druva

Next-Gen Data Protection for VMware

Druva and VMware work together to help enterprises experience the benefits of the cloud (speed, scale, and flexibility) wherever they run infrastructure. VMware's hybrid cloud solutions deliver infrastructure and processes familiar to IT, allowing them to accelerate cloud adoption using familiar processes and controls. Druva aligns to this strategy by centralizing data protection across multiple geographies into a single data pool accessed from a unified management console. This drives significant cost and organizational efficiencies allowing IT, information security, and development teams to speak a common data protection language, which diminishes the perception that protecting hybrid clouds is complex.

Druva Data Resiliency Cloud

The Druva Data Resiliency Cloud brings the simplicity, scalability, and security of the public cloud to enterprise data protection and management. It enables enterprises with VMware to accelerate their cyber, data, and operational resilience without hardware, software, or associated complexity. This results in cost reductions of up to 50% over traditional on-premises data protection solutions.

Druva is a single solution that deploys in 15 minutes, eliminating the need for implementation services to configure and tune infrastructure with its simple setup. Druva protects global VMware environments and centralizes data into a single pool, accelerating backup and disaster recovery operations as well as advanced functions such as long-term data retention, compliance, eDiscovery, and analytics.

Powered by Amazon Web Services (AWS), Druva harnesses and builds on the native technologies and global reach of AWS with these additional features:

- Global, source deduplication reduces cloud-storage costs by ensuring only unique data blocks are stored across multiple SDDCs
- Automated storage tiering reduces long-term retention costs by moving unique data blocks to AWS S3 Glacier Deep Archive
- Multi-layered protection for cyber resilience including air-gapped backup data, built-in multi-factor authentication (MFA), and Druva Rollback Actions to recover deleted backup data

Customers see Druva as an attractive alternative to other backup solutions for the following reasons:

- Legacy data backup solutions are unable to scale with VMware Cloud implementations
- Druva aligns with customers' cloud initiatives
- Rising data protection costs and security concerns for both backup infrastructure and data
- Druva automatically scales to support between a few terabytes and up to tens of petabytes, eliminating the additional costs and overhead of legacy/hardware or appliance-based solutions
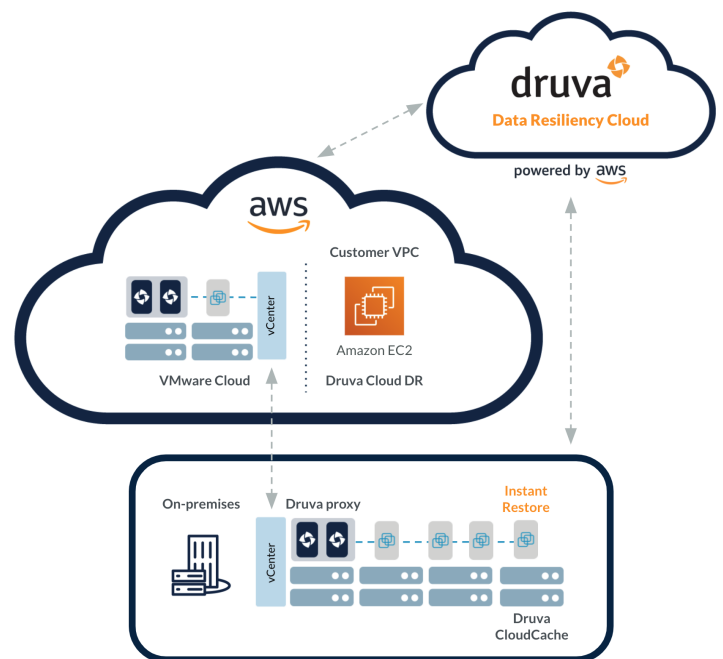
Use cases

vSphere and VMware Cloud backup

Druva integrates with VMware through an agentless proxy that auto-discovers virtual machines and configures change block tracking, including guest processing and data backup. vCenter integration allows customers to use the Druva management console to define virtual machine groups, long-term backup retention rules, and establish recovery time and point objectives (including a local caching option). Customers can also enable change block tracking, application quiesce, and application-aware backups.

VMware cloud disaster recovery

Restore your VMware and VMware Cloud environments in minutes using Druva cloud disaster recovery. Druva enables orchestrated failover/failback of the latest VM snapshots into a customer-owned Amazon Virtual Private Cloud (VPC).

The latest VM snapshots are restored into the VPC from the Druva Cloud and changed into EBS snapshots, aka DR copies. These copies are used to set up the failover and recovery of data within minutes.

When a disaster is declared, EC2 instances are launched from the DR copies and all data volumes are simultaneously restored in approximately 15 minutes. As soon as the VMs are up, your data is available to access. Druva is unique in its ability to automatically failback EC2 instances into any production vCenter environment.

Key features

Data security, ransomware protection, and privacy

- Operational security of cloud backup platform with 24/7 fully managed security operations
- Air-gapped backups are automatically protected in-flight using Transport Layer Security (TLS) protocol, and at-rest with AES-256 bit envelope encryption (customer controlled keys)
- AWS PrivateLink secures data transfers between the Druva Data Resiliency Cloud and customer AWS VPC

Fast and simple backup and restore

- Global, source deduplication and change block tracking accelerate backups and reduce bandwidth and storage consumption
- Agentless and automated protection based on vSphere metadata
- Application-aware backup and recovery
- Instant restore of VMs from Druva CloudCache without hardware lock-in

Cloud disaster recovery for VMware

- One-click failover/failback of on-premises and VMware Cloud environments with RTOs less than 30 minutes
- Automated runbook orchestration, pre-flight validations, and unlimited testing
- Clone DR backups across AWS regions or into different accounts

Benefits

- Simple and scalable backup and DR with ransomware protection
- Flexible restores and workload mobility across on-premises and cloud (15+ AWS regions)
- Integrated long-term retention and optimized storage for a 50% lower TCO
- Business continuity at a fraction of the cost of replication solutions