druva

Druva for Hybrid Workloads

As businesses are adopting a "cloud-first" strategy, data protection is one of the first IT functions to migrate to the cloud. Traditional data protection, whether on-premises or extended to the cloud, is no longer adequate and only increases the security, storage, and management challenges that customers face, as seen most acutely by recent security breaches and ransomware attacks that led to data loss. With Druva, businesses can improve data protection and management of their data center and remote office workloads with a unique cloud-native approach.

Delivered as-a-service, Druva combines high-performance, scalable all-in-one backup, disaster recovery (DR), archival, and analytics to simplify data protection, dramatically reduce costs, and improve data visibility for today's complex information environments, all while enabling organizations to achieve their most aggressive business continuity SLAs (RTO and RPO). With no hardware, no software, and no complexity, businesses can get started within minutes.
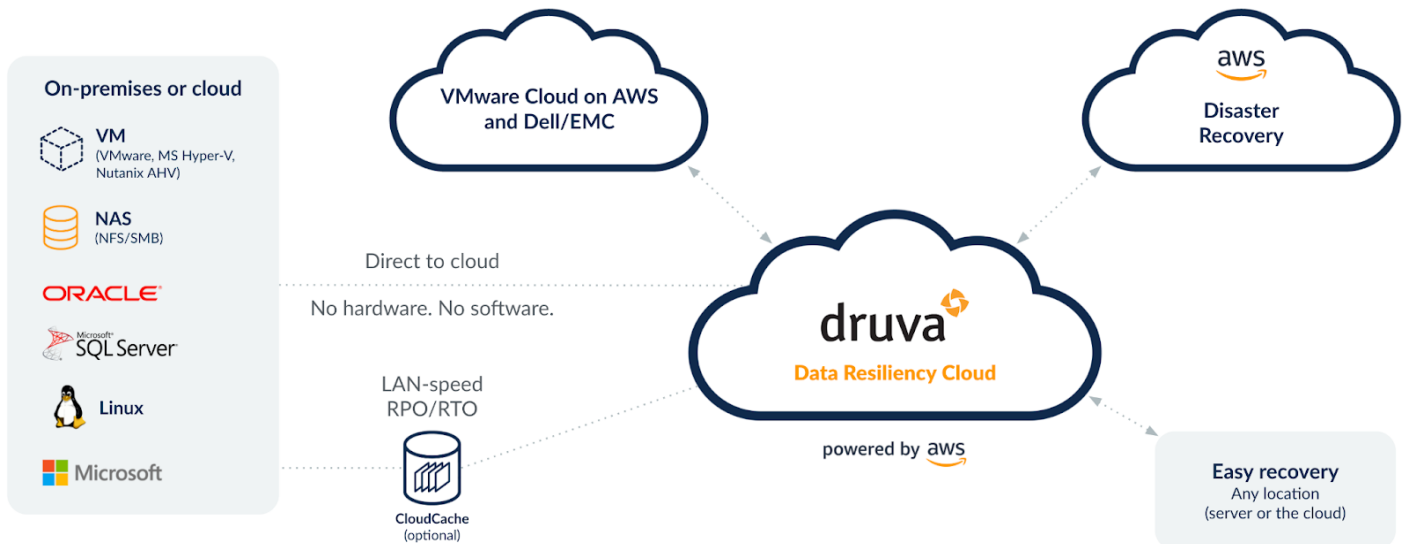
∞
Unlimited backup and restore scalability

50%
Lower TCO

0%
Infrastructure

Why Druva for Hybrid Workloads?

Unified backup, archival, and DR in the cloud

By leveraging the elasticity and scale-out capabilities provided by the cloud, Druva enables organizations to centralize the data protection and management of hybrid workloads that can run in the data center or in any public cloud (e.g., AWS, Azure, GCP). Workload coverage includes physical servers, virtual servers (VMware, Microsoft, Nutanix), hybrid cloud infrastructure (VMware Cloud), any NAS or Cloud file storage, and databases (MS-SQL, Oracle). From a single control plane, IT can easily recover server data down to the file level, failover virtual machines (VMs) for DR with an RPO of one hour and RTO of minutes, easily archive and manage data in the cloud

for compliance, and spin up instances cross-region and account for test and dev (workload mobility) purposes.

Up to 50% lower TCO

With Druva, organizations significantly lower their TCO over traditional or competitive solutions. Being 100% cloud-native, Druva requires no additional hardware or software for data protection, employs an auto-tiering model for cost-efficient storage, provides global scale-out deduplication — reducing bandwidth usage by up to 80% with the smallest storage footprint, and offers customers a true consumption-based model that eliminates wasted resources. Furthermore, there are no restore (egress) charges and customers only pay for data stored, post-global deduplication, in the cloud.

100% Software-as-a-Service, built on AWS

Built from advanced cloud technologies and microservices in Amazon Web Services (AWS), Druva harnesses the native efficiencies and global reach of the public cloud while delivering unmatched storage flexibility, scalability, data durability, and security. Druva delivers new features twice a month, with automated deployment for all customers.

Meeting application data recovery speeds

Druva delivers against stringent RTO and RPO requirements of critical applications while eliminating lock-in, complexity, and the cost of legacy on-premises hardware infrastructure. With one-click disaster recovery and automated runbook execution, Druva eliminates recovery complexities and enables failover within minutes, as well as failback. Customers can easily order an AWS Snowball Edge appliance on-demand, which comes preconfigured with Druva software to meet business continuity SLAs while still taking advantage of cloud-scale, resiliency, and simplicity. Alternatively, IT can also install Druva CloudCache on any commodity hardware of their choice and achieve VM restore speeds up to 600 GB/hr.

Industry-leading data security and privacy

Druva's approach to storing enterprise data utilizes both an advanced data-scrambling algorithm and a unique envelope-based encryption model where the data and metadata are decoupled and encrypted. This guarantees that your data is only accessible by you — a critical component to meeting today's stringent global data privacy regulations. Under no circumstances can Druva access your data. Furthermore, to ensure complete protection against ransomware, Druva provides data isolation from infrastructure attacks, and high-performance restores to minimize downtime of compromised systems.

Key features

Data backup and recovery

- Incremental-forever backup model
- Unlimited full restore points for quick recovery (no restore egress charges)
- Global, source-side, inline deduplication (Petabyte scale)
- High-speed, LAN-based backup and restore with CloudCache for tight RTO/RPO needs
- Search across snapshots for granular and high-performance recovery

Administration

- Cloud-based centralized management
- Unified interface for hot, warm, and cold backups
- Role-Based Access Control (RBAC)
- Actionable and predictive analytics for storage and backup optimization

Data security and privacy

- Immutable, air-gapped backups
- 256-bit AES and TLS 1.2 encryption
- SOC-2 Type II, HIPAA, Privacy Shield certification
- No key management required

Disaster recovery

- RPO of one hour and RTO of minutes for failover — no AMI conversion
- One-click failover and failback
- Automated recovery orchestration and runbook execution
- DR into customer VPC and clone VPC across regions or accounts

Data archiving

- Lower costs for compliance with both auto-tiering to cold storage and direct backups to Druva Archive Storage (50% less)
- No limitation to the number of aged snapshots
- Shared global dedupe index for lower-cost, long-term storage