the IT professionals that must keep all systems going. Let's explore some of the solutions to this problem including host level backups.

SPONSORED BY

Redmond MAGAZINE

VEEAM IT JUST WORKS!™

BY PAUL KORZENIOWSKI

For IT, Availability Has Become Job One

Host level backups gain momentum as enterprises search for ways to limit system downtime

IT Faces Many Challenges in Meeting Today's Business Demands

Increased connectivity and a growing dependence on technology has altered the business landscape. The traditional nine-to-five expectation is no more. Nowadays, enterprises need to operate seven days a week, 24 hours a day, and 365 days a year. With no exceptions.

In addition, downtime is simply not tolerated because the business is so tightly connected to the performance of its Information Technology (IT) systems. This change in outlook has put the Chief Information Officer (CIO) in a bind. While system availability expectations have risen, the work required to keep servers chugging along has become more onerous: Techies find themselves managing solutions that sit in a hodgepodge of data centers and serve customers spread across the globe.

Further compounding the problem, the tools used to oversee data center hardware often do not mesh with

current system design. Even though companies have virtualized the majority of their servers, their products and DR methodologies are often outdated and provide inefficient approaches to keeping those applications available. A few new ways to enhance availability have emerged; host based backup and restoration is emerging as the best option.

The Changing IT Landscape

Nowadays, corporations rely on technology more than ever before. Increasingly, IT has become a corporate differentiator and a key business enabler. Top management's view of its computer systems is shifting from a cost center to a revenue generator.

Product and service differentiation often revolves around how well an enterprise leverages technology. In a growing number of cases, the dividing line between the core business and IT are blurring. Founded in 1994, Amazon.com is an example of a retailer that has branched out beyond its traditional roots and relied on technology to rise to Number 18, with $107 billion in revenue in 2015, on the Fortune 500 list.

The change in mindset is evident in current hot technologies. With the Internet of Things (IoT) becoming more popular, enterprises collect more data than ever before, so data usage is exploding. IDC expects that worldwide data storage requirements will rise from 4.4 zettabytes (one trillion Gigabits) in 2013 to 44 zettabytes in 2020. Businesses are not just collecting data; they are relying on it for competitive gain and to streamline inefficient business processes, deliver more appealing products, and serve customers faster and more efficiently.

With the Internet of Things (IoT) becoming more popular, enterprises collect more data than ever before, so data usage is exploding.

Corporations have been trying to differentiate themselves via customer service and here again technology plays a major role. Users have high expectations. Clients contact firms at all times and expect immediate service. If a customer tries to contact the organization at midnight or 6 a.m., they expect —no demand—a response. Communication is no longer limited to the traditional work week: customers browse and buy at night and during the weekend, so service is a seven day a week, 24 hours a day endeavor.

Downtime impacts the corporate ledger. Gartner pegged the average downtime cost at $5,600 per minute, which extrapolates to well over $300K per hour. Those are hard numbers. Social media creates other problems as word of mouth quickly goes viral. An enterprise may be riding high in the morning, encounter a system misstep at lunch, and spend the rest of the afternoon and coming days trying to recover—sometimes unsuccessfully.

In addition, downtime, in some

WHAT TO LOOK FOR IN A BACKUP AND RECOVERY SYSTEM

With a growing emphasis on system availability, organizations are looking for new restoration and backup options. They need tools that leverage modern data center technologies, such as virtualization and cloud, to create an always-on organization.

Here are four key capabilities:

1. High-Speed Recovery: Nowadays, speed is the new business currency, so faster is better. Solutions that work quickly and efficiently help enterprises position themselves for success.

2. Data Loss Avoidance: With data becoming more vital to gaining a competitive advantage, corporations want to ensure that once generated their information is protected and backed up. Real time data protection is the Holy Grail, so businesses demand systems that offer companies much more granularity than the option of nightly backups.

3. Verified Recoverability: Backups are only as good as the data that they can deliver when needed. They require bullet proof tools that guarantee the recovery of every file, application, and virtual server.

4. Increased Visibility and Proactivity: IT wants to proactively monitor their systems and be alerted about issues before they negatively impact operations. An ounce of prevention is better than a pound of cure.

So when evaluating a restoration solution, firms have a lot to think about. By looking at recovery speed, data loss avoidance, and system visibility, ideally they find a strong solution.

cases, has a very real and negative impact on IT professionals' careers. One in five firms fires an IT employee as a result of network downtime, according to an Avaya survey.

