druva

White paper

Cloud-native data protection for VMware

Reduce costs, accelerate recovery times, and simplify management

Q120-20033

The need to evolve backup strategies for VMware

One size does not fit all when it comes to meeting VMware-based virtual machine (VM) data protection requirements. Costs must be kept in check while shrinking recovery time and meeting recovery point objectives. IT professionals are pressured to deliver near-zero data loss and near-zero application downtime. They also need to comply with data privacy regulations such as the European Union's (EU's) General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA). Finally, the IT professional must develop strategies that protect against increasingly tenacious ransomware attacks.

Many enterprises are turning to the cloud—hoping to cut infrastructure costs, accelerate recovery times, and simplify management.

The problem is that protecting mission-critical VMs with legacy approaches is complex and expensive. It is also likely to be sub-par when it comes to meeting more demanding protection service level agreements (SLAs). Many enterprises are turning to the cloud—hoping to cut infrastructure costs, accelerate recovery times, and simplify management. However, many cloud data protection solutions simply "bolt on" cloud support rather than being truly cloud native, and as a result fall short of the "cloud promise." These solutions cannot provide key functionalities, including on-demand capacity scalability, utilizing in-cloud compute, instant data availability, minimal data loss, workload mobility, and improved data visibility.

The current state of VMware backup and recovery

Most disaster recovery and business continuity solutions for VMware environments on the market today are designed for legacy on-premises environments. Cloud capabilities were not included from the beginning of development. The primary problem is that these solutions do not leverage the cloud intelligently.

The first generation of cloud-enabled backup and recovery solutions effectively leveraged the cloud as a replacement for tape archives—a long-term repository for secondary data copies that in an ideal world never need to be retrieved again. But now, a number of backup solutions have taken the next step in terms of running backup software in a virtual machine in the cloud, using the cloud as a tier for old backups and even providing some in-cloud disaster recovery features. However, these "cloud-washed" solutions still lack a number of key capabilities.

Current solutions lack cloud tiering

The first problem is that some of these "cloud-washed" solutions often require all of the organization's data to be backed up both on-premises and in the cloud. The cloud implementation becomes a mirror copy of the on-premises environment, rather than being used to shrink the on-premises infrastructure footprint. Capacity requirements effectively double because data is being stored twice, both on-premises and in the cloud.

While a few solutions have updated to support using the cloud to store old backups, they typically cannot migrate between cloud storage tiers (such as from Amazon S3 to Amazon Glacier), losing an opportunity to further lower costs by leveraging the least expensive of cloud storage options. The lack of using multiple tiers is a problem because it limits the enterprise's ability to use the most cost-effective cloud storage tier for their data. In addition to not using all the available cloud storage options to lower costs, the lack of intelligent in-cloud tiering also results in another silo of storage that must be managed.

Cloud-washed solutions lack automation

The operations problem is exacerbated by the fact that most cloud-washed solutions do not automate the highly laborious, manual processes required for recovery—which not only adds costs but also slows time-to-recovery. For example, recovering from a disaster, unlike a single server outage, requires multiple virtual machines to be recovered and in a specific order. The problem is most backup solutions don't automate runbook execution, which means that IT must engage in the costly and time-consuming manual testing of processes.

Most cloud-washed solutions do not automate the highly laborious, manual processes required for recovery.

Also, most cloud backup and disaster recovery solutions do not automate workload migration. They require manual intervention on the part of the IT manager to target configuration of infrastructure resources, such as networking, as well as mapping the workload's various dependencies, before it can be migrated. Format conversion may also be required before migrating the workload into the cloud platform. Even native VMware tools such as Site Recovery Manager (SRM) require complex artifact-based integration with infrastructure resources, which increases the likelihood of failures, makes troubleshooting difficult, and adds complexity.

Limited integration with cloud compute

Cloud compute cycles are more cost-effective and scalable than on-premises server implementations. With cloud compute the business pays only for the cycles that it uses, but without proper integration to leverage these cycles on-demand, "cloud-washed" solutions are costing you more than they should. These backup solutions were designed for servers that are 'always-on', so they consume cloud compute resources even when idle.

