Technical Validation

HPE Cloud Volumes: Enterprise Cloud Data Services

Simplifying and Automating Enterprise Storage and Backup in the Cloud

By Kerry Dolan, Senior IT Validation Analyst

January 2021

This ESG Technical Validation was commissioned by HPE and is distributed under license from ESG.

© 2021 by The Enterprise Strategy Group, Inc. All Rights Reserved.

Contents

ESG Technical Validations

The goal of ESG Technical Validations is to educate IT professionals about information technology solutions for companies of all types and sizes. ESG Technical Validations are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objectives are to explore some of the more valuable features and functions of IT solutions, show how they can be used to solve real customer problems, and identify any areas needing improvement. The ESG Validation Team's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments.

© 2021 by The Enterprise Strategy Group, Inc. All Rights Reserved.

Introduction

This ESG Technical Validation documents remote testing of HPE Cloud Volumes, a suite of enterprise cloud data services delivering high-performance, resilient block and backup services that are simple, secure, flexible, and efficient. Testing focused on demonstrating ease of use for running enterprise applications and backup in the cloud, workflow automation, and integration with public cloud compute.

Background

Many organizations today are focused on "digital transformation" to make optimal use of all their data while streamlining costs to drive business forward. According to ESG research, the most-often cited objective for digital transformation initiatives is to become more operationally efficient, followed by improved customer experiences, products, and services (see Figure 1).[1]

Hybrid cloud storage is part of that digital transformation and offers an attractive option that can deliver greater efficiency, elastic scaling, business agility, and payment options—but there are challenges. Many applications require high reliability, availability, data protection, and performance, which are not standard in the cloud. These applications must often be refactored to operate in the cloud, creating complexity and delay. In addition, most cloud solutions charge egress fees to get data back, significantly reducing application mobility. Migration complexity and compatibility issues often result in being locked into a single cloud vendor, reducing agility. Also, managing siloed on-premises and cloud infrastructure separately with different policies is costly and inefficient.

Figure 1. Digital Transformation Objectives

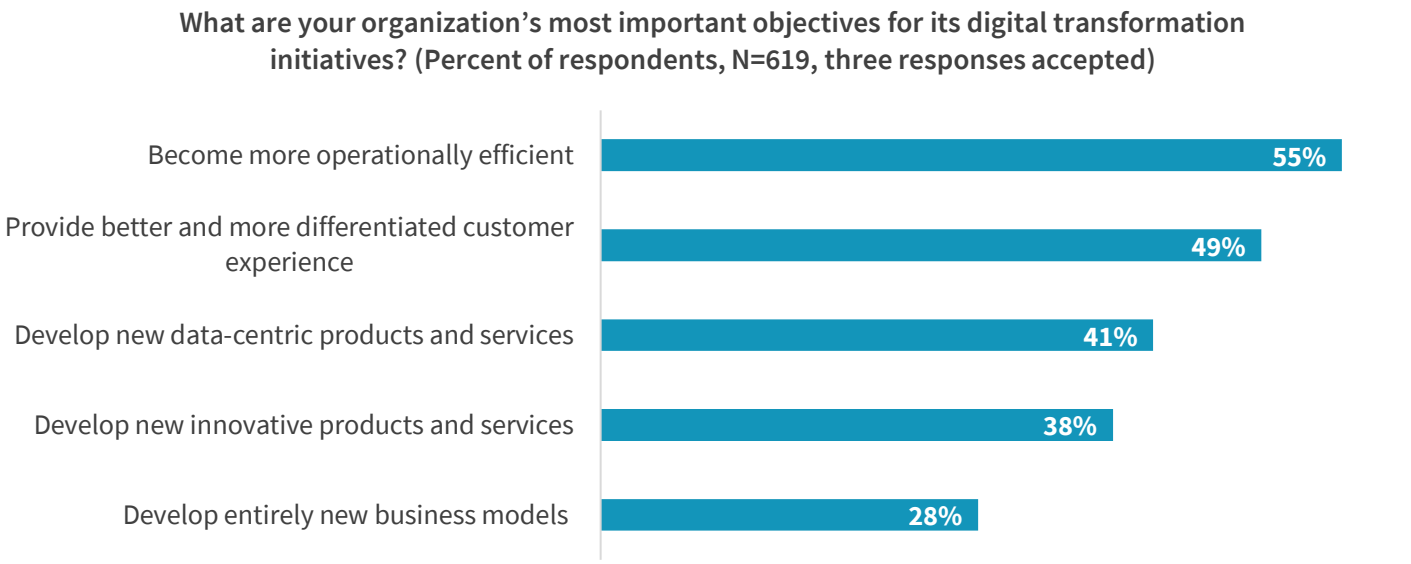

Source: Enterprise Strategy Group

For hybrid cloud storage to be most beneficial, organizations need a unified experience for legacy and cloud-native apps, with simple, automated management. They need enterprise performance and availability, multi-cloud flexibility, and the ability to move data as needed between clouds and between on-premises and the cloud without costly egress charges.

[1] Source: ESG Research Report, *2020 Technology Spending Intentions Survey*, February 2020.

© 2021 by The Enterprise Strategy Group, Inc. All Rights Reserved.

HPE Cloud Volumes: Enterprise Cloud Data Services

HPE Cloud Volumes is a suite of enterprise cloud data services. It provides elastic scale and pay-as-you-grow pricing, while also ensuring that cloud data and applications enjoy the same high reliability, high performance, and protection as on-premises data and applications. Services are delivered using HPE infrastructure in secure, highly available data centers that are located around the world, directly connected to public cloud providers, including Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP). This ensures that customers can easily leverage public cloud compute instances and use low-latency HPE Cloud Volumes storage. HPE Cloud Volumes is supported by HPE InfoSight for predictive analytics to efficiently manage storage.

HPE Cloud Volumes provides a secure, unified data experience with maximum agility. There are two data services: HPE Cloud Volumes Block and HPE Cloud Volumes Backup. Pricing options include monthly consumption-based, discounted up-front payments, and bonus credits.