The end result is uptime demands are increasing. Companies are striving to operate always-on businesses. In the past, two nines (99.0%, which translates to a 1.8 days of downtime per year) availability was commendable, but nowadays, organizations strive for three nines (99.9%, which means 8.75 hours annually) and even four nines (99.99%, which translates to less than one hour per year). As a result, CIOs spend a growing part of their day trying to ensure that their data center infrastructure is resilient.

The Availability Challenges

While organizations recognize the need to service customers around the clock, they run into a number of barriers when trying to improve service availability. As downtime has become less tolerable, systems have become more complex. Software is broken down into smaller elements, so there are more items to track. Enterprises mix premises and cloud based solutions, so hardware is scattered in more locations. In addition, users work with a wider range of devices. They have supplemented desktop PCs and laptops with smartphones and tablets.

Another challenge is many companies rely on traditional system backup and recovery tools. These products were designed for yesterday's solutions: they focus on operating system backup and recovery, but nowadays, the virtualization layer controls system resources. "The virtualization market has matured rapidly over the past few years, with many organizations having server virtualization rates that exceed 75%," stated Michael Warrilow, research director at Gartner.

As the use of virtualization has increased, the need for new backup and recovery tools has become more pronounced. Rather than operating system based tools, organizations require products that operate with leading hypervisors, like VMware and Microsoft Corp.'s Hyper-V. However, many IT departments are ill prepared to meet that mandate. According to a Forrester Inc. survey, 35% of global disaster recovery decision-makers listed mismatched systems and expectations as one of their biggest challenges they face when they try to recover from an IT related major business disruption.

So, the search for new ways to backup data center systems has intensified. Two options, agent based solutions and host level backups, have emerged, and the latter is gaining traction for a number of reasons.

The Virtualized Restoration Duopoly

Initially, enterprises backed up their virtualized data with guest agents, software running on individual Virtual

WHEN DOES THE BUSINESS RECOVER?

The concepts of downtime and system interruption are easy to understand, but given the complexity of the systems and the processes involved, determining their potential impact on the organization can be challenging. A financial system down for two hours may represent a more significant loss than a retail system down for one day. Two terms, Recovery Time Objective (RTO) and Recovery Point Objective (RPO), have emerged to help businesses better understand the various relationships involved in system interruptions.

RTO is the duration of time needed to restore a service before unacceptable consequences occur. For instance after a self-examination, a corporation determines that a 10-hour break in services means millions in lost sales, a number from which the company may not be able to recover. The restoration services are then designed to bring the systems back online within that timeframe.

RPO focuses on recovering the systems to the point they were operating at before the problem occurred. During the day, the business constantly changes (new orders, product shipments, and bills being paid) and those alterations are entered into a number of different applications. When an outage occurs, information is lost. Depending on the application and the backup features in place, the loss could be a few minutes in lost productivity or a few days of work. RPO is the time needed to get the business back to where it was when the problem arose.

System outages impact businesses in various ways. Managers need to understand the implications from downtime and then take steps to mitigate the chances of a catastrophic problem. RTO and RPO help them understand the underlying issues and take steps to lower the risk of potential damage.

Machines (VMs) and devices. The agent-based approach was quite popular in the 1990s when distributed computing started to become more common because it enabled businesses to spread system resources out on enterprise networks.

However, limitations arose as networks and systems grew larger and became more complex. As the number of devices swelled, the volume of software running on different systems increased as well. Techies found themselves overseeing software running on hundreds, thousands, and tens of thousands of devices and generating status information on a seemingly never ending basis. Data center staff often lacked visibility into device performance and the reports necessary to help them quickly determine how the systems were operating.

Cost was another hurdle. The agent approach required that corporations install software on many systems, and vendors typically charged their clients on a per agent basis. As applications became more dispersed and more modular, more agents were needed, and licensing bills rose significantly.

Another approach, host-level backup has recently been gaining ground. This technique does not require software running on all systems. Host-level products back up system images as well as corporate data within VMs. As a result, backup and restoration becomes portable: Virtual Hard Drives (VHDs) can go anywhere in the data center. Another plus is host level solutions are granular: they backup and restore individual items and not necessarily entire hosts. Because the entire state of the VMs, including all of its metadata and snapshots, is captured, bare metal restores become fairly fast and simple to complete.