Cloud-washed solutions also miss another opportunity, which is to leverage scale-out compute in the operation of the data protection solution itself. Data protection requires multiple schedules to be managed and executed, and indexes to be maintained. Additional value-add services like search or legal hold require even more CPU resources. A cloud-washed solution is typically scale-up in nature and can't leverage processing power beyond a single server or node.

The impact of the current state of cloud backup and disaster recovery

In short, cloud-washed approaches to using the cloud for backup and disaster recovery do not help IT teams to reduce cost or complexity. These solutions still require IT to remain actively engaged with business continuity processes including disaster recovery, preventing them from focusing on more meaningful projects for the organization. Meanwhile, they are expensive and time-consuming to scale, because they require the customer to purchase, install, and configure more hardware. With workload requirements becoming less predictable and with data center floor space coming at a premium, the expense of on-premises infrastructure is quickly becoming unacceptable and unrealistic.

Requirements for cloud-native VMware backup and disaster recovery

Truly cutting costs and improving business agility and continuity for VMware-based hybrid clouds requires using the cloud in a more strategic fashion.

Firstly, a cloud-native solution should provide the option to backup data from the on-premises production storage implementation directly to the public cloud. This can help to mitigate—and in some instances potentially even to eliminate entirely—the need for on-premises infrastructure. Not only can reduction in on-premises infrastructure cut both capex and opex costs for the organization, it can also help to avoid hardware vendor lock-in. Meanwhile, cloud resources are inherently more flexible and scalable than their on-premises counterparts. IT professionals can more quickly, easily and cost-effectively expand their disaster recovery solutions as the needs of the organization change and grow.

Truly cutting costs and improving business agility and continuity for VMware-based hybrid clouds requires using the cloud in a more strategic fashion.

IT professionals should look for the ability to leverage cloud compute resources, which are available on-demand and for only as long as needed, alongside cloud storage resources, to more cost-effectively support disaster recovery and to open up additional use cases such as test and development, workload bursting, and reporting.

A cloud-native solution can also provide the ability to fail over directly into the cloud, to fail back to the on-premises data center when needed, and to migrate workloads through converting workloads "on the fly" into the cloud service provider's format. Together, these capabilities provide a few important outcomes:

They provide flexibility in terms of where and how applications are recovered. Recovery point objective (RPO) and recovery time objective (RTO) service level agreements (SLAs) today vary per application. They depend on the application's legal compliance requirements, the amount of data loss and downtime that the application can tolerate, and the levels of performance that the application requires when it is operating in a failover state. As a result, one size hardly fits all from a disaster recovery standpoint.

They open the enterprise to harnessing the cloud for additional use cases such as test and development and workload bursting.

Disaster recovery tests and validations can occur at any time—improving IT's and your organization's confidence in its ability to facilitate business continuity.

They also accelerate time-to-recovery. Application restarts can be "push-start" simple, and dramatically reduce the time that it takes to migrate applications between on-premises and cloud infrastructures. In a world that can tolerate only minimal application outages—if any at all—this is important.

From the vantage point of accelerating recovery time, keeping operational costs to a minimum, and freeing up as much of IT professionals' time as possible, requires a single tool that can be used to centrally manage all disaster recovery environments. A single tool and interface are especially important in the face of the increasingly distributed organization, which has multiple remote data centers. One interface to manage all sites and all infrastructure can go a long way when it comes to saving time in learning and navigating various interfaces.

IT teams should also look for an automated, "set it and forget it" approach, to maximize the time they can free up to focus on innovation as opposed to routine, day-to-day management tasks. For example, setup and configuration of networking and dependencies across the infrastructure stack is a laborious process that could save IT professionals a lot of time by being automated.

The benefit of cloud-native VMware backup and disaster recovery

To summarize, a backup and disaster recovery implementation that uses the cloud strategically and natively radically cuts costs because it eliminates entirely the need for a secondary data center. Additional infrastructure does not need to be purchased, deployed, and managed. The amount of data center floor space, which is at a high premium today, that IT needs to procure is minimized. IT staff also can focus their time spent managing, testing, provisioning, and troubleshooting hardware on the production data center. They also have more time to devote to more revenue-facing initiatives for the business.

How can Druva help?

Druva's cloud-native backup solution was designed to greatly simplify and cut the cost of protecting the VMware-based hybrid cloud.