Figure 2. HPE Cloud Volumes Overview

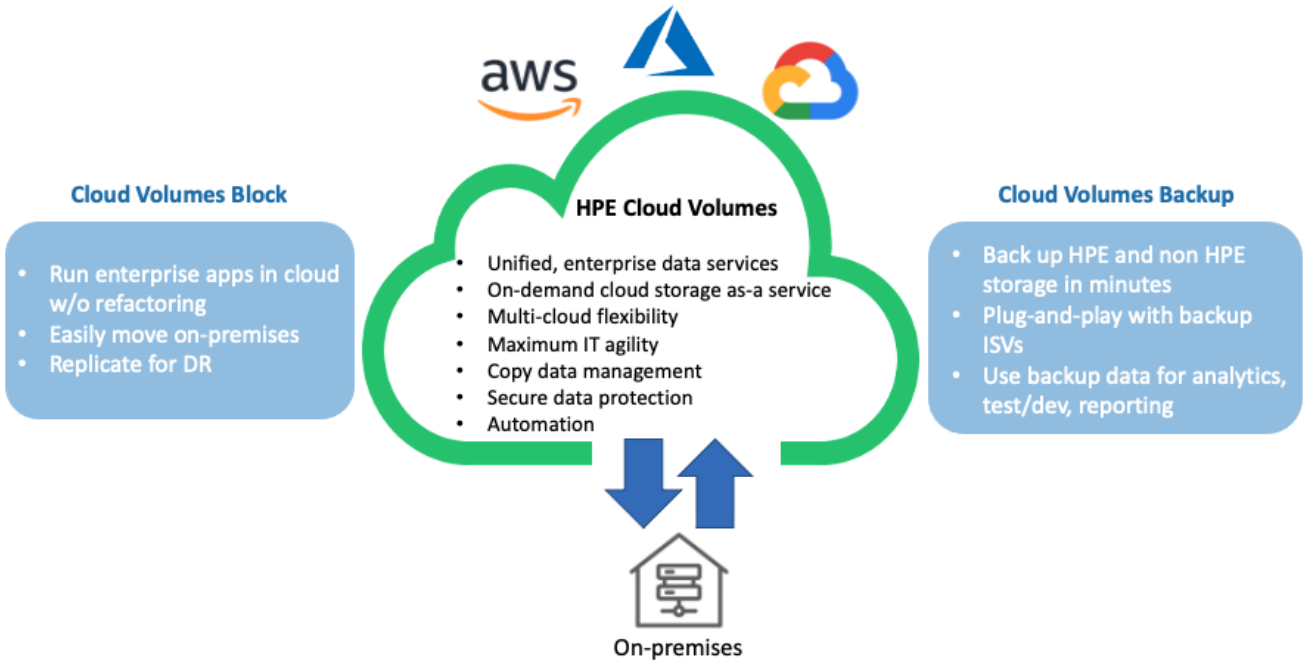

Source: Enterprise Strategy Group

HPE Cloud Volumes Block

HPE Cloud Volumes Block provides highly reliable, high-performance cloud storage that can be provisioned in minutes to run enterprise applications. Data is secure and encrypted in enterprise-grade storage, with the same enterprise features and the same management in the cloud as on-premises. These include the ability to create highly available clusters, copy data management features including instant snapshots and clones, and easy-to-use replication from on-premises storage to HPE Cloud Volumes Block for disaster recovery.

© 2021 by The Enterprise Strategy Group, Inc. All Rights Reserved.

Hybrid Cloud Challenges

Mission critical, enterprise applications need enterprise data services—high availability, data protection, efficiency. But these advanced data services are not available by default in the cloud, and the architectural differences of the cloud require hours or days of rewriting code to run these apps.

This service offers a fast on-ramp to the cloud. HPE Cloud Volumes Block lets organizations run enterprise applications in the cloud with fast provisioning, high reliability, fast performance, and no code rewrites because the architecture is the same on-premises and in the cloud. For cloud-native and microservices-based applications, customers can leverage AWS, Azure, or GCP compute without having to move data. For example, developers can start application development with public cloud compute using Cloud Volumes Block storage; if they choose to bring the application back on-premises for production, the data can be easily migrated with no egress charges. Integration with container orchestrators such as Docker and Kubernetes enables HPE Cloud Volumes Block to be used in fully automated CI/CD pipelines.

Managing HPE Cloud Volumes Block is simple, and the self-service portal lets organizations move data back and forth from the cloud to on-premises easily and enables automated workflows. Examples include using snapshots and clones to make data available to multiple public cloud providers, enabling failover cluster instances across different zones in a region for fault tolerance, and speeding container deployments.

HPE Cloud Volumes Backup

HPE Cloud Volumes Backup is an enterprise cloud backup service, which delivers a simple, efficient, and flexible way to store backup data. Instead of adding backup infrastructures that are complex to deploy and costly to scale and manage, organizations can back up to the cloud. This cloud-native backup service is simple and automated and does not require additional gateways or virtual/physical appliances to backup data to the cloud. HPE Cloud Volumes Backup breaks down typical backup silos, by protecting HPE Storage (including Nimble, Nimble Storage dHCI, 3PAR, Primera, and SimpliVity), using HPE Recovery Manager Central (RMC)[2], and both HPE Storage and non-HPE storage, using leading backup applications from Veeam, Veritas, Commvault, MicroFocus, and others.

With HPE Cloud Volumes Backup, organizations can move backup infrastructure to the cloud regardless of storage type or backup applications *without changing existing workflows*. This service allows customers to avoid public cloud lock-ins and eliminate egress charges, with the flexibility to recover on-premises for operational restores or in the cloud to run other workloads. Cloud-based backups can transform backup data into a business asset to be used for test/dev, reporting, analytics, etc.

ESG Technical Validation

ESG viewed remote demonstrations of both HPE Cloud Volumes Block and HPE Cloud Volumes Backup. Testing was designed to demonstrate how HPE Cloud Volumes makes it easy to deliver all the features needed to run enterprise applications in the cloud. These include ease of use with the single HPE Cloud Volumes portal, automation of workflows, and the simplicity of using HPE Cloud Volumes storage with public cloud compute. Note that while we used the HPE Nimble Storage array for these demonstrations, backup tasks can also be done with other HPE Storage, using HPE RMC, and with HPE Storage and non-HPE storage, using leading backup applications.

[2] HPE Recovery Manager Central (RMC) software integrates HPE 3PAR, HPE Primera, HPE Nimble Storage, and HPE Nimble Storage dHCI with HPE StoreOnce systems and HPE Cloud Volumes Backup to provide a flash-integrated snapshot, replication, and backup service that augments traditional backup approaches either on-premises or in the cloud. Combining the performance of snapshots and replication with the protection of backups, RMC enables fast, efficient, reliable, and simple protection of business-critical applications.