As they adopt host based backup and recovery, enterprises improve their uptime, and become more responsive.

Lower costs are another potential benefit. Licensing is usually per-host or per physical CPU basis rather than on each element, so expenses often are less than the agent approach. There is more cost consistency: expenses do not change as VM configurations are altered.

The various market drivers are forcing enterprises to closely examine their backup and restoration capabilities. Because of the potential pluses, a growing number of businesses are adopting host based restoration.

Host based Systems Challenges

Host based solutions do present firms with a few hurdles. VM-level backups aren't the same as traditional methods. Since training and new business processes are required, a firm must include time for employees to gain new knowledge.

Sometimes, businesses must pause a VM in order to back it up from the host. Ideally, such a pause is typically very brief.

Benefits of Moving to New Restoration Options

As they adopt host based backup and recovery, enterprises improve their uptime, and become more responsive. The new solution delivers high-speed recovery. With problems quickly rectified, companies are better able to sell products and service customers, and meet their system and business objectives (see sidebar).

By making the change, businesses leverage their investments in virtualization technology, storage, and cloud technologies. They deploy modern data center infrastructures that help organizations save time, mitigate risks, and reduce capital and operational costs.

Uptime has become a major issue for IT departments. Many businesses have been a few steps behind in updating their restoration solutions, so they mesh with their modern computer infrastructure. Consequently, they experienced trouble in meeting uptime objectives. New solutions, such as host-level restorations, are becoming more popular and helping organizations leverage technology for competitive gain. R

Paul Korzeniowski is a freelance writer who specializes in computing issues. He has been covering technology issues for more than two decades, is based in Sudbury, MA, and can be reached at paulkorzen@aol.com and followed at #PaulKorzeniowski

BY PAUL SCHNACKENBURG

Better Backups with Hyper-V in Windows Server 2016 TP5

ADCs serve a new generation of business applications, with specific benefits when implemented in the cloud.

With the Technical Preview (TP) 5 of Windows Server 2016 recently released, I thought I'd take the opportunity to cover all the new goodness coming in Hyper-V. TP 5 is apparently feature complete; it's also the last preview before Release to Manufacturing (RTM), expected sometime in the second half of 2016.

Some of the new features were covered in my interview with Ben Armstrong, principal program manager on the Hyper-V team, published in six parts here on virtualizationreview.com. Part 1 is here; it contains links to all the other parts.

In the past, Microsoft would develop Windows Server "in the back room" with a public beta late in the development cycle, more for fixing bugs than as a source of ideas for improvements or new features. This version of Windows, on the other hand, has been "in the open" since TP 1 back in October 2014, all the way through to TP 5, released April 27.

Overview

Given the long development time, expectations are high. Fortunately, this version doesn't disappoint, with many

new features, including:

- A new type of checkpoints
- A new backup platform
- Rolling cluster upgrades
- VM compute resiliency
- Storage QoS
- Storage Spaces Direct
- Shielded VMs
- Windows and Hyper-V containers
- Nano server and PowerShell Direct

I'll go over each of these in brief in this article. There are also other improvements to existing features such as Shared VHDX, Hyper-V Replica, more online operations for VMs, a better Hyper-V manager console and more.

Backup and Checkpoints

Backups in Hyper-V can sometimes be a bit shaky, due to a reliance on the underlying Volume Shadow Copy Services (VSS) system. Windows Server 2016 instead makes change tracking a feature of Hyper-V itself, making it much easier for third-party backup vendors to support Hyper-V.

Snapshots and checkpoints are dangerous for production workloads. They have a convenient workflow: take a snapshot; make some changes in the virtual machine (VM); if those changes turn out badly, simply roll back to the snapshot.

The problem is that if it happens on a Domain Controller (DC) or database server that's replicating with other servers, it's now out of sync; and there's no easy way to tell, nor any easy way to fix it. Microsoft made changes in 2012 for Active Directory (AD) DCs to make them safer, but this still doesn't cover any other workloads in danger from a wrongly-applied snapshot.

Production checkpoints in Windows Server 2016 (the classic checkpoints can still be used) uses VSS inside the VM; when you apply them, the VM will assume it's been restored from a backup and reboot, rather than be restored to a running state. This eliminates the danger while retaining the convenience of snapshots.

