Datrium DRaaS with VMware Cloud on AWS

Disaster Recovery as a Service and Cloud Backup

WHITEPAPER

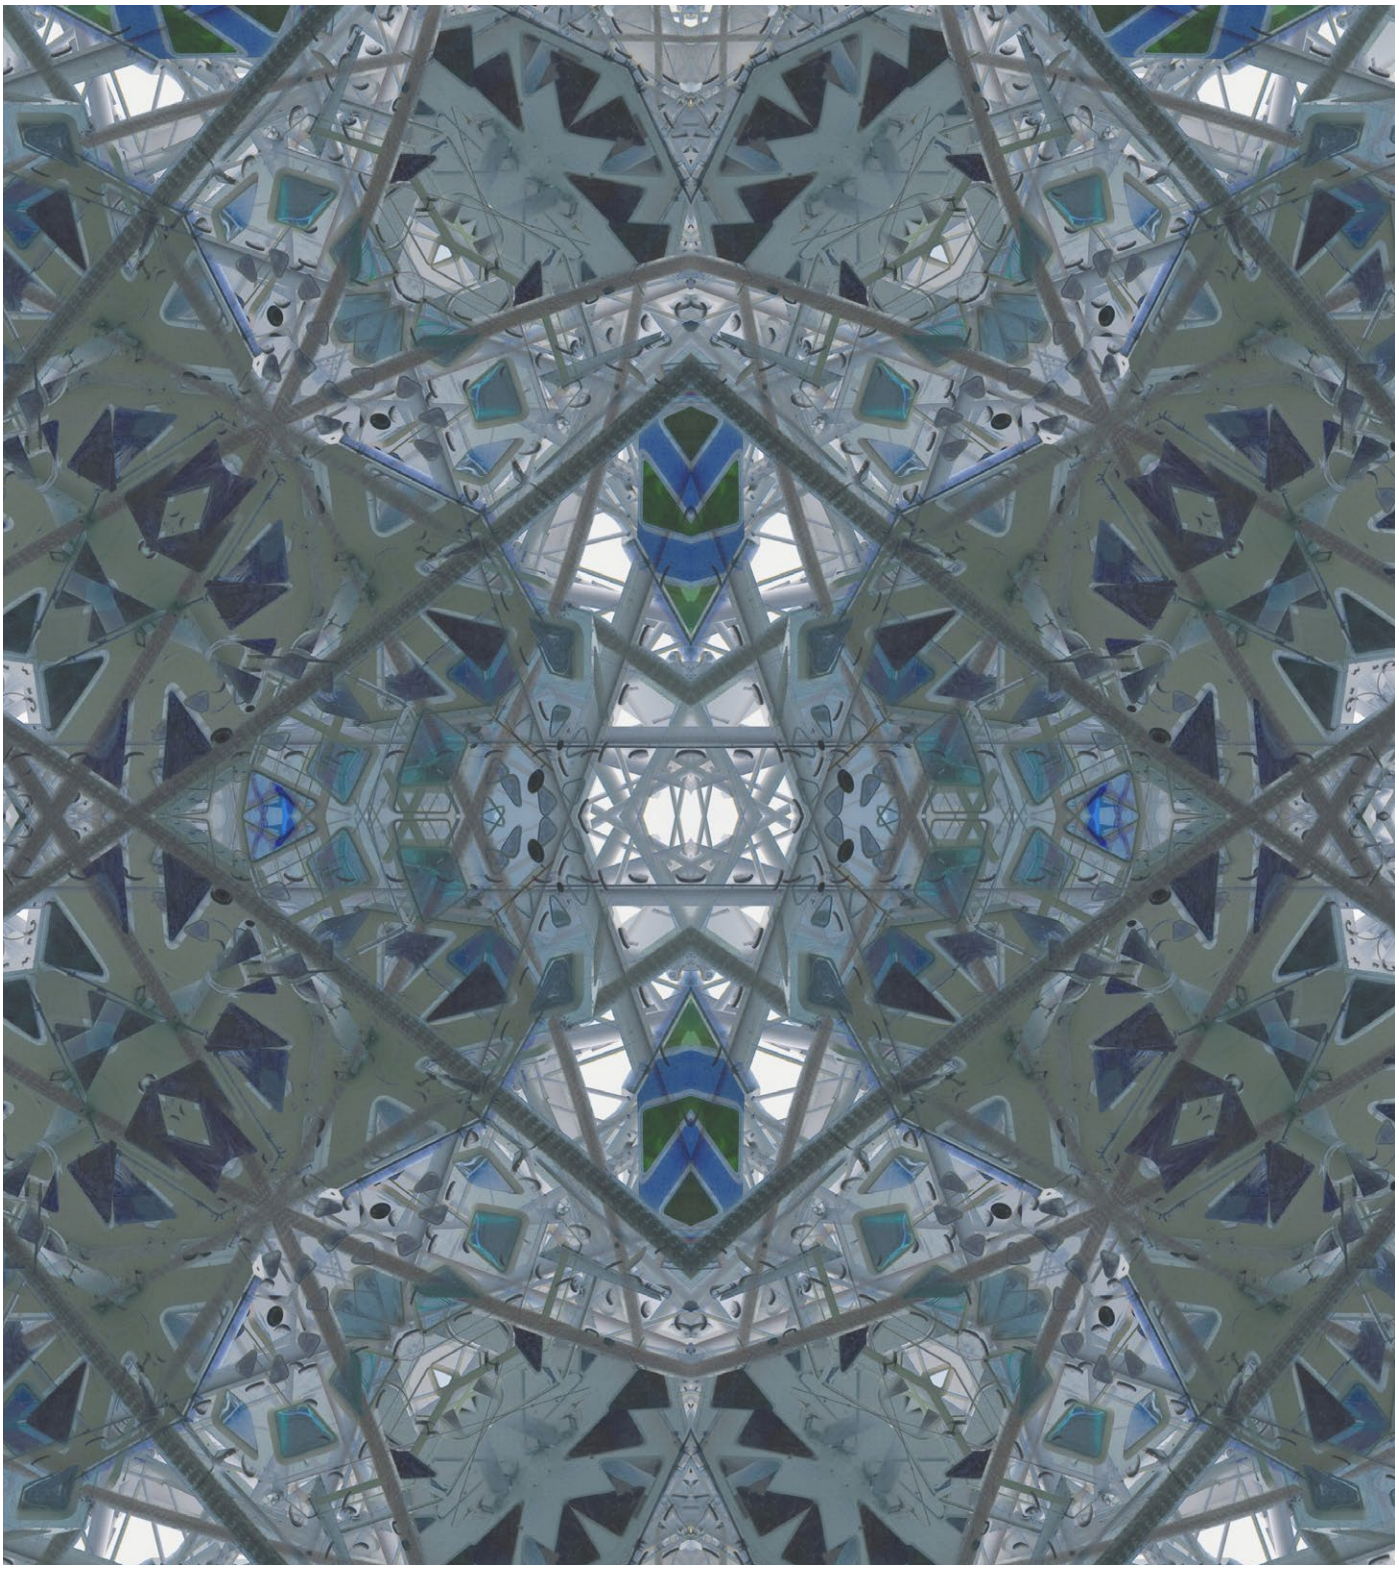

Introduction

Until now, organizations have had to deal with a complex maze of Disaster Recovery (DR) solutions by integrating components and products from several different vendors. This issue has led to increased complexity, ample opportunity for misconfiguration, and staggering levels of resource inefficiencies due to multiple data transformations. DR has caused a financial burden by requiring two data centers, with the DR site managed by the organization or a Managed Service Provider. Teams were forced to contend with the complexity of multiple tools and user manuals, difficult upgrades, and a steep learning curve. And DR solutions were built to address on-premises workloads and did nothing to address workloads running in the cloud. As Steve Duplessie, Founder and Senior Analyst at ESG said, "Until now, Disaster Recovery has been more disaster than recovery."

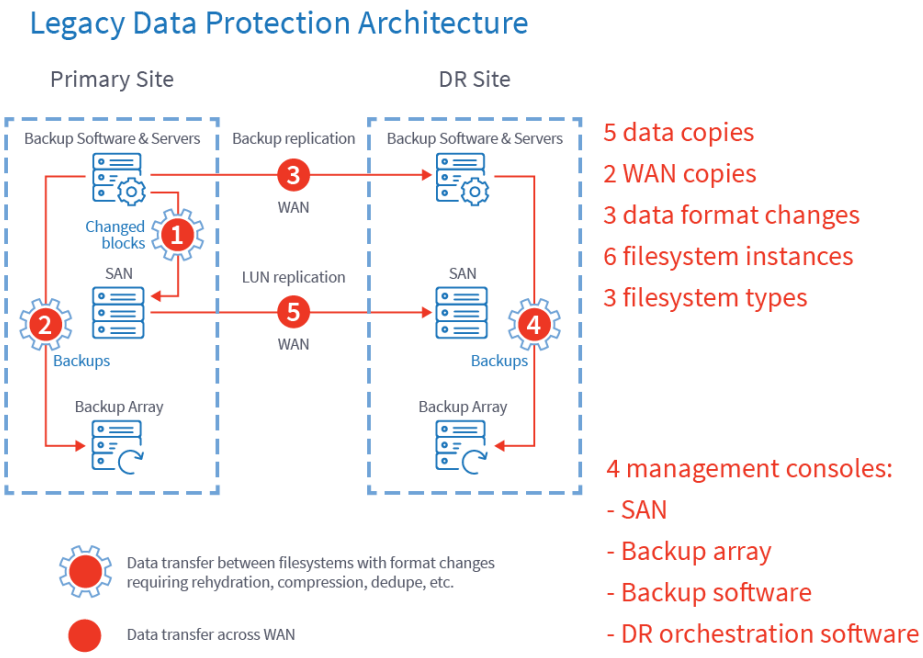

Datrium DRaaS with VMware Cloud on AWS changes this situation. It provides failproof, on-demand DR for all VMware workloads on VMware Cloud or on any VMware workload in on-premises data centers. It includes cloud backup, DR orchestration, and DR sites in VMware Cloud on AWS as fully managed services from Datrium. It's a pure SaaS product with a cloud-native design that eliminates all the complexity of packaged software.

In this whitepaper, we'll discuss the key features of Datrium DRaaS, including:

- 10x lower DR costs from an on-demand data center in the public cloud
- More efficient IT with consistent VMware operating environment in the cloud and on premises
- Instant restart of any VMware workload from cost-effective S3, after a ransomware attack or other disaster
- Recent snapshots or old backups to restore data with instant RTO and no rehydration overhead

Why Cloud-Native SaaS for VMware Cloud on AWS. Legacy DRaaS requires custom, private data centers, such as SunGard and iLand, that have capital-intensive business models with significant financial risk. Public clouds at scale don't carry financial risk and offer a wide selection of regional data centers worldwide. Organizing those facilities using a purely software-defined model such as Datrium DRaaS is the safest, long-term approach for IT. VMware Cloud on AWS eliminates the overhead involved in converting to a new type of cloud and administrative model. And Datrium takes all the guesswork out of getting started by offering a single point of administration and billing for all services.

