Datera Snapshot and Clone Technology

Datera snapshot and clone technologies enable high performance data protection and copy data management with optimal resource utilization

Benefits

- Datera snapshots deliver the highest service levels RPOs/RTOs

Efficient

- An efficient Redirect on write (RoW) implementation reduces production workload impact and data sharing for Copy Data Management (CDM)

Resilient

- Not only is the primary data resilient to hardware component, storage node or site failures, but snapshot-based recovery points are also resilient

Ecosystem Support

- Integration with VMware vCenter plug-in, Kubernetes (CSI), OpenStack and CloudStack

Contents

Overview. 2
Technical Implementation 2
Copy Data Management Use Cases 4
Cloud Backup (Datera2Object) 5
Ecosystem Support. 6
Features and Benefits 7
About Datera 8

Overview

Datera Services Platform (DSP) is a software defined storage platform built with a unique architecture that enables the rapid and flexible adoption of the latest industry-standard server and storage media technologies. The platform autonomously delivers a cloud-like user experience and business model on-premises at enterprise scale, with linear performance, continuous availability, and agility.

Software Defined Storage has matured, and now offers the full range of services traditional storage arrays have offered for two decades. Datera customers enjoy deduplication, compression, encryption and fully integrated snapshot and clone capabilities. These services are even more powerful on a platform with autonomic, intent and policy driven control of workloads and volumes. Datera = SDS flexibility + full enterprise storage services + industry leading automation.

Storage integrated snapshots and clones have become standard features in many enterprise storage products. Snapshots can be the foundation of a data protection strategy, providing an instant point-in-time copy of data that can be used for recovery in the case of incidents like operator errors, application data corruption, or malware attacks. Snapshots are used with application data where recovery service levels are more stringent than what can be met with traditional disk or tape-based backup solutions. Clones are similar to snapshots in that they can be created instantaneously and represent a virtual copy of data. Whereas data in a snapshot is immutable, data in a clone can be changed and can be used for operations like data analytics or test/dev. Datera relies on a unique combination of technologies to provide snapshots and clones that are high performance, efficient, resilient and provide flexible management.

Technical Implementation

The DSP architecture is designed for scale-out and efficiently distributes I/O across storage nodes for resiliency and performance. The system uses a multi-replica copy technique to ensure maximum data resiliency. The system supports up to five replicas of data which are distributed across storage nodes, racks and metro geographies based on policy configuration for any given application or volume. A patented lockless distributed-coherence protocol is utilized to deliver low latency performance in the largest scale-out storage environments. The system continuously optimizes data placement based on available resources and application needs.

The node local block Store is implemented based on augmented B+ trees for its meta-data. This augmentation allows sharing of data blocks between B+ trees, but also sharing of sub-structure. A snapshot merely results in the B+ tree for a volume, whose snapshot is being taken, becoming a shared B+ tree with the volume and the snapshot sharing it. All the data and metadata associated with a snapshot is replicated across nodes like the source data to ensure the same level of resiliency. This replication all happens quickly and efficiently due to the use of timestamp-based lockless coherency algorithm which also guarantees that all nodes take snapshots consistently.

When the volume receives writes/overwrites after the snapshot is created, the B+ trees are partly unshared as necessary, by dynamic copy-on-write (CoW) of the meta-data, while the data itself is redirected-on-write (RoW). This is in order to ensure that the snapshot keeps references to the data blocks in the volume at the time of the snapshot while the volume can reference new data blocks.

This diagram illustrates whose data is represented to the client as well as snapshot and backend storage.