Druva simplifies data protection by providing a single data model for endpoints, on-premises production data center infrastructure, and cloud applications. At the same time, Druva drives down total cost of ownership (TCO) by up to 50 percent when compared to legacy data protection architectures. Notably, it supports data mobility and tiering across multiple cloud storage tiers, as well as compression and deduplication, for cost efficiency. Specifically, Druva deduplicates source data in order to reduce network bandwidth requirements and to contain the amount of data that must be stored in the Amazon S3 cloud storage service over time. Savings are furthermore enhanced with the application of deduplication globally across all environments. As an additional cost-efficiency measure, the Druva software is applied dynamically as needed.

Because it is truly cloud-native, Druva can scale on demand and infinitely.

Because it is truly cloud-native, Druva can scale on demand and infinitely—from both a capacity standpoint and also from the standpoint of compute and the data protection services themselves. Customers have scaled their Druva implementations into multi-petabyte ranges.

Druva has also built in a number of additional security provisions, that complement the investments made by Amazon Web Services (AWS) to provide failsafe, enterprise-grade functionality. Firstly, it is important to note that Druva has invested in meeting regulations, including HIPAA, SOC-2 and FIPS, beyond those that are supported by the AWS infrastructure itself. Additionally, Druva applies an isolated data model that separates metadata and customer data. Metadata describes the customer's data and is used to locate it as well as for collaboration on files. Separating metadata from customer data can help to avoid metadata inconsistencies, ensuring that data can be properly located. It also can help to ensure that customer data can still be accessed and restored in the event of a cyberattack.

Another trait unique to Druva is the fact that its architecture applies digital envelope encryption, or a two-factor encryption model. Digital envelope encryption provides better security, data privacy, and manageability when compared with alternative methods such as a key and data or an escrow model. Digital envelope encryption generates and encrypts a unique key that is then turned into a token and stored. It also encrypts the data itself—this is the two-factor model. The token can only be decrypted and accessed as a result of a designated administrator or user providing their unique credentials. Any individual would require those credentials in order to access the data. This approach not only provides heightened security and privacy, but it also provides a tremendous amount of visibility into and control over who is accessing data.

Conclusion

Without a doubt, the cloud is a requirement to meet today's data protection demands. However, in order to truly reduce the data protection cost structure and at the same time introduce new, strategic capabilities to the enterprise, the cloud must be leveraged strategically and intelligently.

IT professionals should look for a solution that was purpose-built to run in the cloud and to fully take advantage of the cloud's potential value. Storing data on-premises should not be a requirement, to avoid the need for an expensive on-premises secondary data center. What's more, there should be the ability to intelligently tier data across various cloud storage services for additional cost-efficiency. The disaster recovery solution should also go beyond simply storing data in the cloud. It should take advantage of elastic cloud compute cycles for faster and less expensive failover and to open up the opportunity for additional use cases.

To enable customers to implement the cloud in this way, Druva invested in cloud disaster recovery, which was purpose-built to run in the cloud. Data can be backed up directly to, and it can be booted up directly in, the AWS-based Druva implementation. Data is tiered automatically across various classes of cloud storage services. To further cut the cost structure, we include application-aware compression and deduplication. Additionally, Druva adds a number of capabilities that support security and business agility. The platform itself is agentless and integrates directly with vCenter, for "set it and forget it" automated virtual machine discovery, provisioning and load balancing. Declaring a disaster and failing over to a recovery state is as easy as pushing a button.

Visit druva.com and see how you can evolve your backup strategies for VMware.

druva

Sales: +1 800-375-0160 | sales@druva.com

Americas: +1 888-248-4976
Europe: +44 (0) 20-3750-9440
India: +91 (0) 20 6726-3300
Japan: +81-3-6890-8667
Singapore: +65 3158-4985
Australia: +61 1300-312-729

Druva™ delivers data protection and management for the cloud era. Druva Cloud Platform is built on AWS and offered as-a-Service; customers drive down costs by up to 50 percent by freeing themselves from the burden of unnecessary hardware, capacity planning, and software management. Druva is trusted worldwide by over 4,000 companies at the forefront of embracing cloud. Druva is a privately held company headquartered in Sunnyvale, California and is funded by Sequoia Capital, Tenaya Capital, Riverwood Capital and Nexus Partners. Visit Druva and follow us @druvainc.