Responsibility	Datrium	Customer
Service contracts and billing, including AWS[1], VMware Cloud accounts	✓	
Single point of support for whole stack	✓	
Managing DR orchestration uptime and compliance	✓	
Backup vaulting as a service (Datrium Cloud DVX on S3)	✓	
Automating user-defined DR plans	✓	
Terms of Service for backup and DR plan execution	✓	
Creating DR plans and internal DR SLAs		✓
Purchasing the right plan for SLA needs		✓

1 *Optionally, AWS contracts and billing can be managed by customers directly. Please contact your Datrium representative for more details.*

Key Building Blocks

Cloud Backup

Cloud Backup is automatically built into DRaaS. It can be used to automatically backup any VMware workload and uses S3 to persist a durable repository of backups stored in a native-compressed and deduplicated form. Global deduplication ensures that data is stored only when it is unique across origins. When a restart is required on VMware Cloud, Cloud Backup includes instances with flash which can project a live datastore based on its snapshots, so the SDDC can boot an image with instant RTO and no rehydration overhead.

ControlShift

ControlShift is a DR orchestration service that runs in AWS and executes DR plans from new or old backups. Within DRaaS, ControlShift provisions and monitors SDDCs in VMware Cloud on AWS. DR plans and states are stored in a highly-available plan database (Amazon Aurora) replicated across multiple availability zones (AZs). Built-in self-healing ensures that, in the event of public cloud unavailability, all affected Datrium services automatically migrate to a healthy AZ without any data loss.

DRaaS Connect

DRaaS Connect, a feature of DRaaS, is a downloadable, lightweight VM that enables customers to protect any VMware workload in just minutes with no new software or infrastructure to install. DRaaS Connect enables DRaaS to orchestrate failover from a VMware Cloud SDDC in one AWS AZ to another AZ or from any on-premises vSphere infrastructure, including SAN, NAS, HCI, and DHCI to a Datrium-managed VMware Cloud on AWS.

DVX

DVX is Datrium's market-leading DHCI solution. It combines the robustness of SANs with the simplicity of HCI systems. Its built-in backup, primary storage, and encryption capabilities, as well as its ability to scale compute and performance independently, make it a top infrastructure choice for leading enterprises.

VMware Cloud on AWS

VMware Cloud on AWS provides an on-demand VMware Software-Defined Data Center (SDDC), which is used by DRaaS as a cloud DR target. ControlShift, a cloud DR orchestrator and component of Datrium DRaaS, can provision an SDDC with different trade-offs in runbook RTO and prices (see VMware Cloud SDDC Deployment Modes below). A provisioned SDDC incurs hourly charges. Upon DR test completion, the SDDC can be decommissioned in the ControlShift UI. ControlShift performs automated network configurations for both AWS and VMware Cloud to make S3 backups available for spin-up in SDDC. The SDDC is managed in the familiar vCenter interface.

Production Site Options

Customers can use three deployment models to address different production site types for DRaaS, as shown below, either with DRaaS Connect or using Datrium DVX native functionality. Now, all VMware users have access to cost-effective, reliable, cloud-based DR with instant RTO regardless of where their workloads are running.

VMware Cloud on AWS Workloads	VMware Workloads in On-Premises Data Centers (SAN, NAS, HCI)	VMware Workloads in On-Premises Data Centers on Datrium DVX DHCI systems
DRaaS Connect for VMware Cloud replicates vSphere snapshots to Cloud Backup in S3 for later failover from an SDDC in one AWS AZ to another AZ.	DRaaS Connect for vSphere On Prem replicates vSphere snapshots to Cloud Backup in S3 for later failover to VMware Cloud and failback to on-premises data centers.	DRaaS Connect is built into Datrium DVX and replicates VM snapshots to Cloud Backup in S3 for later failover to VMware Cloud and failback to on-premises data centers.

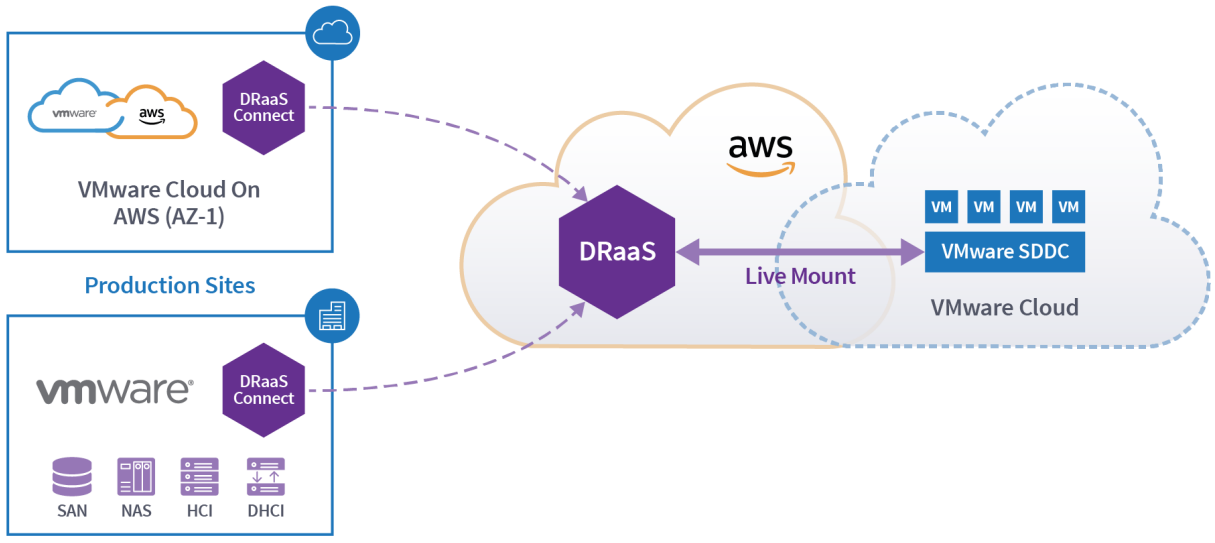

DRaaS Features

Failproof DR for All Your VMware Workloads

Get DR with instant RTO for all of your VMware workloads wherever they're running. Automatically fail over from an SDDC in one AZ to another with DRaaS Connect for VMware Cloud. Get DR with instant RTO for all vSphere on-premises infrastructure, including Datrium DVX, as well as third-party SAN, NAS, and HCI systems.