© 2021 by The Enterprise Strategy Group, Inc. All Rights Reserved.

HPE Cloud Volumes Block

HPE Cloud Volumes Block lets organizations quickly and easily provision enterprise-grade storage in the cloud with HPE's Site Recovery Engineering team maintaining the infrastructure. To run enterprise applications, storage must deliver advanced features, including:

- Dynamic scalability to meet changing needs.
- High availability, often configured using HA clusters that maintain uptime despite component failures.
- Data snapshots and replication for robust data protection.
- Cloning to make additional use of data such as for text/dev and reporting.

HPE Cloud Volumes Block provides all of these features for storage in the cloud with the same procedures that are familiar to users on-premises.

ESG Testing

First, ESG viewed several demonstrations of tasks that are important for running enterprise application in the cloud:

- Creating HA clusters in the cloud. We viewed how easy it was to create fault-tolerant cloud storage in multiple zones in the cloud. This process with HPE Cloud Volumes Block was simple to configure, using the same process as Nimble Storage on-premises. Two Linux nodes were handling an MS SQL Server workload, and the database cluster was backed by HPE Cloud Volumes Block shared storage, housed in two different zones in the US West region. ESG watched as one node was powered off; the workload continued to run, immediately failing over to the other node, without interruption.
- Multi-cloud app migration. In this demonstration, a stateful, containerized MS SQL Server application using HPE Cloud Volumes Block storage was moved between AWS and GCP. The application running on AWS Elastic Kubernetes Services (EKS) was moved to Google Kubernetes Engine (GKE), including the persistent storage, without disruption and without egress charges. This capability enables multi-cloud production capability, as well as hybrid cloud disaster recovery and continuous integration/continuous development (CI/C) pipelines.
- Hybrid Cloud CI/CD Pipeline. This demo started with a microservices-based, GKE production workload on HPE Nimble Storage dHCI on-premises; HPE Nimble dHCI was continually replicating to HPE Cloud Volumes, which created a clone. This clone was used in a new cluster as part of a Jenkins pipeline, which then automatically deployed back on-

Customer Success

A large financial services organization was working on a cloud-first strategy but having some difficulty. In addition to a year-long data migration, this company wanted, in the cloud, the same configuration on which they depended for fault tolerance on-premises: high availability failover clustering for enterprise applications, configured so that any failover would not impact name resolution or clients reaching that access point. While this was easy to configure on-premises, their cloud provider could not accommodate that.

The customer chose HPE Cloud Volumes because it provided the same capability in the cloud as on-premises. The company felt much more comfortable with HPE Cloud Volumes because the experience was the same as on-premises, compared with piecing together different services from another cloud provider.

In addition, HPE Cloud Volumes delivered:

- A simpler, faster data migration path that eliminated the multiple hops of another cloud provider, also eliminating egress charges.
- A file share failover cluster across multiple zones with the DNS, Active Directory, and load balancing configuration they wanted.

© 2021 by The Enterprise Strategy Group, Inc. All Rights Reserved.

premises into production. In today's world of continuous integration, the development environment is no longer separate but rather is really part of production. HPE Cloud Volumes Block delivers storage that can be easily integrated into that lifecycle.

Replication Workflow

The first step is getting data into the HPE Cloud Volumes cloud. HPE Cloud Volumes Block provides native replication directly from HPE Nimble (and HPE Nimble dHCI) all-flash storage. This replication provides a hot DR copy and enables data usage faster than having to restore. Data can be restored to the same or a different array or can be cloned for presentation to any of the hyperscalers. For other HPE storage, seeding data in the cloud requires simply backing up to HPE Cloud Volumes Backup and then restoring to HPE Cloud Volumes Block. In this testing, we demonstrated the replication workflow, the cloning workflow, and how to present the clone to a public cloud compute VM.

First, ESG reviewed a demonstration of the replication workflow from an on-premises HPE Nimble storage array to the HPE Cloud Volumes Block service. From the HPE Cloud Volumes portal, we clicked ***Replication Stores*** in the left navigation bar, and then clicked ***+New Store*** to create a location for the data to land in HPE Cloud Volumes Block. We entered a name for it (*sqldatabases*) and selected 1024 GB capacity. Customers can replicate single or multiple volumes, from single or multiple arrays, into a replication store. Next, we selected the type of flash, as customers would, based on their workload SLAs. HPE Cloud Volumes storage is backed by HPE hybrid and all flash arrays and offers Premium or General-Purpose flash. Note that once we began configuring this replication store, a monthly price estimate was displayed, which changed according to our selections.

