druva

Druva CloudCache for Hybrid Workloads

All the benefits of the cloud, without compromising enterprise SLAs

The challenge

Organizations want to embrace cloud data protection for its efficiency, scalability, and flexibility, but need to have local copies of data, whether in the data center or remote offices, to address strict service-level agreements (SLA). Commonly cited roadblocks for adopting cloud-based data protection include the following:

- Limited bandwidth at remote sites
- Large volumes of data to be backed up and restored from the cloud
- Aggressive recovery time objectives (RTO) and/or recovery point objectives (RPO) for mission-critical workloads
- Compliance requirements for backup copies on-premises

As organizations look to modernize their data protection strategy with a SaaS platform, they need to solve for strict RTO and RPO requirements while also removing the restrictions and overhead of traditional backup infrastructure.

Business challenges

- Limited bandwidth
- Large data restores
- Strict RTO/RPO SLAs

Integration benefits

- Instant restore of VMware VMs
- Up to 600GB+/hour restores
- Easily meet enterprise SLAs
- A no charge feature, fully integrated in the Druva Cloud
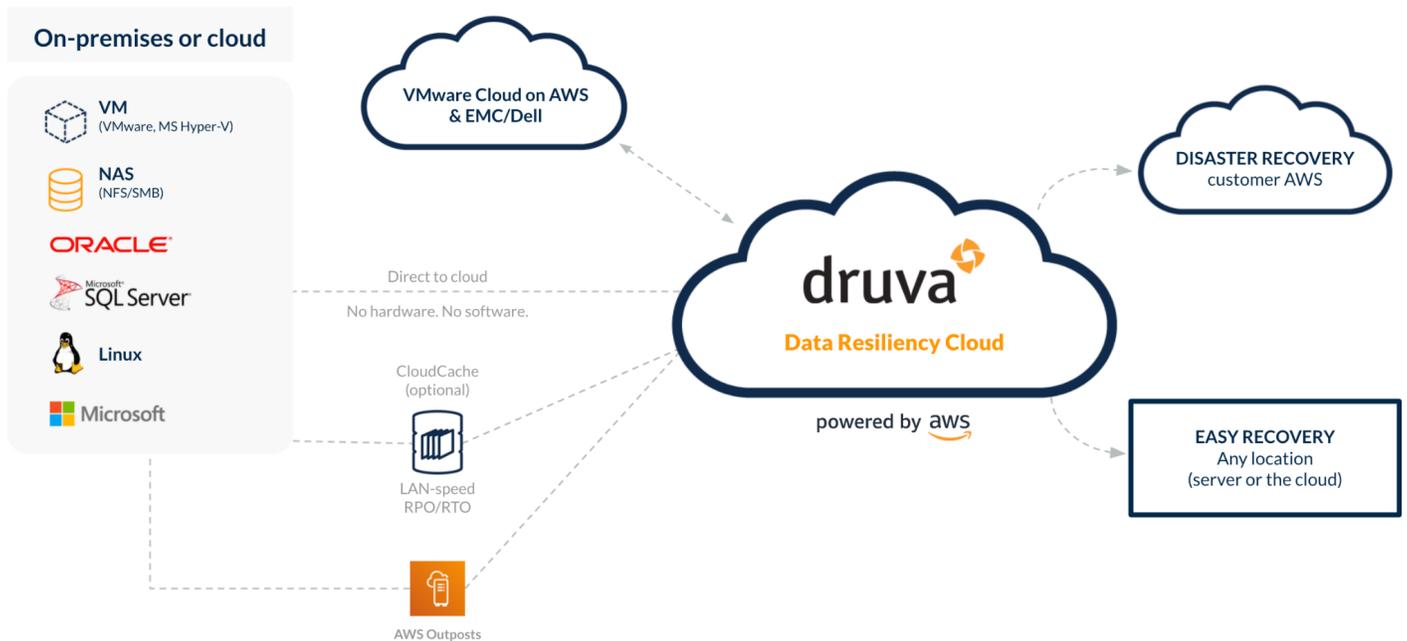
The solution

Druva CloudCache enhances the protection of hybrid workloads both in data centers and remote offices by providing the flexibility to meet strict RTO and RPO requirements regardless of available WAN bandwidth. Druva CloudCache is a software appliance that provides local backup and recovery points. It is fully integrated in the Druva Data Resiliency Cloud and can be deployed in minutes as a VM and scaled up or down as needs change.