- **DR for Your Cloud-Based Workloads**
  Automatically fail over from an SDDC in one AZ to another with DRaaS Connect for VMware Cloud. Snapshots of running VMs in the active AZ are stored in the DRaaS repository on AWS S3. In the event of a disaster, these snapshots can be instantly restarted on ESX hosts in a different AZ – all based on well-defined runbook policies.
- **DR for Non-Datrium On-Premises Infrastructure**
  DRaaS Connect for vSphere On Prem extends Datrium DRaaS to any vSphere on-premises infrastructure and provides efficient replication of local vSphere snapshots into DRaaS' deduplicated, compressed, encrypted snapshot storage on Amazon S3. It's managed by a DRaaS cloud-based control plane to define VM protection groups and their frequency, replication, and retention policies. On failback, DRaaS will return only changed blocks back to vSphere and the local on-premises infrastructure through DRaaS Connect.

10x Better Economics with On-Demand Data Center in the Cloud

Eliminate the cost of a second data center and use the public cloud on demand for your DR needs. In steady-state, pay only for backups stored on low-cost S3. Backups are further optimized on the wire and in S3 with data deduplication and compression, which reduces egress bandwidth charges, so the only ongoing cost is for storing backups on S3. DRaaS enables failover from these VM backups in S3 to an on-demand VMware Cloud SDDC using DR plans (or runbooks) managed by ControlShift. This functionality lets you store backups in S3 at low cost and pay for the DR sites only when you need to test or run DR plans.

Fail over to on-demand VMware Cloud on AWS SDDC targets using ControlShift DR plans and store forever-incremental, native backups in S3. DR sites are ready when needed, in the cloud, whether for testing or actual failover. Pay for SDDC only when you need it while your data is ready to use on AWS S3. Efficiently fail back with minimal AWS egress charges by transferring only the changed and globally deduplicated data. Select from a broad range of DR sites around the world.

DR with Instant RTO for 1000s of VMs from Immutable Backups

You don't need to rehydrate data from backups before recovering your workloads. Even though backups are stored long term on low-cost AWS S3, you get instant RTO for DR. All the backups in S3 are instantly made executable on a live cloud-native NFS datastore mounted by ESX hosts with caching on NVMe flash. Instant recovery from older backups is especially useful for ransomware recovery. Enterprises can recover thousands of VMs from the same point in time to restart their entire data center with instant RTO, unlike any other DR solution for VMware Cloud on AWS.

Failproof DR and Efficient Failback

ControlShift automatically checks your plan for health and compliance every 30 minutes, so you can be confident your DR plan is going to work when you need it. Efficiently fail back with minimal AWS egress charges by transferring just data-reduced changes. Similar to failover, failback is fully automated.

Data changes that occur while executing in the cloud are captured and stored as a Cloud Backup snapshot in S3. ControlShift orchestrates the transfer back to the on-premises data center, which includes just the data that changed. Cloud egress charges are minimized by delta transfers, which are further data-reduced.

A single data stack DRaaS solution eliminates the risks associated with multiple data transformations and misconfigurations. Datrium controls protected, backup, and recovery-site endpoints, and it orchestrates all data movements. It also automatically performs end-to-end integrity checks to verify backup fidelity regardless of data location or past replication history. Datrium employs an efficient algorithm to calculate cryptographic hashes of backups (and from DVX primary storage) to continuously validate data integrity across the entire distributed environment, both in the cloud and on premises.

Consistent Operating Environment

Simplify the management of your production and DR sites. DRaaS maintains VMs in their native vSphere format and eliminates the need for brittle and time-consuming VM disk format conversions. DRaaS doesn't use translations, so you can be sure that your VMware workloads will run seamlessly once deployed to your DR SDDC, just like they do in production. Operate both your cloud DR site and production site with vSphere, so you have access to familiar abstractions such as clusters, resource pools, data stores, virtual switches, and port groups following a failover.

While VMware ESX hypervisor dominates on-premises private cloud deployments, public clouds use several other incompatible hypervisors and related administrative tools: AWS relies on Xen and, more recently, on KVM similar to Google Cloud; Azure relies on the Microsoft Azure Hypervisor, a customized version of Microsoft's proprietary Hyper-V hypervisor. Complex vSphere enterprise environments rely on many other virtualization abstractions which have no immediate analogs in the public cloud: clusters, resource pools, data stores, virtual switches, port groups, etc. vSphere also offers a set of widely used services based on these abstractions that have no equivalent in the public cloud: vSphere HA, FT, vMotion, DRS, etc.

VMware Cloud on AWS finally makes the transition between private and public clouds robust by presenting an execution environment in AWS that is similar to the on-premises execution environment – without VM conversion.