Figure 3. Create New Replication Store

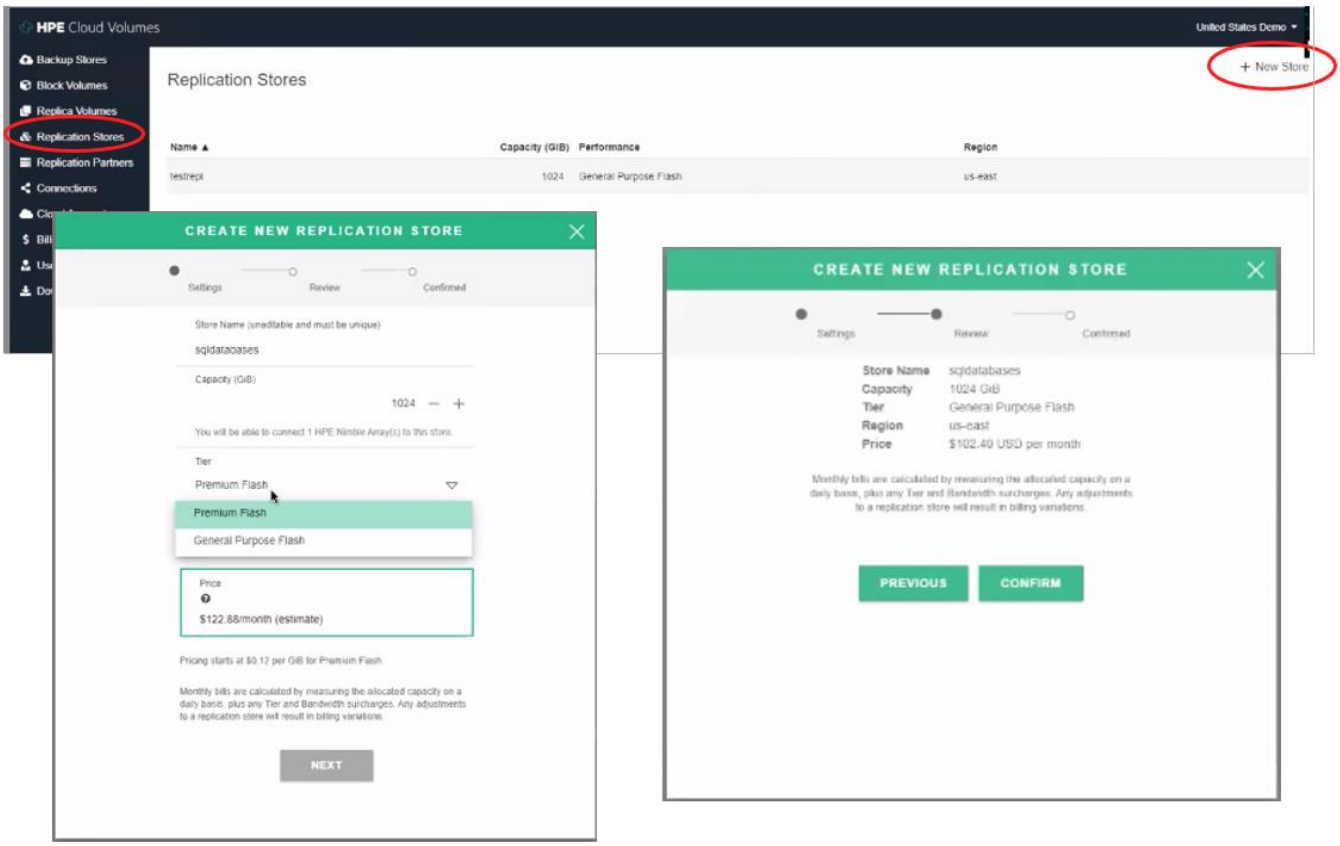

Source: Enterprise Strategy Group

© 2021 by The Enterprise Strategy Group, Inc. All Rights Reserved.

The next task was to create a replication partner (i.e., to choose the on-premises array and volumes from which to replicate). Back on the HPE Cloud Volumes portal, we clicked on the ***Replication Partners*** tab; here, users can create a new partnership or modify a current one to link to the replication store just created. We provided a host name (for the management endpoint for the on-premises HPE Nimble array) and entered the credentials for that endpoint. Credentials are not stored here; we clicked ***Submit***, and that triggered a secure API call from HPE Cloud Volumes Block to the on-premises array to enable replication.

At this point, the connection between the on-premises array and HPE Cloud Volumes Block had been established; the next step was to tell the array which volumes to replicate. On the HPE Nimble Storage UI, we clicked on ***Replication Partners***, where our new replication store was already displayed and identified with a ***Cloud*** tag. We selected ***Volume Collection***, which defines volume-level consistency groups and enables scheduling of local and remote snapshots. Application integration enables application-consistent copies.

We selected a previously created Volume Collection called *developer1-AQL.* To configure HPE Cloud Volumes Block replication, we clicked the drop down menu by ***Replication Type***, selected ***Periodic Snapshots***, selected the new ***Replication Partner*** *(sqldatabases)*, and configured the retention time for 90 days. Now, the next time the local snapshot executes, it will also automatically synchronize with HPE Cloud Volumes Block using a secure SSH tunnel. Figure 4 shows the configured HPE Cloud Volumes Block replication, application consistency, schedule, and retention.