druva
Data Resiliency Cloud

Production workloads

CloudCache
SW appliance

Instant restore for VMware

Data Centers or Remote Offices
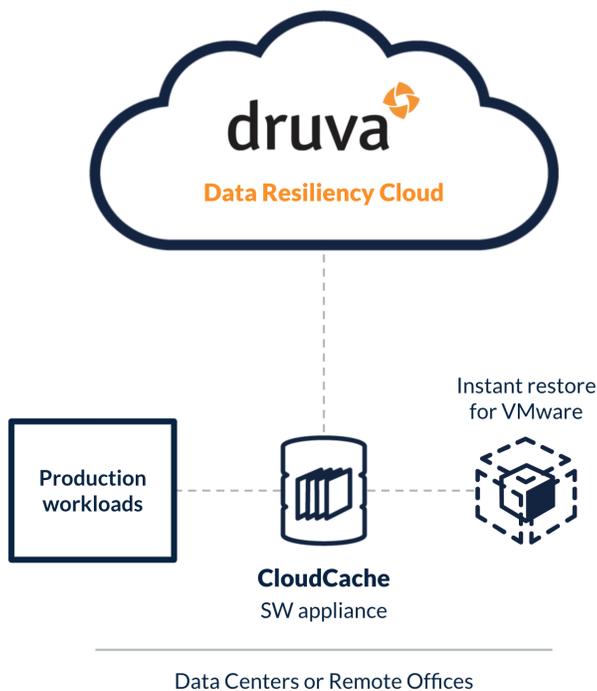
Key features

- Fast backup and recovery regardless of available WAN bandwidth (up to 600GB/hr)
- Global, source-side deduplication reduces LAN and WAN bandwidth consumption
- Bandwidth throttling and scheduled replication to control WAN consumption
- VMware Instant Restore of multiple VMs simultaneously
- Enhanced resiliency, ransomware and accidental deletion protection via a golden copy of all data and metadata in the Druva Cloud
- Encrypted local data storage

Druva CloudCache helps organizations meet tight RTOs with high-speed, LAN-based restores and support for VMware Instant Restore. For demanding RPOs and/or limited WAN bandwidth, customers can back up locally to CloudCache and schedule replication to the Druva Data Resiliency Cloud during off-peak hours.

Key benefits

- **Easy and flexible deployment** - Software appliances on commodity hardware or VMs can be scaled up or down as needed to meet different use cases (e.g. initial deployment vs. ongoing operations).
- **Meet demanding RPOs/RTOs** - Provides strong performance for the data center and remote office without hardware lock-in.
- **Built-in resilience** - Stores golden-copy of metadata and data in the cloud.
- **Simplified management** - Manage CloudCache from within the Druva Cloud console.
- **Optional, no-charge feature** - Pay only for storage consumed in the cloud (not on-premises with CloudCache).
- **Automated long-term retention** - Store an unlimited number of snapshots with local retention up to 30 days and infinite cloud retention. Druva can automatically tier data across storage tiers (from warm to cold) to reduce your storage costs.

How it works

Druva CloudCache deploys on the same local network as protected production systems. Specified policies dictate whether to back up to a CloudCache instance or directly to the cloud. Global, source-side deduplication reduces both LAN and WAN consumption as well as local storage requirements. CloudCache replicates data to the cloud on a scheduled basis while allowing for network-based backups and restores to occur as often as required.

Data resiliency is paramount and enabled by the Druva Cloud Architecture, which delivers an infinitely scalable cloud file system with self-healing, global deduplication, and auto storage tiering. All backup metadata exists in the cloud while CloudCache holds the latest snapshots for up to 30 days.

The Druva CloudCache architecture ensures that data is always available, once synced to the cloud, even with the loss of a CloudCache server. Data transfers from CloudCache to the Druva Cloud can be configured during off-peak hours with bandwidth throttling to ensure optimal network usage. After 30 days, data is automatically purged from the local CloudCache, but remains available in warm storage in the Druva Cloud. For data sets that require longer retention (e.g. 1yr+), customers can enable Druva long-term retention to realize storage savings of 20 percent. With LTR enabled, Druva automatically tiers data from warm to cold storage (AWS S3 Glacier Deep Archive) after a set retention period to optimize storage usage and reduce costs.

Unlike expensive purpose-built hardware appliances, Druva CloudCache is a highly flexible software appliance, available at no additional charge, which can be provisioned using commodity servers and storage in any number of sites. Customers can dedicate as much or as few resources as needed at any time. Embrace the scale, security, and efficiency of cloud-native data protection with the flexibility to meet strict RTO and RPOs without hardware lock-in.

Supported workloads

- Windows, Linux
- VMware vSphere, Microsoft Hyper-V
- Microsoft SQL and Oracle databases[1]
- Network Attached Storage (NAS)

---

[1] Expected 1H 2022 for Oracle direct-to-cloud backup method. Druva also offers Phoenix Backup Store for Oracle databases in an integrated NFS server used as an RMAN backup target location for Oracle databases on-premise or in the cloud (i.e. Oracle on AWS EC2).