As part of creating a DR plan, users map their production virtual infrastructure abstractions (networks, resource pools, folders, data stores, IP addresses, etc.) to the corresponding entities in the destination VMware Cloud for AWS AZ to use for DR. The native on-premises VM geometry is fully preserved, as are all virtual hardware devices. The existing in-guest OS drivers continue to function the same way following migration to the cloud – eliminating all risks of VM conversion between different hypervisor types and the associated virtual hardware and guest OS driver changes.

DRaaS for Any Situation, from Power Failure to Ransomware

DRaaS with VMware Cloud on AWS includes Cloud Backup and ControlShift technologies, providing complementary features for any DR situation.

Cloud Backup	+	ControlShift – Orchestration
Immutable cloud backups in AWS S3, with global deduplication, blanket encryption, and built-in verification		Flexible private and hybrid cloud topologies
Searchable catalog with deep history		Disaster and cybercrime recovery from AWS S3 backups
Recovery of arbitrary sets of VMs, containers, and files		Comprehensive orchestration plans with automated compliance checks
Recovery from ransomware using backups that are months old		Automated runbooks and audit reporting
Snapshot and restore 1000s of VMs with point-in-time consistency		Workload cloud mobility and instantiation
Forever-incremental native replication, to and from the cloud		DR testing in isolated environments

Modern SaaS Experience with Cloud-Native Design

Datrium DRaaS elements are built and managed using modern cloud-native services, including AWS Lambda, Aurora, and S3, which are delivered as SaaS applications in Datrium-managed AWS and VMware Cloud accounts. With automated deployment, configuration, maintenance, upgrades, and failure recovery, you can focus on managing your backups as well as creating and executing your DR plans. There's no manual on-premises software installation.

Simplicity of a Single Data Stack for Recovery Compliance Objectives

ControlShift eliminates the need for parallel hardware and software backups and DR stacks by integrating all components and aspects of the backup and DR into a single system with unified management. A unified cloud orchestration service manages a protected DVX and one or more accompanying DVX systems deployed at another location or in the cloud.

DVX integrates primary and secondary storage, making it possible to use a single management console to establish backup and replication policies and to configure, test, and execute DR plans. Both backup policies and DR plans operate on precisely the same abstractions: backups for VMs and groups of VMs. Because snapshots are at the storage level, ControlShift delivers consistent point-in-time backups across many VMs executing on different servers. This advanced functionality requires native storage integration that is not available from third-party backup software that relies on hypervisor APIs to take snapshots and copy snapshot state into backups.

The system's built-in health checks can pinpoint problems anywhere in the backup and DR stack. For example, replication failures due to network connectivity losses will automatically flag all affected DR plans. ControlShift also automatically performs DR plan compliance checks to ensure that the changes in the execution environment don't invalidate DR plans.

DRaaS offers a 30-minute Recovery Compliance Objective (RCO) with autonomous compliance checks and just-in-time creation of VMware SDDCs. The service also delivers a Recovery Point Objective (RPO) from 5 minutes to multiple years, supporting primary and backup data simultaneously.

Simplified Support

With DRaaS' autonomous deployment, configuration, maintenance, upgrades, and healing from component failures, you can focus on managing your data protection requirements. Datrium monitors and supports all components of the system, including AWS and VMware Cloud SDDC, aided by Datrium partnerships with both companies.

Datrium's staff is fully certified on VMware Cloud and ancillary products, so you get a one-stop support experience. Datrium's unified platform, which includes primary, backup, and DR, dramatically reduces issues. However, if there is an unforeseen issue, the Datrium Support team can quickly address it.

DRaaS Architecture

Datrium DRaaS provides "shared nothing" isolated secure cloud storage and compute environments with no cloud infrastructure sharing between customers. The AWS account hosts backup and DR orchestration services, and the VMware Cloud account hosts cloud DR targets.

All Datrium services are deployed as Amazon Machine Images (AMIs) into a Datrium-created Virtual Private Cloud (VPC) and Subnet. VPC endpoints used to access all other external services required by ControlShift and Cloud Backup are created automatically. All components are monitored and restarted for high availability and resilience. Also, all required state is replicated to ensure resilience.

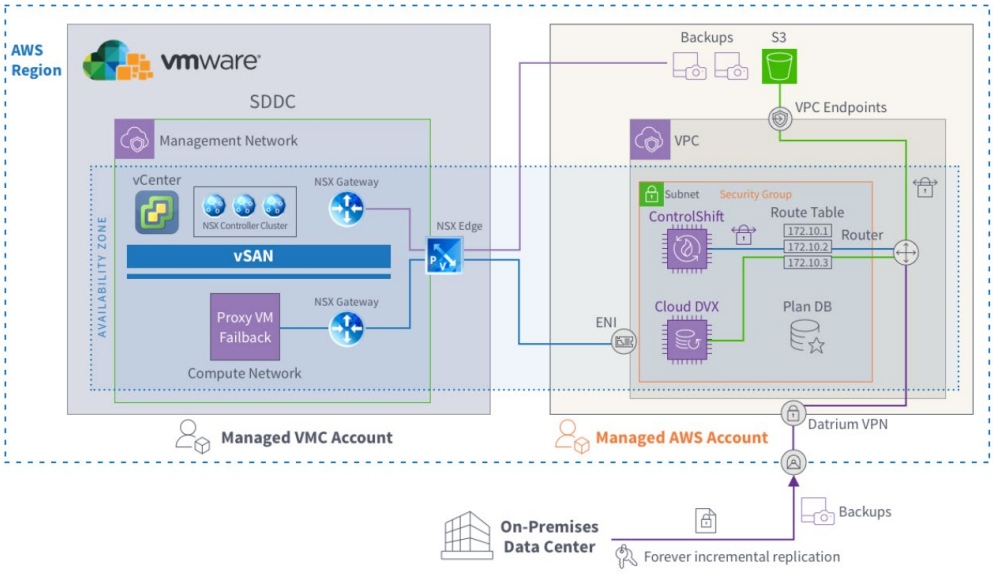

ControlShift (DR Orchestration)

ControlShift (offered as part of the DRaaS service) provides end-to-end orchestration for workload recovery to the cloud or other on-premises sites with comprehensive DR plan definition, workflow execution, testing, compliance checks, and report generation. In this single system context, both backup policies and DR plans operate on precisely the same virtual machine (VM) abstractions, using backups for individual VMs and groups of VMs.