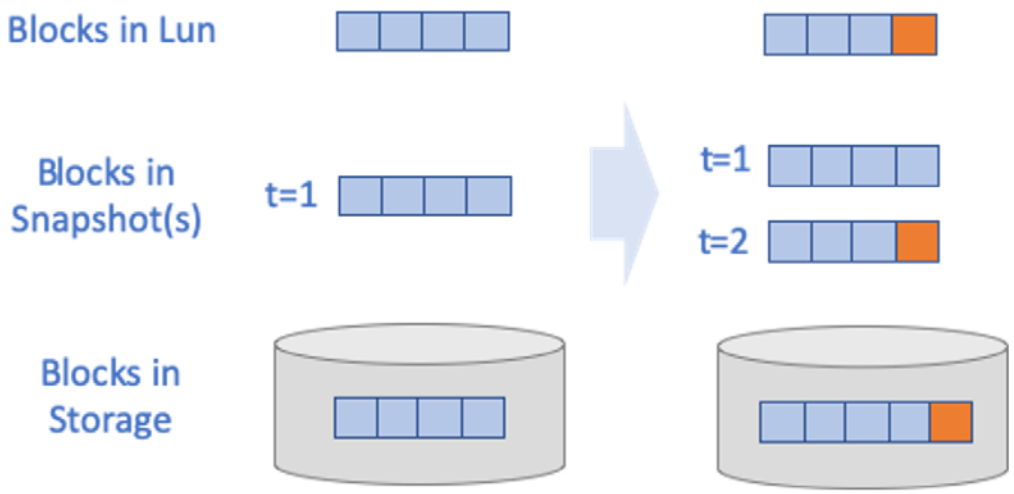

Redirect on write (RoW) implementation reduces storage system resource impact by minimizing the number of writes required when using snapshots. New data written after a snapshot is created is written to a new block. Previously written data in a snapshot does not have to be read and rewritten to storage like some implementations.

Over time the meta-data for the volume diverges from the meta-data for the snapshot even though at the time of the snapshot they were not just identical, but actually the same meta-data (same augmented B+ tree).

The technical implementation of clones are very similar to snapshots, using the same B+ trees and space efficient pointer mechanisms. The major difference is that clones are exported as new read/write volumes. Hence, new user writes can arrive for the original volume or the clone, unlike snapshots which are read only. Writes to a clone result in the meta-data for the original volume and the clone to diverge. Another difference is that snapshots belong to the same volume and Application Instance as the source, whereas clones are created as a different volume in a new Application Instance and Storage Instance.

One important advantage of Datera clones over most other implementations is that the clones are completely independent of the source volume, even though they initially share all the same data. Clones can also be made of clones with this same independence. This allows source volumes to be deleted without consideration of clone usage. Another advantage is that there is flexibility to change the media or placement policy for the clone for automated migration of data to improve performance or minimize cost.

Data Protection Use Cases

Snapshots are an essential component of a complete data protection solution for mission critical applications that have the most demanding service level requirements. This is especially true with large data sets that would take a long time to backup with traditional methods that move a lot of data and have a negative impact on system performance. Taking snapshots does not require data to be moved. A set of pointers is added to the B+ tree structure and this happens in less than a second even on the largest multiple terabyte volumes. These operations have minimal impact on the system so these snapshot recovery points can be created much more frequently than a traditional host file system based backup. More frequent snapshots provide a lower Recovery Point Objective (RPO) which means that the potential for data loss is much less. The other advantage of snapshots is that the Recovery Time Objective (RTO) process is also much lower due to the pointer based implementation, so that application time down can also be minimized.

Snapshots are generally not considered "backups," since the physical snapshot data is shared with the source data resulting in a single point of failure. However, Datera snapshots can be thought of a little differently. First, although data is shared between the snapshot and source volume, both the source and snapshot data is replicated across multiple storage nodes to avoid there being a single point of failure. With traditional storage arrays, a failed system would result in loss of both the source data and the associated snapshot recovery points, making recovery impossible. In the case of a Datera, single storage node failure would not result in data failure. In fact, a common Datera configuration with a replication factor of three can tolerate up to two node failures with no data loss.

Many applications also require off-site backup copies for resiliency or compliance purposes. This is an area where Datera stretched cluster solution can have the additional benefit of site resiliency. With a stretched cluster both the source data and all snapshot recovery points are synchronously replicated and stored across two physical sites to avoid a single point of failure.

The best data protection solutions combine the use of snapshots for low RPO/RTO recovery points with a host file system based backup for application consistency and long-term data retention. Datera is the perfect complement to a backup application whether the workload is file system or database on a bare metal server, VMware, Kubernetes or other cloud application.

Copy Data Management Use Cases

Industry data indicates that for some applications such as SAP there are 20 non-production copies of data for every production copy of data. The idea of Copy Data Management (CDM) involves taking a holistic approach to re-using copies of data (virtual copies) for multiple use cases such as:

- Disaster recovery
- DR test
- Operational recovery (backup)
- Data analytics
- Test/dev

CDM leverages Datera embedded technologies like snapshots and clones The user points to source data and specifies “virtual copies” of data. DSP supports clones that can be used to quickly make space efficient copies of data for test/dev, analytics and other copy data management workflows. For example, a single remote copy of data can be used for site disaster recovery, DR test and even operational recovery if the recovery points (snapshots) are retained long enough. With read/wrote clones the same copy of data can be used for test/dev or DR test – taking advantage of host application servers that are idle waiting for DR.

The fact that Datera clones are independent of the source volume and can be easily migrated to different media provides the ultimate flexibility and can enable a broad set of new and interesting use cases like analytics or archive.

