CUSTOM REPORT

VIRTUALIZATION & Cloud Review

DRAAS FOR VMS IN THE MODERN DATA CENTER

By Trevor Pott

KeepItSafe®

www.keepitsafe.com

Data protection in today's organizations increasingly involves either a hyperscale public cloud provider or a regional service provider. Cloud backups have become the new normal, and Disaster Recovery-as-a-Service (DRaaS) is being embraced by organizations of all sizes. Not all approaches to cloud-backed data protection are equal, however, and an in depth understanding of capabilities often provides superior results.

For those ready to embrace the cloud for data protection, perhaps the best advice that can be given is to learn to love the offerings of managed cloud service providers. While the hyperscale public clouds are more feature rich – and in some cases less expensive – service provider clouds tend to come with actual customer service, a selling point that can seem irrelevant until the day comes when one has to rely on their data protection solution.

To illustrate, let us explore one of the most common cloud-backed data protection solutions currently in use: Veeam Cloud Connect as a solution that backs up an on-premises VMware infrastructure to a service provider cloud. Veeam is well-known as a data protection company, particularly in the VMware space, however they also provide data protection offerings for bare metal servers, the Hyper-V hypervisor and even offer a lightweight Office 365 backup solution.

BACKUP ALL THE THINGS

Any discussion about data protection should begin with a discussion about the objectives of data protection. Experienced administrators will be familiar with the terms Recovery Time Objective (RTO) and Recovery Point Objective (RPO), both critical concepts as relates to data protection.

RTO is the length of time it takes to restore a workload or file from a backup. RPO is the maximum allowable amount of time that can elapse between when the last backup was taken and the restore event. As such, RPO is a measure of the maximum allowable data loss.

All organizations want both RPO and RTO to be zero, however, excepting in very rare cases, this is economically unlikely to occur. While an RTO of less than 30 seconds is entirely doable today, RPOs at or near 0 often require the participation of the target workload in the backup process.

RPO 0 typically involves the creation of application clusters, which are quite common with database applications. These are often difficult to configure, touchy to maintain, and expensive because they require purchasing multiple copies of the application software in question, and keeping those copies online at all times as part of the cluster.

The speed of light plays a role in

Cover: Everything Possible / Shutterstock.com

determining the viability of RPO 0 as well. The greater the distance between cluster members, the higher the latency that workload will experience. While a metro-area cluster of workloads (workloads separated by no more than 100km) is considered by most organizations to be acceptable, the speed of light renders more geographically dispersed clusters unusable for many use cases.

True RPO 0 solutions are thus capable of providing protection against the failure of an individual data center, but offer little protection against events that impact an entire region. Hurricane Sandy's impact on the east cost of the USA during the closing months of 2012 has become the canonical example. Sandy impacted multiple data centers in the New York area alone. Some organizations which had planned failovers to data centers in New Jersey also discovered that having their second site in the next state over was not geographically distant enough to ensure continued operations, as New Jersey data centers were also heavily impacted.

A CAREFULLY CONSIDERED DATA PROTECTION STRATEGY IS INCREASINGLY VITAL TO ORGANIZATIONS OF ALL SIZES.

There is thus always a judgement call to be made about which RTOs and RPOs are acceptable for which workloads. There are many considerations to be balanced.The cost of creating RPO 0 clusters for workloads that support it. The cost of internet and/or WAN bandwidth to the Disaster Recovery (DR) site, and the impact of backups on production infrastructure and the geographic resiliency required must all be balanced in making these determinations.

As the amount of data under management and the number of workloads grow, so too does a data center's complexity, as well as the expertise required to design a viable data protection strategy. A carefully considered data protection strategy is increasingly vital to organizations of all sizes. Organizations may wish to consider seeking expertise and guidance from vendors, consultants and service providers.

It is not uncommon for organizations to create a single RPO 0 metro cluster for their core database application, with a low-but-non-zero RPO additional layer of backup providing geographic resiliency. Other workloads in that same organization may have RPOs as high as a month: these workloads may simply be Operating System Environments (OSEs) and an application that processes data, but not contain any actual data themselves. If failed over to a month-old backup, the worst case scenario is usually that some updates need to be run to bring them back up to snuff.

Being able to manage this diversity of requirements is the hallmark of a mature data protection solution.

CONTINUOUS DATA PROTECTION

Astute administrators may have objections |to some of the characterizations of data protection challenges discussed above. There

exist data protection vendors that offer Continuous Data Protection (CDP). Veeam, for example, is expected to be releasing this capability with Veeam 10, and it should work with vSphere 6.5 or newer.

CDP monitors workloads in real time, capturing all writes made by that workload, and sending those writes to the data protection destination. While this certainly sounds like an RPO 0 technology, it isn't. To understand why one must understand the difference between application consistent and crash consistent backups. In turn, this involves understanding the complexities of the storage path that underlies modern IT infrastructure.