Figure 4. Configuring and Scheduling Replication

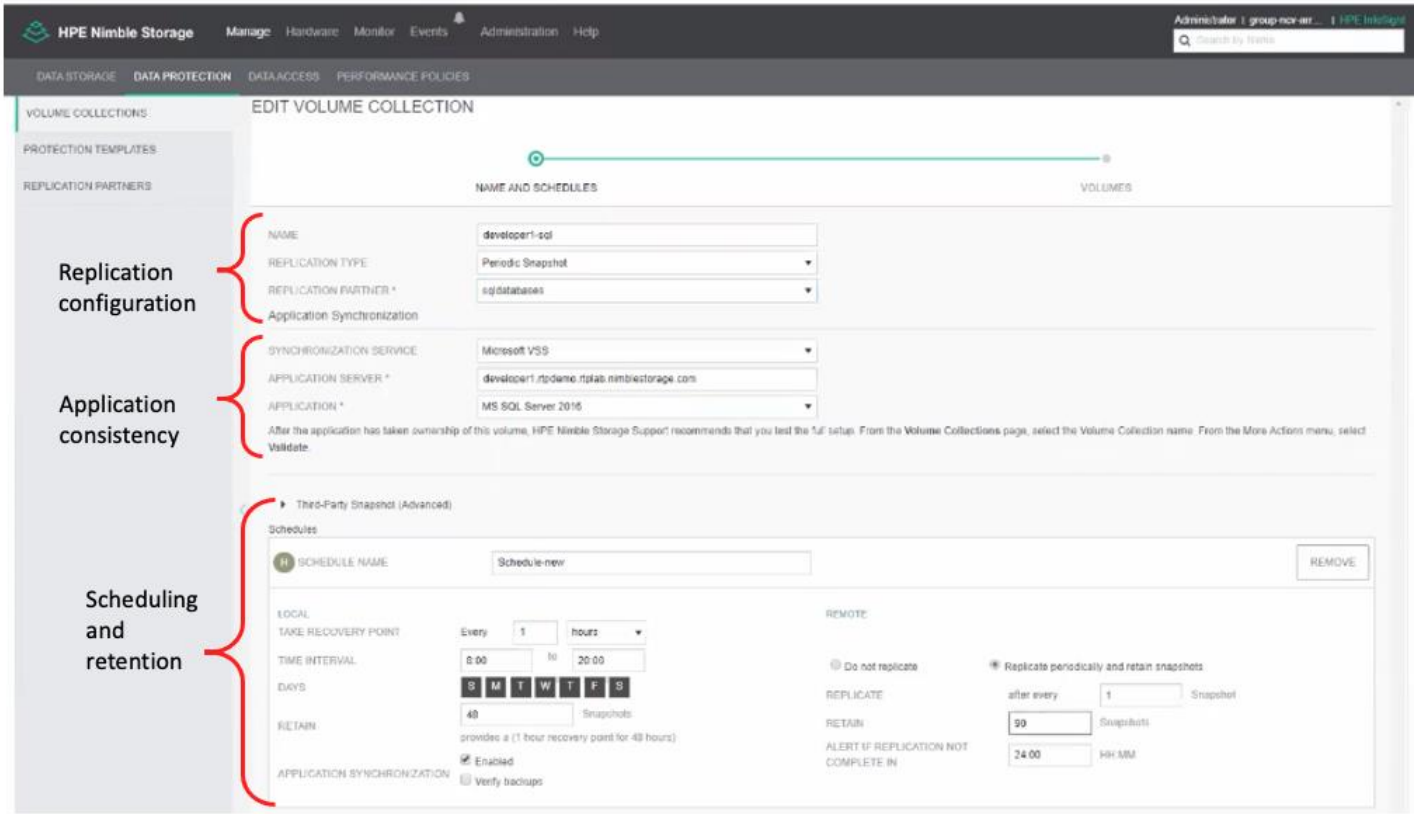

Source: Enterprise Strategy Group

© 2021 by The Enterprise Strategy Group, Inc. All Rights Reserved.

Clone Workflow

Next, we viewed the cloning workflow, which makes a copy of this replica for various use cases such as DevOps, analytics, and reporting. From the HPE Cloud Volumes portal, we viewed ***Replica Volumes***, showing the volumes in the cloud, some with incoming replication updates and some without. We selected a SQL database snapshot and, in a few clicks, configured clone capacity and performance; note that performance is not inherited from the parent. Customers can dictate both capacity and performance independently. For example, the production volume may be configured for 20K IOPS, but a DevOps clone can be created with only 1K IOPS, saving money.

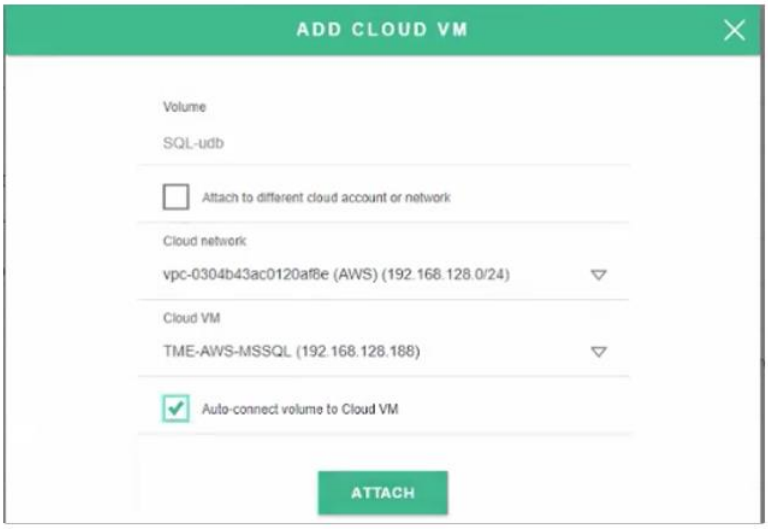

Next, we chose an AWS account and network to which to present the volume. HPE Cloud Volumes Block does this by automatically creating a dedicated iSCSI direct connection to the VPC for AWS, express route for Azure, or interconnect to the VPC for GCP. Once this connection is complete, the customer can select the specific machine to connect to within the VPC.

Next, we created a snapshot schedule to protect this new clone. At this point, the clone was viewable in the list of volumes in HPE Cloud Volumes Block, showing the cloud provider, region, network, and subnet.

Finally, we assigned the ACLs to attach this volume to an AWS compute instance. If we wanted to change to Azure, that would be as simple as clicking a box and choosing an Azure VNet. We chose an AWS VM, selected ***Auto-connect volume to Cloud VM***, and clicked ***Attach,*** and once we refreshed the AWS instance, the database was available.

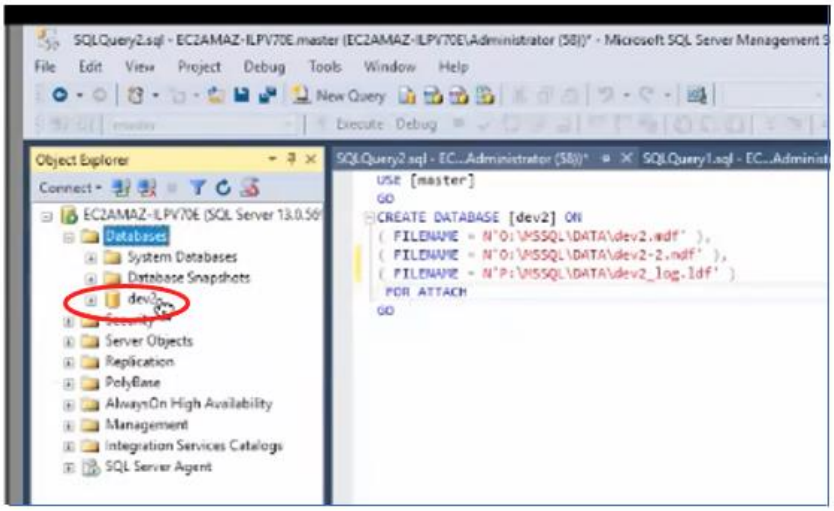

Note that we were able to not only present the volume to the compute instance using the HPE Cloud Volumes Block portal but also to make the required API calls to bring the volume online and make it available for use. This is enabled during initial HPE Cloud Volumes account creation, when customers set up delegate access to their cloud providers.

© 2021 by The Enterprise Strategy Group, Inc. All Rights Reserved.

Why This Matters

Despite the cloud's tempting agility, scalability, and flexible payment benefits, many organizations have avoided it because of the need to change processes, refactor applications, and learn new ways of doing everyday tasks. Many cannot afford the pain and business interruption of change. Others have moved to the cloud without realizing the difficulty of adding enterprise storage features and then repatriated applications back on-premises.

With HPE Cloud Volumes, customers don't have to manage storage infrastructure—they can just consume it and still get the features needed to run enterprise applications safely and securely. ESG validated the ease of creating highly available clusters in the cloud, app portability across clouds, and implementing a hybrid cloud CI/CD pipeline. We also validated that in a few minutes, customers can easily migrate workloads to the cloud to free up on-premises capacity and create a DR copy. With a few clicks, we created a replication partnership between an on-premises array and HPE Cloud Volumes with an automated workflow. We cloned a database volume and presented it to an AWS compute instance without complexity.

This level of ease and automation delivers a simple on-ramp to the cloud for enterprise applications and also enables customers to bring data back on-premises without egress fees.