Cloud Backup (Datera2Object)

Datera2Object replication is a solution for backing up and restoring Datera primary snapshot data to and from a remote public cloud or on-prem object store. The solution uses replication to create copies of Application Instances and volumes in remote object storage solutions such as Datera system with S3 object services for backup. Backups can also be replicated to Amazon Web Services (AWS) S3, the Google Cloud Platform (GCP) and generic S3 Object Storage systems.

Datera2Object operates by taking a primary storage snapshot and replicating its data to an object store. The remote snapshot objects can be retained long-term as an off-site copy. After the initial, baseline, data transfer and subsequent updates are incremental, block level transfers. The replication takes place directly between the Datera source system and an object storage system, making the process very efficient with low resource impact. The replication takes place directly between the Datera source system and an object storage system, making the process very efficient with low resource impact.

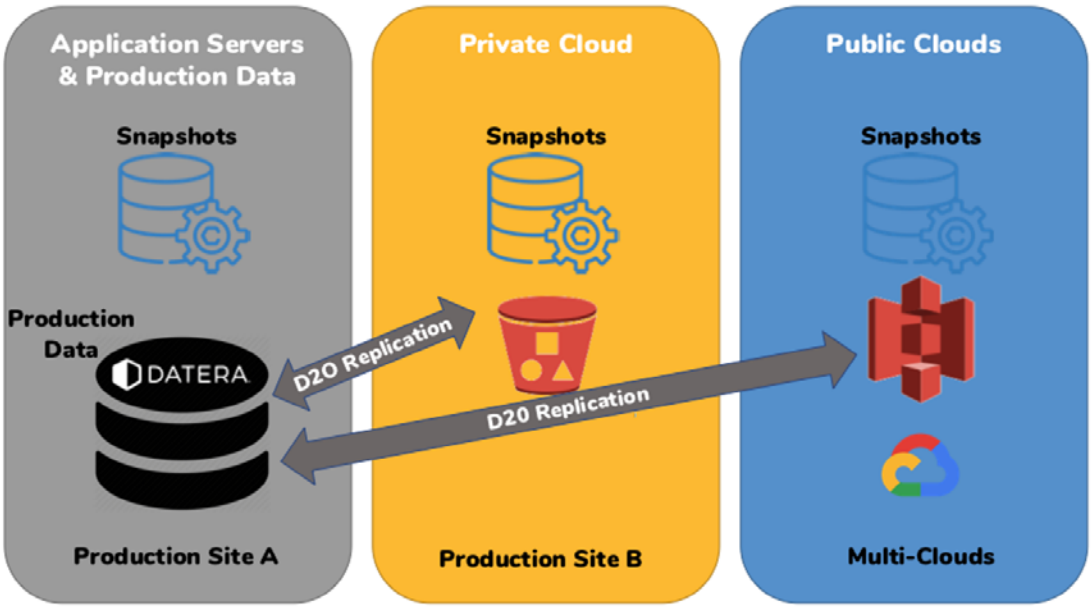

The restore process is similarly efficient with the ability to directly restore back to the source storage system with only incremental data transfer. Data can also be restored to an alternate Datera system for operational recovery or remote disaster recovery.

Datera2Object can be used with Datera local snapshots to provide a comprehensive data protection solution. This combination provides low RPO/RTO snapshots on the source system for instantaneous short-term recovery as well as efficient long-term backup and recovery with a remote public or private cloud system.

Datera2Object supports these primary use cases:

- Creating space efficient backups of Datera Application Instances and volumes in a remote object store
- Restoring Datera Application Instances and volumes from a remote object store to either the source system or and alternate Datera system
- Migrating Datera Application Instances and volumes to remote object store with the ability to migrate back to the same or another Datera system

Ecosystem Support

Traditionally, all storage provisioning and management operations happen at the storage system level, with no visibility for the VI admin. The Datera plugin for vCenter provides the ability to manage its storage directly from vCenter. Using the Datera plug-in for vCenter, new volumes can be created for use in a cluster from end-to-end with a simple wizard. The vCenter GUI allows the configuration of storage snapshot schedules when provisioning new VMware datastores. On-demand snapshots can also be triggered from the GUI, and supports the automation of recovery workflows from Datera storage snapshots. The snapshot creation and restore processes are near instantaneous to provide VI administrators with peace of mind.

Datera plugin for vCenter also enables snapshot data to be cloned and mounted to other ESX hosts to enable granular recovery or support test/dev workflows. Again, the plug-in exposes the power of storage clones to VI admins with a friendly user interface to automate workflows.