Storage solutions are slow. Even the fastest NVMe all-flash storage is thousands of times slower than the volatile main DRAM memory of a server. As a result, it is common practice to add DRAM caches to storage in various places in order to speed up access.

These caches often lie to applications about the security of their data. Applications attempt a write, and receive confirmation that the write has been committed. In reality, that write has only been made to the volatile DRAM cache and not to non-volatile storage. The write will be flushed from cache to storage as quickly as the storage subsystem allows, however, there can be several milliseconds – or in extreme cases several seconds' – worth of lag between when the application thinks it has made a write and when the write is actually committed to non-volatile storage.

A few milliseconds may not seem like a lot of time, but in that few milliseconds dozens of sales could be processed from an online store, or a critical patient record could be updated. For some workloads – again, typically databases – this gap, however small, matters.

Write caches can exist at multiple levels. Applications themselves can employ caching, OSEs typically employ caching, and caching can exist at the hypervisor level, and potentially in multiple places within the storage path below. CDP solutions typically monitor writes that are visible to the hypervisor or OSE, and one or more layers of caching usually exist above where the CDP solution can monitor.

This means that CDP solutions offer – at best – a perpetual crash consistent backup. A crash consistent backup is a backup equivalent to turning the power off on the server while it is in operation. At current, no cloud-backed CDP solutions manage even this tight of an RPO.

The ability to restore data is just as important as the ability to back it up. Make sure you pick the right backup consistency for your workload, and that you test that you can restore from the backups you've made. Data which cannot be restored from a backup does not exist!

Because of the time required to compress the data recorded and then transmit it to the data protection destination, cloud-backed CDP solutions typically have an RPO of between 5 and 30 seconds. This is a far cry from the true RPO 0 solutions provided by clustered applications. Clustered solutions do not accept a write as committed until it has been written to all members of the cluster, which is why the speed of light and geographic proximity play such a strong role in the latency of these solutions.

Application consistent backups involve the application, OSE, and potentially more

elements of infrastructure in the backup process. When a backup is to be taken, the workload is quiesced (or "stunned") so that no new writes take place for a brief period of time. All write caches are flushed, and a snapshot of the workload is taken. This snapshot is then sent to the data protection destination.

Application-consistent backups can only ever be snapshots. A pseudo-CDP solution can, in theory, take a series of application-consistent snapshots in rapid succession, however, as the frequency of these increase the impacts on the production workload become more noticeable.

RTO 0

With a thorough understanding of RPO completed, a brief exploration of RTO is warranted. As mentioned above, RTOs very near o are entirely possible. This is thanks to virtualization.

Workloads which have been backed up to an appropriate infrastructure can be instantiated directly from the data protection destination. A VM backed up to a cloud provider, for example, can simply be turned on at the cloud provider, allowing for the backup copy of that workload to be online in the time it takes to boot the VM. Similarly, if organizations have a local backup repository in addition to their offsite destination it is likely that the local backup repository will have the same capabilities.

Organizations using Veeam with a local repository can make use of this functionality. Veeam repos can have a vPowerNFS folder on it. This feature, called Instant VM Restore, takes the workload's backup (which is compressed and deduped) and makes it available via an NFS share. It performs the rehydration of the VM on the fly, meaning that it does not have to be unpacked prior to being launched. Administrators can then storage vMotion during a maintenance window.

Cloud-side solutions work similarly, with the exception that getting the workload back onto one's production infrastructure isn't quite so seamless. Here, the proximity of data protection destinations can play a role.

WORKLOADS WHICH HAVE BEEN BACKED UP TO AN APPROPRIATE INFRASTRUCTURE CAN BE INSTANTIATED DIRECTLY FROM THE DATA PROTECTION DESTINATION.

Service provider clouds are often of use because they are located geographically proximate to an organization's primary data center, but may also offer additional data centers at a geographically distant location as well. One advantage that regional service providers can sometimes provide over the hyperscale public clouds is the ability to, in cases of emergency, load all of an organization's data onto a hard drive and simply drive it down the road. This helps when a complete restore is required, and pulling that data down across the internet or WAN connection would take too long.

VEEAM, VMWARE AND SANS, OH MY!

With a thorough understanding of RPO and RTO out of the way, the details of how Veeam can interact with one's VMware environment

will offer insights into how to achieve the performance desired for the workloads required. The Veeam software consists of three roles, all of which can be installed on a single server, but which should be installed separately for larger deployments.

The first role of a Veeam server is the management role. The management role handles job scheduling and backup coordination. It is a reasonably lightweight role, and even in large deployments there is no reason not to simply virtualize this role.