HPE Cloud Volumes Backup

Through the same portal, HPE Cloud Volumes Backup enables backup using HPE RMC for HPE Storage or leading backup applications from Veeam, Commvault, Veritas, Microfocus, and others for both HPE and non-HPE storage. This provides a secure, offsite, backup target that is invisible to ransomware because of highly secure HPE Catalyst protocol. Ransomware can't infect and encrypt what it doesn't see; HPE Cloud Volumes Backup creates backup stores which are not directly accessible by the operating system. Backup images are invisible and inaccessible to ransomware, thereby ensuring data integrity and making restores possible in the event of a ransomware attack.

Data is encrypted in flight and at rest when residing in HPE Cloud Volumes Backup. The backup process includes source-side deduplication and compression for efficiency; customers are charged only for consumed capacity in the cloud. HPE RMC-based backups can be fully automated and restored to either on-premises arrays or to the HPE Cloud Volumes Block service, delivering both data protection and backup data reuse. Backups using other applications leverage the HPE StoreOnce appliance and can be restored automatically to on-premises arrays or manually to HPE Cloud Volumes Block.[3]

ESG Testing

The HPE Cloud Volumes Backup demonstration used the HPE Nimble array and RMC, which orchestrates local snapshots and replication between on-premises arrays as well as backups to HPE Cloud Volumes Backup. Testing included creating a backup store, restoring from HPE Cloud Volumes Backup to HPE Cloud Volumes Block, and integrating with ISVs.

Create Backup Store

From the HPE Cloud Volumes Portal ***Backup Stores*** view, we created a backup store in the cloud similar to the way we created the previous replication store. In just a few clicks, we provided a name and description, chose a region, and created the store. No capacity configuration was needed; HPE Cloud Volumes Backup creates and expands storage automatically as needed.

[3] HPE also offers Cloud Bank Storage, a licensed feature of the HPE StoreOnce appliance that provides a cloud-based object store for long-term retention/archiving of data in the cloud. HPE Cloud Volumes Backup offers direct-to-cloud backup-as-a-service with or without the StoreOnce appliance, and data is stored on HPE enterprise-grade hybrid or all-flash storage. This results in two backup tiers with different retention periods and price points.

© 2021 by The Enterprise Strategy Group, Inc. All Rights Reserved.

Different from some data protection-as-a-service solutions, with HPE Cloud Volumes Backup, customers are charged only for consumed capacity; there is no upfront license charge or egress charges when restoring from HPE Cloud Volumes to on-premises. Source-side deduplication and compression reduce the capacity in the cloud, minimizing costs up to 20x, according to HPE. In addition, backups are stored on fault-tolerant, highly available, hybrid and all-flash HPE storage, rather than on hyperscaler storage with lower SLAs.

After creating the backup store, we downloaded authentication credentials; this information is used to configure the backup applications. Data is encrypted at rest, and these credentials are needed to decrypt the data. Next, we configured the backup application to write to this backup store. From the RMC UI, we clicked ***Add Storage Device*** and selected ***Cloud Volumes Backup***, the region, the backup store, and the authenticated credentials. In those few steps, we added HPE Cloud Volumes Backup as a repository available to HPE RMC.

Once this was complete, we created a new copy policy (i.e., backup job) in HPE RMC. Copy policies govern creation, scheduling, and retention of local and remote recovery points; we created one that included automatic creation of an HPE Cloud Volumes Backup copy to the repository we just configured (see Figure 5).

Figure 5. Configure HPE RMC Copy Policy to Write to HPE Cloud Volumes Backup

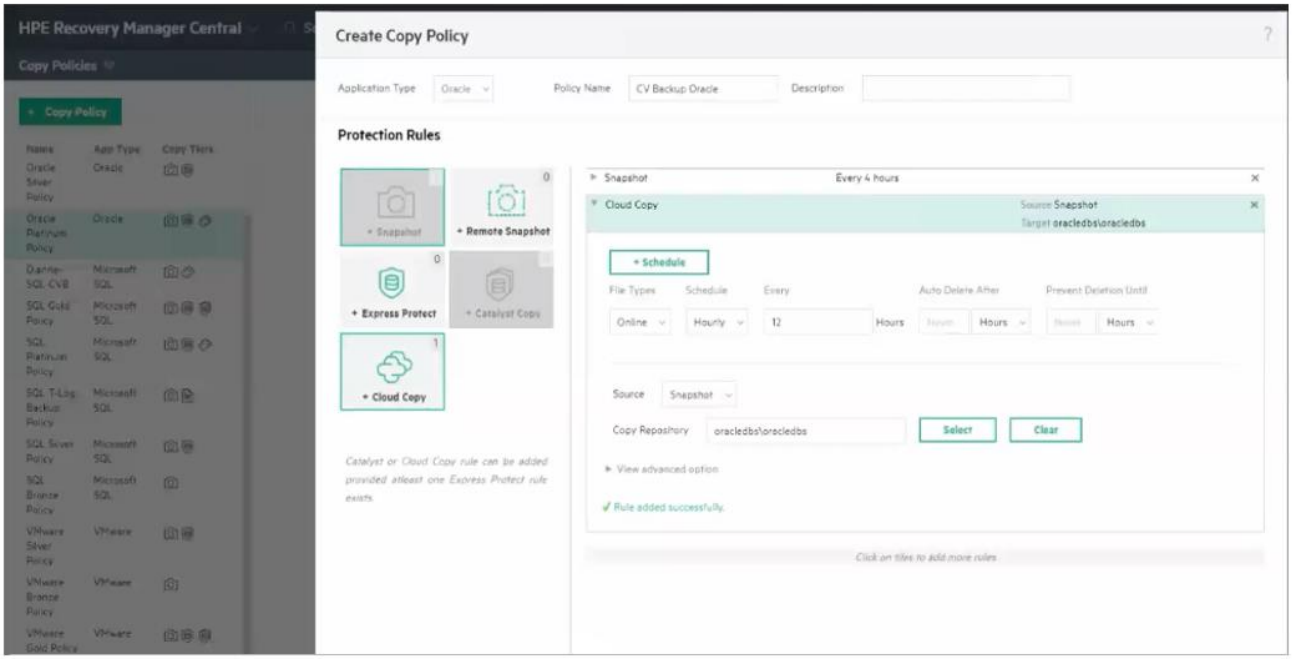

Source: Enterprise Strategy Group

Restore from HPE Cloud Volumes Backup to HPE Cloud Volumes Block

We also viewed a demonstration of restoring from HPE Cloud Volumes Backup to HPE Cloud Volumes Block, using a previously created SQL database backup. Back at the HPE Cloud Volumes portal, we selected a backup store and clicked **Restore to CV Block.** We selected the volume from which we wanted to restore, the specific snapshot to restore, and the volume on HPE Cloud Volumes Block to restore to. Once that was complete, we connected that restored volume to an AWS SQL Server instance for multi-purposing, exactly as we had done previously.