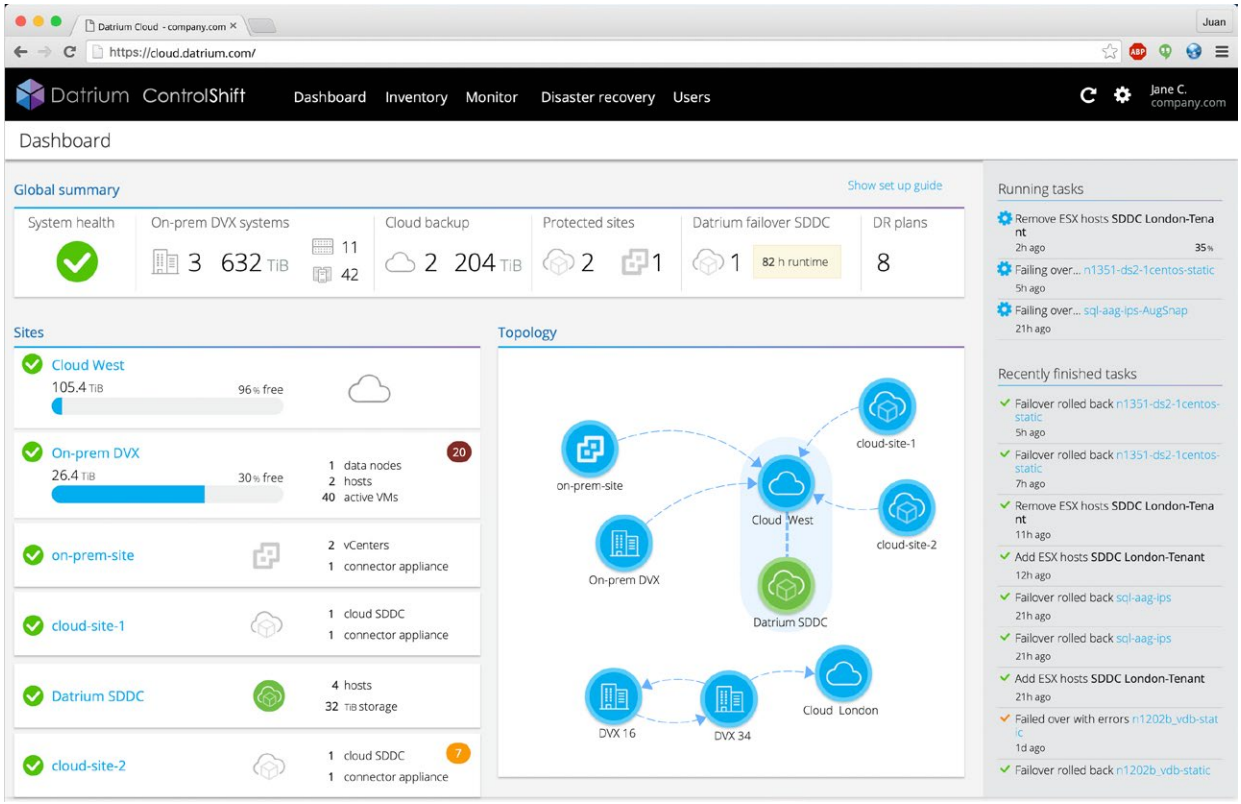

Traditional DR orchestration software products are complex systems composed of dedicated DR orchestration servers and internal databases often augmented with third-party array software agents. These servers and databases are provisioned per site and need to be licensed, secured, monitored, managed, and upgraded – which requires additional maintenance and special operational skills. The initial installation and configuration of DR products often require professional services engagements making the overall solution expensive. DR rollout and upgrade processes are lengthened due to the intricacies of the interactions of multiple cross-vendor products and components.

ControlShift is a pure SaaS solution: there's nothing to install or manage. The ControlShift orchestration engine runs as an AWS-based service and leverages the public cloud infrastructure to achieve high availability for its internal operation. DR plans and execution states are replicated across multiple AZs with automatic failover to a healthy AZ, without any data loss in the event of a disaster affecting the public cloud. Monitoring and upgrades are automated and performed by Datrium as part of the service offering.

When DraaS is activated, ControlShift becomes operational immediately. Users can focus on designing and testing their DR plans instead of having to manage the internal complexities of the DR orchestration software. ControlShift includes all necessary network connectivity and encryption software, and it establishes a secure bidirectional channel between protected sites and the orchestration engine – eliminating the need for an external VPN.

VMware Cloud SDDC Deployment Modes

Just-in-Time Deployment

Replacing an on-premises DR site with a virtual site hosted in the public cloud is attractive for many reasons. Still, it doesn't necessarily reduce the total cost of the overall DR solution because of the recurring charges for maintaining a cloud DR site. The DR costs are merely shifted from on-premises capital and operational expenses to the recurring costs of maintaining an always-on cloud DR site. A detailed total cost of ownership (TCO) analysis is needed to ensure that the overall cloud DR solution is priced competitively with the original on-premises DR solution.