The second role of a Veeam server is the repository or "repo" role. The repo is the backup destination. The repo is basically a big box of disks that can occasionally light up the odd VM for recovery as discussed above with the Instant VM Restore functionality.

The Veeam Cloud Connect solution allows a cloud service provider to be used as the repo. While this is helpful, it is recommended that organizations maintain both a local repo and an offsite copy of their data. This allows for Instant VM Restore with the ability to storage vMotion over live workloads back to the production cluster using the local data store, as well as protects against failure of the production data center by taking advantage of the cloud service provider.

The 3-2-1 backup philosophy has become the industry standard approach to backups. The 3-2-1 backup philosophy states that organizations should maintain 3 copies of their data, on two different backup mediums with one copy off site. Remember: if your data doesn't exist in at least two places, then it does not exist!

The third role of a Veeam server is the proxy role. The proxy talks to Veeam guest agents and to the management application which controls an organization's infrastructure. In the case of VMware, this is vCenter. The proxy's purpose is to get the bits from the workloads to the repo, and it has several means of doing so, the mode used ultimately determining whether or not the proxy should be installed on a physical server, or can be left inside a virtual machine.

In network mode the proxy can be either a virtual machine or a physical system. A Veeam proxy operating in network mode will connect to VMware servers using the VMKernel. This is a universally compatible approach to performing backups, but it is also limited to a percentage of the link speed. At 1GbE the cap is very slow, though at 10GbE it is more usable.

The VMware hypervisor (ESXi) caps the rate at which VMs can be backed up using the VMKernel method, so as not to crash ESXi. This places an upper limit on the speed at which backups can occur in network mode, even with the most powerful storage system underlying both the VMware cluster and the Veeam repo.

Network mode is the default deployment type for Veeam proxies. As a result, many administrators new to Veeam often feel that Veeam is slow when compared to other backup solutions. This can be remedied by engaging one of the other proxy modes.

In Hot Add mode, the proxy must be a virtual machine and it must be located on the same cluster as the VM that is being backed up. This means that one Veeam proxy per vCenter will be required.

A Veeam proxy in Hot Add mode takes a snapshot of your VM to be backed up and then attaches that snapshot to the proxy. The proxy then uses the storage network instead of the

VMKernel network to send the data to the repo server, allowing for significantly faster backups.

The downside to Hot Add mode is that things can go a little awry if someone reboots any of the servers involved while the snapshot is attached to the Hot Add proxy. This is increasingly a rare issue, however, as Veeam has been continually refining the Hot Add code to address the problem.

Veeam proxies can also operate in Direct SAN mode. Direct SAN proxies should be physical servers. In Direct SAN mode, the Veeam proxy would have an HBA installed and thus the ability to talk to the Storage Area Network (SAN) storage solution.

Veeam will orchestrate a snapshot with VMware. When VMware is done taking a snap the SCSI ID is handed off to the proxy so that it can back that snapshot up to the repo. This offers the best backup performance of the various proxy modes, however, restores can be a different story.

Thick provisioned VMs can be restored directly via the SAN. Thin provisioned VMs, however, must be restored through the VMware VMKernel method, meaning that they can be just as slow as network mode.

Direct SAN mode also comes with a caveat that one must pay attention to the automount settings of the Windows Server on which the Veeam proxy is installed. If windows automounts the snapshot then it can mangle the metadata, ruining the VM. Veeam will disable the automount setting during install, so do not re-enable it, and check that the organization's Active Directory GPOs do not enable it either. (This should have been disabled for security reasons anyways.)

Similar to Direct SAN mode, Veeam can also operate in Direct NFS mode. Like Direct SAN mode, Direct NFS proxies should be installed on physical servers. Direct NFS accesses NFS shares used to store VMs directly offering a very similar level of performance to Direct SAN.

Organizations with Veeam Enterprise will be able to take advantage of SAN snapshots. This integration requires Veeam to integrate directly with the SAN via the SAN's API. This allows for the quickest and most accurate snapshots.

When using SAN snapshots Veeam will ask vCenter to quiesce the VM, order the SAN to take a SAN snapshot, and then tell vCenter to release the VM. SAN snapshots can be configured up to take SAN snaps without quiesceing the VM, but for the application consistency reasons discussed above, this capability is rarely used.

Veeam proxies can use the VMtools installed inside a VM to communicate with Windows' Volume Shadow Service (VSS). This allows for application-consistent snapshots for VSS-enabled applications such as Microsoft SQL, or Exchange.

VEEAM CLOUD CONNECT

Once the data is off the VMware infrastructure and on to a repo, it's time to send it off site. This is where Veeam Cloud Connect comes in.

Cloud Connect is a product that allows organizations to see a list of available service providers, select one, and immediately start backing up their IT infrastructure to that destination. While it is designed to be push-button simple, the solution is feature rich, and integrates with the core Veeam backup offering.