© 2021 by The Enterprise Strategy Group, Inc. All Rights Reserved.

Figure 6. Restore from HPE CV Backup to CV Block and Use on AWS

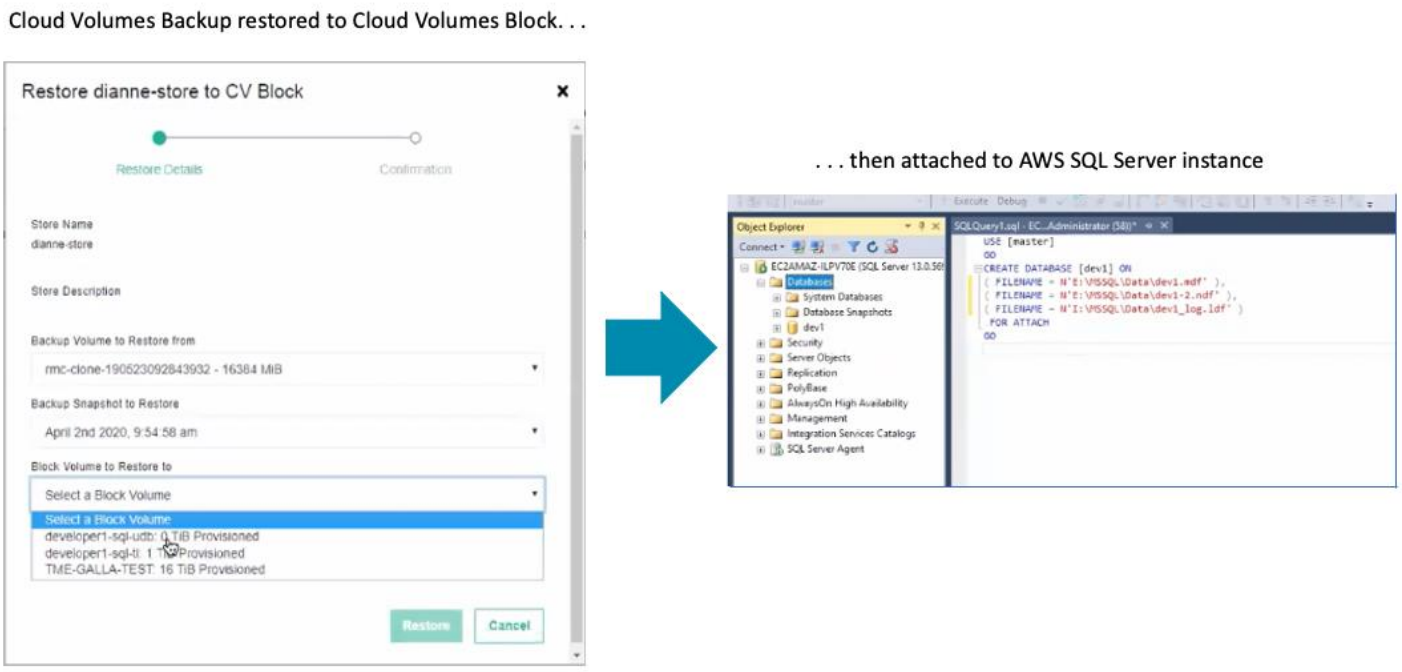

Source: Enterprise Strategy Group

Backup ISV integration

Finally, we demonstrated how easy it is to integrate with other backup applications. From the HPE Cloud Volumes portal ***Backup Stores*** view, we downloaded the secure client. This is not necessary with HPE RMC because HPE built the encryption processes into it, but for other backup applications, a universal plug-in handles that encryption. After downloading the secure client to a Linux endpoint (which can be either on-premises or in the cloud), we extracted it, copied the files into their own dedicated directory, and ran it.

Once that is complete, customers simply navigate to their backup application UI to make it aware of the HPE Cloud Volumes Backup target, which they do in the same way they connect to other targets. They add the IP address of the secure client, add credentials for authentication, select the repository, and apply. Then this target is available to be added to any backup policies.

© 2021 by The Enterprise Strategy Group, Inc. All Rights Reserved.