DR related activities don't contribute to the company's top-line performance, but they are necessary to mitigate risk (of disasters or ransomware) to the top-line. Optimizing the costs of DR is, therefore, an important TCO consideration. Just-in-time deployment of a cloud DR site presents an attractive alternative to continuously maintaining a warm standby cloud DR site. With just-in-time deployment, the recurring costs of a cloud DR site are eliminated in their entirety until a failover occurs, and cloud resources are provisioned.

The on-demand nature of public clouds allows DRaaS to drastically reduce the operating costs of DR by deploying the bulk of the DR infrastructure programmatically following a DR event. During steady-state operation, DRaaS maintains a minimal, low-cost AWS cloud footprint to accommodate cloud backups with no ongoing charges for the cloud DR site. The backups are sent to the cloud backup site, and after some processing, land in a cost-effective compressed and deduplicated form in an S3 bucket. In the just-in-time deployment mode, a cloud DR site is created only following a disaster. VMware Cloud SDDC, a Cloud DR site with a significantly larger server footprint and associated costs, is only deployed immediately before executing a DR plan.

To make that possible, DRaaS leverages the space and cost efficiencies of Cloud Backup. The protected site replicates VMs or protection groups in their forever-incremental format to Cloud Backup, which in turn stores them in a compressed and deduplicated native format within low-cost S3. During normal operation, the costs of data protection are limited to the costs of the Cloud Backup service and the cost of the S3 media.

Following a DR event, ControlShift deploys a new VMware Cloud SDDC and orchestrates the failover to this SDDC as part of a DR plan execution. This process uses a fast, high-bandwidth network link from VMware Cloud SDDC to AWS S3 to get access to backups. ESX hosts in the VMware Cloud SDDC mount an NFS datastore that contains the required backups on S3. Using an EC2 instance for storage I/O processing and EC2 Instance Flash for caching, Cloud Backup software makes the backups immediately runnable with no rehydration. This capability ensures instant RTO restarts from these backups, whether they are recent or years old. The recurring charges for the Cloud DR site start accumulating only after the SDDC deployment. The just-in-time deployment of SDDC reduces DR TCO by more than an order of magnitude.

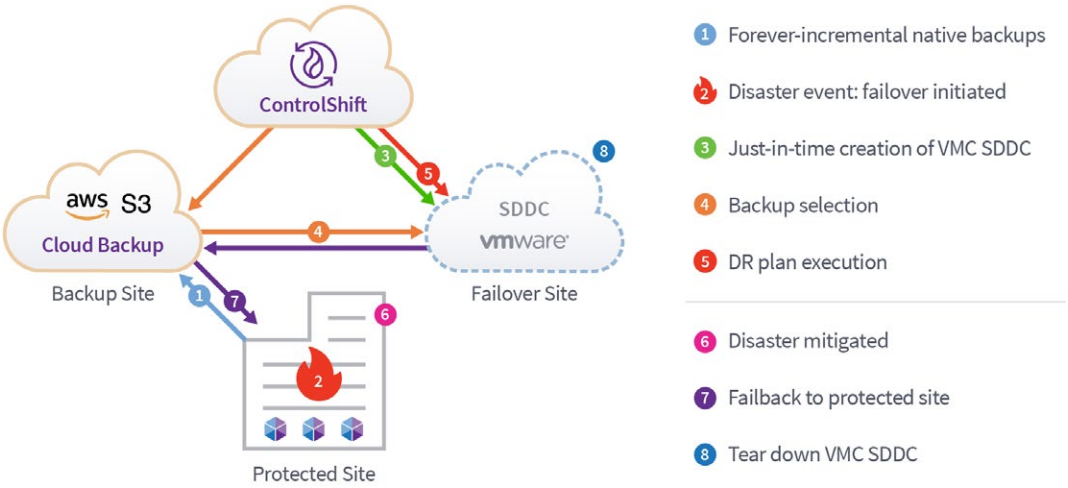

ControlShift supports an efficient, orchestrated failback following an on-premises site recovery. If upon recovery, the on-premises site retains some pre-disaster data, only the data changes that occurred while executing in the Cloud DR site are transferred back to the on-premises protected site.

Ahead-of-time vs. just-in-time provisioning of SDDC is a trade-off between costs and RTO. With ahead-of-time SDDC provisioning, SDDC creation latency could be eliminated. Just-in-time SDDC provisioning dramatically lowers the costs but increases the RTO by deploying SDDC only after a failover.

Ahead-of-Time Deployment

In cases where a DR site has the secondary function of executing non-DR workloads during normal operation, an SDDC can be provisioned before failover.

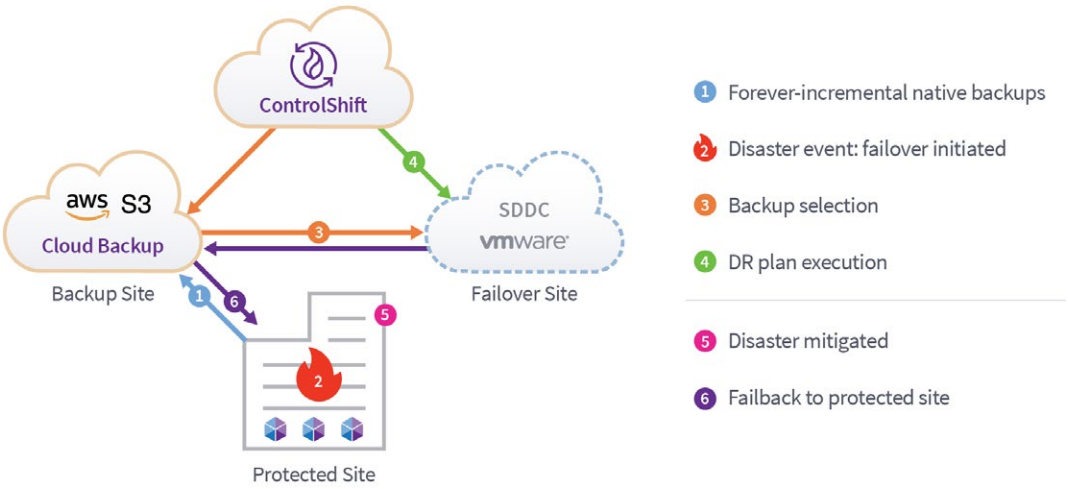

If the sole purpose of the Cloud DR site is to take over workload execution in the event of a disaster and it remains otherwise unutilized, further significant cost savings are possible with the just-in-time deployment.