Cloud Connect is predominantly a backup solution, and while it has certain disaster

recovery capabilities, many administrators will quibble about whether or not is should be viewed as a fully-fledged disaster recovery offering. Cloud Connect lacks some of the automation that advanced users will expect from a more complete disaster recovery solution, relying on manual administrator intervention in many cases.

Restoring workloads backed up to a service provider using Cloud Connect is simple. Restores are simple and can be done in one of two ways. If your Veeam infrastructure at the destination you wish to restore to is intact, then you simply select the VM you want to restore and proceed from there. If one wishes to restore to a site that does not contain your extant Veeam infrastructure, one simply instantiates a Veeam Cloud Connect virtual appliance, enters username and password, and they can then immediately begin to pull data back down to the on-premises data center.

Veeam restores can thus require some manual intervention. The Cloud Connect approach however, allows administrators to install a Cloud Connect appliance in a different data center other than their original production data center – or the same data center, but with all-new equipment – and immediately begin restoring workloads. This makes it highly useful in hurricane Sandy-like situations.

Cloud Connect also offers complete network extension capabilities that make restoring VMs on the cloud provider side effortless. The Veeam network extension appliance creates either a layer two bridge and/or serves as an automatable firewall. This means that organizations don't need to address anything at the cloud provider end, a traditional problem with data recovery solutions, especially in the cloud.

Cloud Connect in full failover mode recreates the entire networking environment, but requires all VMs to failover. This can be used not only to fail over a production environment as part of a disaster recovery, but it can be used to clone an entire environment for dev and test purposes as well. Full failover mode only requires a Cloud Connect appliance VM on the cloud service provider side.

Cloud Connect does not need to do a full failover. If network extension appliances exist both on-premises and at the cloud service provider, then Cloud Connect can Partial use layer 2 extensibility to connect the on-premises network to the cloud provider's network. This allows for partial failover of VMs without having to re-address those VMs, or any other part of the network.

CLOUD CONNECT IN FULL FAILOVER MODE RECREATES THE ENTIRE NETWORKING ENVIRONMENT, BUT REQUIRES ALL VMS TO FAILOVER.

Cloud Connect can also create a stateless clone environment that doesn't commit any changes to the replicated environment, which is useful for testing that backups are functioning as intended. Using this capability is called "rolling back the failover".

CONCLUSION

Cloud data protection solutions offer a

diversity of approaches to protecting one's workloads. Traditional problems, such as the need to re-address VMs, have largely been solved. Workloads can be backed up using a number of different methods with varying consistency levels designed to meet a variety of RPOs.

KEY TAKEAWAYS

1. Not all clouds are created equal. Hyperscale cloud vendors do not provide individualized strategy, guidance and assistance for each customer. For those organizations seeking assistance in achieving desired RTOs and RPOs, managed service providers should be engaged.

2. Choosing a data protection provider is complex. Ensure they have the full range of services required. Bear in mind the differences between backup types, and the complexity of disaster recovery. Look for data protection providers with a proven track record.

3. Data is valuable, storage is cheap! Basic cloud backups are inexpensive, but proper disaster recovery capability does come with a premium. "You get what you pay for" is as applicable to data protection as anything else. Take care when selecting your vendors!

RTOs approaching zero are feasible, both on-premises and in the cloud. Those organizations looking for guidance and support on their data protection journey will find a rich ecosystem of cloud service providers that offer fully managed services, and professional consultation, and that ecosystem of providers is growing every day.

Disaster protection is no longer a complex and expensive endeavour fraught with difficulties. Cloud backups are commonplace and disaster recovery can truly be delivered as a service. DRaaS is the new normal, and solutions like Veeam Cloud Connect allow organizations of all sizes to not only achieve peace of mind, but serves to help connect those organizations to service providers which can deliver knowledge and capabilities beyond data protection as required.

ABOUT KEEPITSAFE

KeepItSafe provides global cloud data availability through its Backup-as-a-Service (BaaS), Disaster Recovery-as-a-Service (DRaaS), endpoint protection, and cloud SaaS application backup. Backed by a $1.2 billion public company, j2 Global®, Inc. (NASDAQ: JCOM), KeepItSafe meets data-security protection regulations with ISO 27001, SOC 2, HIPAA, and PCI compliance in 20+ data centers across three continents. KeepItSafe's holistic approach leverages its global footprint and best-of-breed technologies to deliver comprehensive data availability and as-a-Service solutions by offering custom managed and monitored services with 24/7 live support. KeepItSafe's secure enterprise-class data centers support virtual-, physical-, and cloud-to-cloud solutions with 256-bit encryption and multi-cloud scalability via a global network of service providers, system integrators, and cloud resellers.

Find out more:
www.keepitsafe.com

KeepItSafe®