Rolling Cluster Upgrades

The upgrade story from Windows Server 2012 to Windows Server 2012 R2 was pretty good, enabling live migration of VMs from the old to the new. But you still had to stand up a separate Windows Server 2012 R2 cluster to start the process, which wasn't ideal.

Going from Windows Server 2012 R2 to Windows Server 2016 is a lot easier: simply evict one cluster node, format and install Windows Server 2016, and add it back into the cluster. It now acts as a Windows Server 2012 R2 host, so VMs can be Live Migrated to it; that means you can take another host and clean install it. Rinse and repeat as many times as required. When all nodes are upgraded and you're sure you're not going to add any down-level nodes, you use PowerShell to upgrade the cluster functional level, similar to the way you do AD upgrades.

VMs have had version numbers internally since the first version of Hyper-V. Because of the rolling cluster upgrade scenario, they're now visible, so you need to be able to upgrade the configuration files for each VM. This is also done using PowerShell. Once upgraded, a VM can only run on Windows Server 2016 hosts. Each VM uses the new .vmcx file format for configuration and .vmrs for runtime state data; both are binary files and do not support direct editing (unlike the current XML file type).

VM Compute Resiliency

Clustering hosts together provides continuous VM availability for planned downtime; simply Live Migrate VMs from the host first, then perform the maintenance. For unplanned downtime, VMs on a failed host are automatically restarted on another host in the cluster, providing for high availability with a few minutes downtime for the restart. So far, so good.

There are times, however, when host clusters can cause issues by themselves. A short network outage between the hosts can cause them to initiate a failover of many VMs, when, in fact, the network could right itself after a few seconds. Such a failover could cause more downtime, with numerous VMs restarting simultaneously.

Clustering hosts together provides continuous VM availability for planned downtime; simply Live Migrate VMs from the host first, then perform the maintenance.

In Windows Server 2016, if a host loses connectivity to the cluster, the VMs will keep running for four minutes (this can be changed) in "isolated mode." If it's longer, normal failover will occur. If a host has numerous disconnections over a 24-hour period, it will be quarantined and its VMs Live Migrated off as soon as possible.

Today, if a VM has an outage to the shared storage where the virtual disks are housed, it'll crash if the outage is longer than about a minute. In Windows Server 2016, if storage connectivity is lost, the VM will be paused, pending reconnection to its

virtual disks, avoiding the likely data loss in a crash.

You can now specify priority for VMs—high, medium and low—when failover occurs. TP5 allows admins to create sets of VMs, define dependencies between them, and let this dictate the order in which VMs are started.

Storage QoS

In the current version of Hyper-V, you can set a min or a max (or both) value for IOPS for virtual hard disks. This works fine as long as the backend storage can actually deliver the combined IOPS requirement for all running VMs; if it can't, there's no way for the individual hosts to manage IOPS requirements.

Windows Server 2016 brings a centralized storage IOPS "cop," sitting on the Scale Out File Server (SOFS) nodes. It's managed either through PowerShell or with Virtual Machine Manager (VMM), and provides a way to create policies that can be applied in aggregate across VMs or to individual VMs. It also monitors the IOPS actually used by each VM, giving you a more comprehensive view of the way your applications use storage.

Storage Spaces Direct

Microsoft's implementation of Software Defined Storage (SDS) took shape in Windows Server 2012. SOFS nodes act as the front end of a SAN (but simpler to set up and manage); SAS JBOD (just a bunch of disks) disk trays with HDDs and SSDs provide the data storage.

In Windows Server 2016, Microsoft takes the next logical step by offering Storage Spaces Direct (S2D), which provides pooling of local storage (SAS, SATA and NVMe HDDs and SSDs) in each host, and offers it up as VM storage. This can either be disaggregated with storage nodes in one cluster and Hyper-V nodes in another, or hyper-converged where each host is both a storage node and VM host.

New in TP 5 is the ability to have fewer than four nodes, along with support for SATA disks (previous previews required SATA disks to be connected through a SAS adapter).

One basic problem with any hypervisor is that host and/or fabric administrators have to be as trusted as the highest level administrators in an organization.