Pilot Light with Cloud Bursting Deployment

In Pilot Light mode, DRaaS enables a smaller subset of SDDC hosts to be deployed ahead of time for recovering critical applications with lower RTO requirements.

This deployment mode allows organizations to reduce the total cost of cloud infrastructure by keeping a scaled-down version of a fully functional environment always running in warm-standby while ensuring that core applications are readily available when a disaster event is triggered.

With Pilot Light mode, DRaaS presents an option for administrators to add extra SDDC hosts through Cloud Bursting and fail over the remaining applications. Expanding the SDDC by adding hosts happens in minutes, providing a lower RTO for all applications than the just-in-time deployment RTO at a fraction of the cost of the ahead-of-time deployment. A full SDDC deployment is a more time-consuming operation with a higher RTO impact than SDDC expansion. Pilot Light mode is an efficient solution with a range of options to balance costs and RTO.

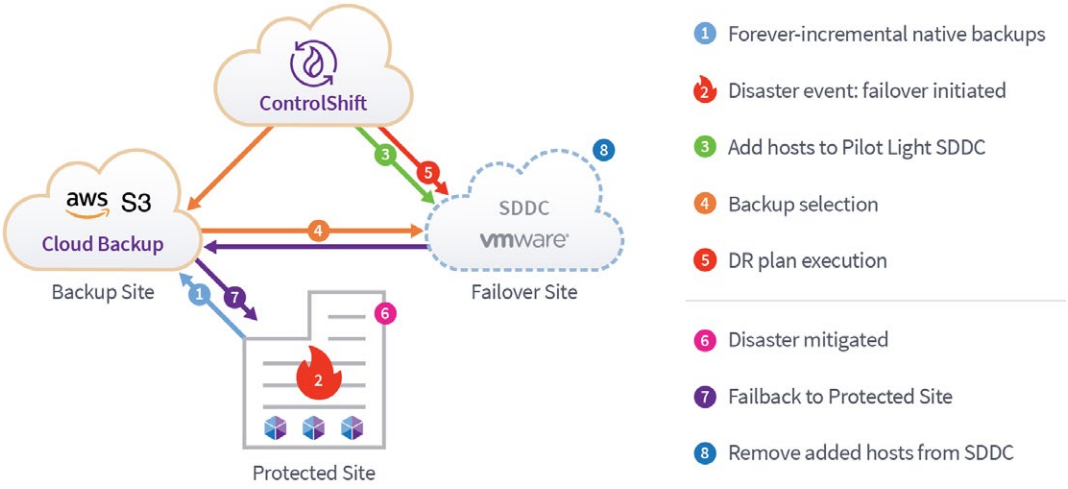

Cloud Backup (Backup Service)

Dedicated tape or disk-based backup and archival systems have traditionally been a key part of enterprise data protection strategies. However, many companies have found the cost and effort of system setup, purchasing, and backup media and software management to be a burden.

Recently, systems built on public cloud object storage services have become an attractive alternative for long-term, off-site data retention use cases. There are three key advantages that public cloud-based systems offer over traditional tape or disk-based backup and archive products:

- Unlike disk or tape systems that need long-term sizing forecasts and upfront CAPEX, public clouds offer pay-per-use pricing and enable scaling to massive levels on demand with zero media management.
- Unlike alternative media options, especially tapes which can take days to make data accessible, data in the public cloud is accessible within a few minutes to hours, greatly accelerating data recovery time. Cloud Backup enables DRaaS and ControlShift to operate in tandem, allowing seamless VMware Cloud SDDC failover and failback.
- On-premises systems send only the incremental changes since the last backup. And these incremental changes are further deduplicated against existing data in Cloud Backup, as explained in the previous section. On-premises systems never send full copies of existing backup data; only the incremental changes are transferred, minimizing the bandwidth and storage usage.

Summary

Datrium DRaaS with VMware Cloud on AWS is a comprehensive cloud-based backup and DR service that protects VMware vSphere environments on premises and in the cloud. It leverages the execution and operational efficiencies of a single integrated data stack to automate and orchestrate all aspects of DR. The solution is much more streamlined and significantly less resource-intensive than legacy DR solutions, resulting in lower RPO and instant RTO for cloud and on-premises environments.

DRaaS delivers a single stack backup and DR service for enterprises without contract management overhead. The solution has a single provider and billing, and enterprises can use VMware Cloud SDDC on AWS for automated and user-defined DR plans with failover and failback from the public cloud.

DRaaS leverages S3 for cost-effective backup, while the integrated data and orchestration stack enables consistency checking of the entire environment, which dramatically reduces errors when a disaster occurs.

With DRaaS, you get a complete solution that delivers comprehensive support, simplified purchasing, and billing, which eliminates the cost and friction caused by multiple point solutions. You get everything you need for failproof, on-demand DR for all of your VMware workloads with instant RTO in one solution.