Figure 7. Integration with Backup ISVs

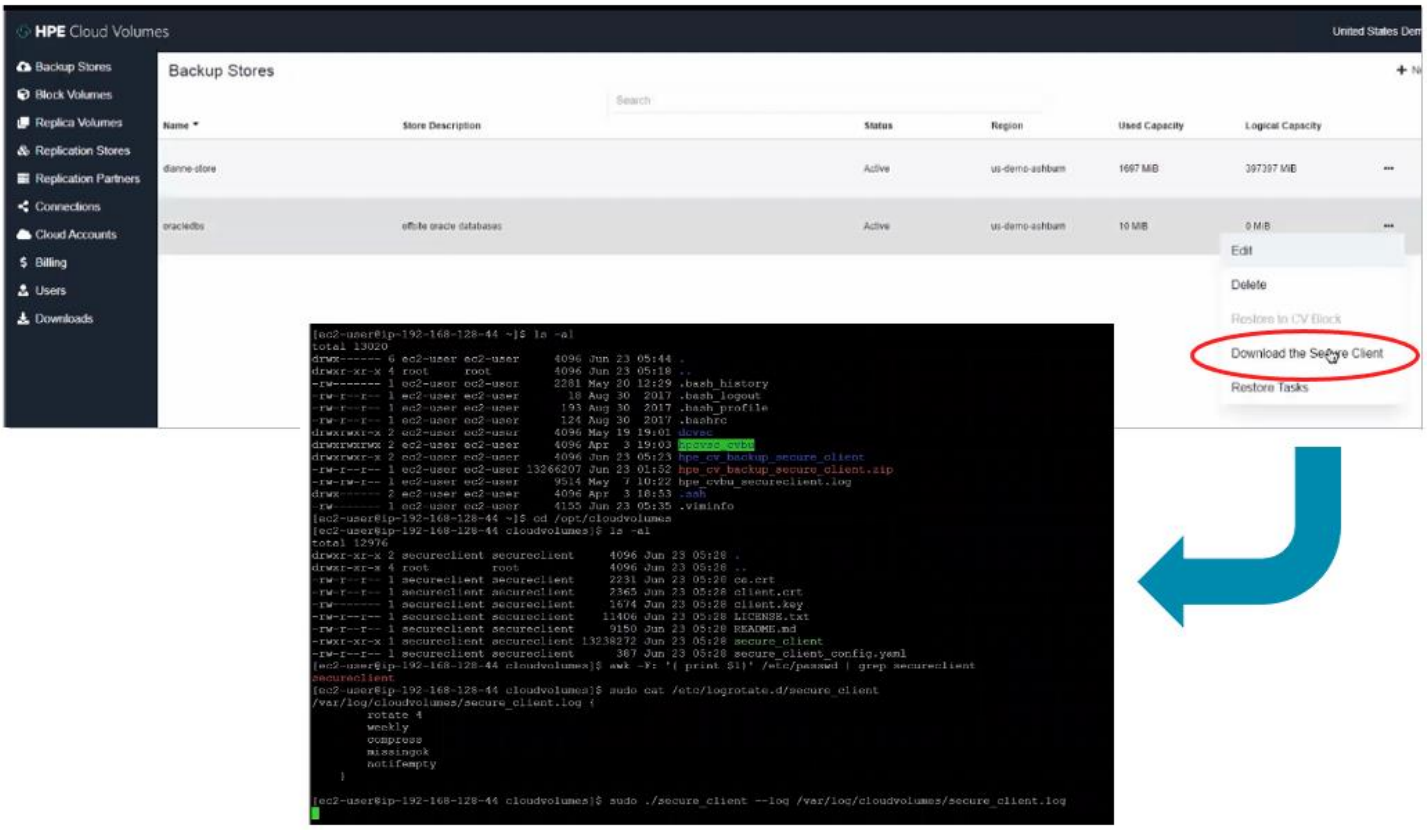

Source: Enterprise Strategy Group

Why This Matters

For many organizations, once backup processes are in place, they don't want to change them. So backing up to the cloud, while offering attractive benefits, can be intimidating. In addition, having separate backup applications and policies for on-premises and cloud data increases complexity and cost.

HPE Cloud Volumes Backup provides simple, automated, pay-as-you-grow backup to the cloud, integrating easily with existing on-premises applications and workflows. This cloud backup service eliminates complexity by freeing customers from the day-to-day hassles and operational costs of backup infrastructure with consumption-based pricing and efficient data mobility. Organizations can spin up storage capacity in the cloud in minutes and integrate with existing data workflows without having to initialize, configure, manage, or tune any physical or virtual infrastructure. In addition, they can restore data back on-premises without typical cloud egress charges and can make backup data into a business asset by using the public cloud for test/dev, reporting, analytics, etc.

ESG validated the simplicity of HPE Cloud Volumes Backup, including creating a backup store, restoring to HPE Cloud Volumes Block, and integrating with backup ISVs. Customers simply log in, provision the HPE Cloud Volumes Backup target with a few clicks, and add it to their existing protection schedules.

© 2021 by The Enterprise Strategy Group, Inc. All Rights Reserved.

The Bigger Truth

IT administrators and application owners are accustomed to the benefits of on-premises enterprise storage, with all the low latency, high performance, high availability, and copy data management features built in. They configure highly available clusters that failover to prevent disruption, use copy data management features for data protection, and count on dynamic scalability of capacity and performance to match the needs of growing, changing businesses. When they go to the cloud, they often do not realize that public cloud storage typically does not offer these enterprise-grade features by default; nor do they realize that code rewrites are required to run in the cloud. If you want to create a clone on premises with, say, HPE Nimble storage, it's an instant process. Doing that in the cloud? That requires first taking a snapshot from cold block storage and fully restoring it to be functional, incurring egress charges as well as delays.

HPE Cloud Volumes Block and HPE Cloud Volumes Backup offer all the features of enterprise-grade storage, delivered as a service in the cloud. That means data security, low latency, easy data movement without vendor lock-in, instant snapshots/clones, and integrations with not only public cloud compute services but also container orchestrators. All of this is quickly provisioned, scalable, and fully managed by HPE.

ESG validated that:

- *HPE Cloud Volumes Block* enables the configurations that enterprise applications need to run in the cloud, such as highly available clusters, multi-cloud app migration, and hybrid cloud CICD pipelines. It also enables automatic replica creation for use in disaster recovery or for data reuse, such as reporting, DevOps, and analytics. Many organizations use multiple public clouds; for example, they might have a production pipeline in Azure and do CI/CD in AWS. Instead of creating, managing, and paying for storage (without enterprise services) separately in these clouds, HPE Cloud Volumes enables IT to provision storage in HPE Cloud Volumes and present it to any public cloud compute instance. No data migration, no complexity, *and no egress fees* to bring your data back on-premises.

- *HPE Cloud Volumes Backup* simplifies and automates cloud backup from any HPE or non-HPE storage array, integrating with existing workflows using HPE Recovery Manager Central or leading backup applications from Veeam, Veritas, Commvault, MicroFocus, and others. This service frees up administrators from daily backup tasks. Data can be restored to any on-premises array or to HPE Cloud Volumes Block for backup data reuse (test/dev, reporting, analytics, etc.), with no egress fees to restore data from the cloud.

HPE Cloud Volumes Block and HPE Cloud Volumes Backup provide an easy, fast way to get your data to the cloud with the security, high performance, availability, and protection of enterprise storage. Organizations can now run enterprise apps in the cloud without sacrificing these enterprise attributes; those working on a "cloud-first" strategy can get there without compromise. There are some additional automation integrations that could increase the solution's value. But if you're looking for an easy on-ramp to the cloud and enterprise storage and backup as-a-service that you don't have to manage, ESG suggests taking a good look at HPE Cloud Volumes Block and HPE Cloud Volumes Backup.

© 2021 by The Enterprise Strategy Group, Inc. All Rights Reserved.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.

Enterprise Strategy Group is an IT analyst, research, validation, and strategy firm that provides market intelligence and actionable insight to the global IT community.

© 2021 by The Enterprise Strategy Group, Inc. All Rights Reserved.

www.esg-global.com

contact@esg-global.com

508.482.0188