Shielded VMs

One basic problem with any hypervisor is that host and/or fabric administrators have to be as trusted as the highest level administrators in an organization. If VMs are hosted elsewhere, in a public cloud, for example, you have to have a lot of trust; a rogue fabric administrator can inspect the memory of a running VM, take an offline copy of the virtual disks, mount these and steal secrets such as passwords or perhaps mount an offline attack against an AD database.

There are a few building blocks for Shielded VMs: generation 2 VMs now come with a virtual TPM chip, which enables Bitlocker for Windows VMs, and dm-crypt on Linux VMs for the virtual hard disks. Generation 2 VMs also provide Secure Boot for both Windows and Linux VMs, as they start from a virtual UEFI.

On separate physical servers that are part of an isolated administrative forest, there's a Host Guardian Service which attests to the health of Hyper-V hosts. There are two models for this: Administrator Attestation and Hardware Trusted Attestation. The first relies on trusted hosts being in a particular AD group; the second one uses new TPM version 2 chips in each host to protect the hypervisor from tampering.

In TP5, a new mode called Encryption Supported supports vTPM, disk encryption and Live Migration encryption; but it still provides less assurance than a true Shielded VM. In TP5, you can also convert normal generation 2 VMs to Shielded VMs, while a new recovery environment allows for troubleshooting of a Shielded VM.

The end result of a Shielded VM is that the fabric administrators have no access to the VM. They can turn it off, but they can't access its memory or connect to it with VM Connect; if they copy the virtual hard disks, they can't access them because they're encrypted.

Windows and Hyper-V Containers

Containers are all the rage in the IT press, and although I think it's going to take a lot longer than the pundits believe before we're all "containerized," Microsoft is now in the running with two flavors of containers. Each container can either run Nano server or Server Core (not the GUI version); for developers, the flavors are identical. The difference comes in the deployment phase: in your own datacenter where you (probably) trust the code running in each container, you can rely on the weaker isolation of a Windows Container, but if you're deploying your code in a public cloud or at a hosting provider, the Hyper-V container gives you the

same isolation the hypervisor provides.

Nano Server

The biggest change in Windows Server since NT was conceived is undoubtedly Nano server. It's a minimal disk footprint, low resource, GUI-less, no local logon server where you add only the functionality needed. The only roles supported today are Hyper-V host server, SOFS server, and as an application platform for modern applications. The benefits here are very small attack surface, low overhead and less frequent reboots due to fewer patches.

PowerShell Direct

PowerShell Direct is a very useful feature in Windows 10 and Windows Server 2016. If you have the credentials, you can run cmdlets inside one or more VMs from the host without having to set up PowerShell remoting first.

New in TP 5 is the ability to run PowerShell directly on a Nano server, along with cmdlets for working with local users and groups.

The biggest change in Windows Server since NT was conceived is undoubtedly Nano server.

Other Improvements

There are quite a few new features in TP 5 coming to both Windows 10 and Windows Server 2016. Even with the Hyper-V role installed, you can now use Connected Standby power state.

The ability to connect a VM directly to a PCIe hardware device is interesting. It doesn't work for every device; more information is available here. If you want to try it out, see these instructions. At this stage, the main aim is connecting VMs directly to NVMe superfast storage, but GPU support is also coming.

Host resource protection is a feature Microsoft added in response to Azure vulnerabilities. In these cases, VMs with hostile code would try various attack methods to starve the host of resources. Host resource protection detects this and limits resources available to the VM.

Hyper-V Manager now lets you enter alternate credentials when connecting to remote Hyper-V hosts, and also save these credentials. The manager can also manage both Windows 10 and Server 2016 as well as Windows 8 and 8.1, and Windows Server 2012 and Windows Server 2012 R2 hosts. The console is operating using the WS-MAN protocol over port 80. WS-MAN also makes it easier to enable a host for remote management.

Hyper-V Replica now supports shielded VMs, provided the destination replica server is authorized to run the replicated VM(s). To support containers, Hyper-V now supports nested virtualization, with a VM being a Hyper-V host with other VMs running inside it; several levels of this nesting is possible.

The new version of Hyper-V brings many unique features as well as important improvements to existing ones. Make sure to check out TP 5 yourself. R

Paul Schnackenburg, MCSE, MCT, MCTS and MCITP, started in IT in the days of DOS and 286 computers. He runs IT consultancy Expert IT Solutions, which is focused on Windows, Hyper-V and Exchange Server solutions.