Datera vCenter Plug-in for VMware Snapshot List

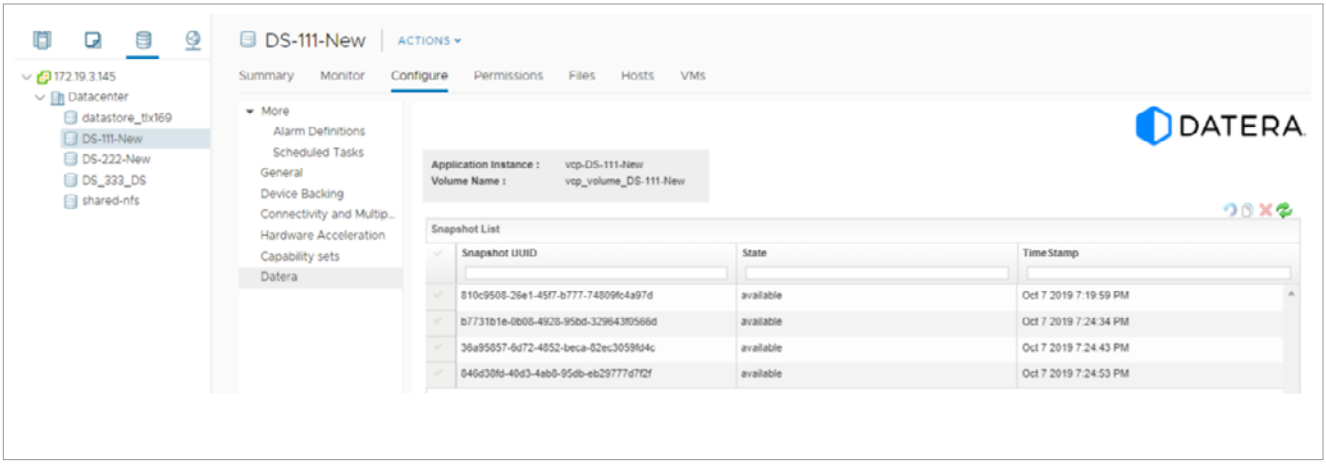

The Datera CloudStack storage plug-in provides Datera storage plugin support for CloudStack Primary storage. The driver is based on iscsiLUN and currently it supports Xen Hypervisors.

The Container Storage Interface (CSI) is a standard for exposing arbitrary block and file storage systems to containerized workloads on Container Orchestration Systems (COs) like Kubernetes. The Datera CSI plug-in can be deployed to expose new storage systems in Kubernetes without ever having to touch the core Kubernetes code. Datera's CSI driver deeply integrates with the K8s runtime. It allows deploying entire stateful multi-site K8s clusters with a single K8s command, and pushing application-specific telemetry to the Datera policy engine, so that it can intelligently adapt the data fabric. Datera's powerful storage classes, and policy driven workloads are a natural fit with Kubernetes and integrates with native snapshots for creating recovery points and doing instantaneous restores.

Features and Benefits

High Performance

A pointer based mechanism enables snapshots to be quickly created in seconds without moving data. This provides very low Recovery Point Objectives (RPO) to meet the most demanding application data protection requirements. The same pointer based mechanism also enables instantaneous restores from snapshots.

Efficiency

An efficient Redirect on write (RoW) implementation reduces production workload impact by minimizing the number of writes required with snapshots. Snapshot data also leverages deduplication and compression to maximize storage space utilization. Similarly, data is restored from snapshots almost instantly, again with minimal resource impact on the system.

Resiliency

Datera offers incredible data durability by combining powerful data services, such as snapshots, with a policy-based, application-driven, shared nothing architecture. The combination of integrated snapshots, clones, replicas, stretch clusters and synchronous replication enable you to build massive resilience and efficiency into your private, hybrid or public cloud composable software defined data center. Not only is the primary data resilient to hardware component, storage node or site failures, but snapshot-based recovery points are also resilient.

Flexibility

Snapshots can be scheduled or invoked on-demand via REST API or a web-based GUI. There is the ability to take a snapshot of an individual volume or at the Application Instance level to provide a multi-volume consistency group. Application templates can be used for large scale, high velocity deployment of snapshot schedules.

Ecosystem Support

Datera also provides ecosystem drivers for Kubernetes (CSI), OpenStack and CloudStack to provide cloud orchestration of storage snapshots and clones. Datera also provides a VMware vCenter plug-in that enables virtualization administrators to configure schedules and access storage snapshot and space efficient cloning capabilities from a friendly GUI interface.

About Datera

Datera is the global leader in autonomous data services software for the open cloud, delivering server-based software defined storage solutions. Enterprise companies around the globe use the Datera Data Services Platform for self-driving data operations at hyper-scale, turning commodity servers into high performance, scale-out storage, and radically lowering infrastructure costs.

https://datera.io/