Dell APEX Block Storage for AWS

Business and Technical Case

February 2024

H19726

White Paper

Abstract

This document describes how to use APEX Block Storage for AWS to bring your enterprise storage experience to the cloud to meet the most demanding needs of your mission critical applications.

DELL Technologies

Copyright

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2023-2024 Dell Inc. or its subsidiaries. Published in the USA February 2024 H19726.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Contents

Executive summary....4

Business Case - Dell APEX Block Storage for AWS....7

Technical Case - Dell APEX Block Storage AWS-specific architecture 17

References....40

Executive summary

Overview

APEX Block Storage for AWS is a deployment of the Dell software-defined block storage innovation in the public cloud. This offer allows customers to experience the same benefits of enterprise-class storage services in the cloud as they have come to appreciate on-premises. Dell APEX Block Storage for AWS uses a scale-out distributed mesh-mirror architecture which provides higher performance, larger volume sizes and improved resilience beyond what is currently available in the public cloud. In fact, it is the most scalable, flexible, and reliable cloud *"software first"* block storage on the market today[a].

Every business has unique needs and requirements with different dimensions that meet those needs. This paper will focus on both the business use case and the technical use case. Analysis from all perspectives, both business and technical, are critical for a successful implementation.

Audience

'C'-level to platform operations administrators, along with DevOps and cloud integration admins can benefit from the content within this document. Any technical decision maker with a multicloud focus will find the use cases compelling for Dell APEX Block Storage for AWS.

The premise of this paper is to define both the business and technical outcomes of a successful Dell APEX Block Storage for AWS implementation. This paper is not a deployment guide but does offer a fully vetted document to provide successful research and planning of the required resources.

Revisions

Date	Part number/ revision	Description
September 2023	H19726	Initial release
March 2023	H19726	Added SDC operating system requirements matrix
February 2024	H1926	Added the Operating system requirements

We value your feedback

Dell Technologies and the authors of this document welcome your feedback on this document. Contact the Dell Technologies team by email.

Author: Robert F. Sonders

Contributors: Anthony Foster, Roy Laverty, Brian Dean, Michael Richtberg, Simon Stevens

Note: For links to other documentation for this topic, see the APEX Info Hub.

[a] Based on Dell analysis of storage software deployable on AWS, Azure, and Google Cloud, May 2023.

Terminology

The following terms and abbreviations are used throughout this document.

Table 1. Terminology used in this document

Acronym	Description
AMI	Amazon Machine Image. A supported and maintained image provided by AWS that contains the information required to launch an instance
AWS	Amazon Web Services
AZ	Availability Zone. A discrete data center with redundant power, networking, and connectivity in an AWS region or Azure region
Azure	Microsoft cloud
CIDR	Classless Inter-Domain Routing
CSI	Container Storage Interface
CSM	Container Storage Module
DAS	Direct Attached Storage. (Both the instance store and the EBS volumes are treated as DAS.)
DDVE	Dell Data Domain Virtual Edition
Dell SDS	Dell Software-defined Storage
EBS	Elastic Block Storage volume of a gp3 type. Software-defined storage volumes from AWS
EC2	AWS - Elastic Compute Cloud
Fault set	A logical entity that contains a group of SDSs within a Protection Domain
Hyperscaler	Any large public cloud service provider that can provide services such as compute and storage
IAM	Identity and Access Management
Instance - AWS	EC2 compute instance. A basic building block of AWS
Instance - Azure	Series type virtual machine – predefined sizes
LIA	Lightweight Installation Agent. A required APEX Block Storage component for many maintenance operations
MDM	Metadata Manager. An APEX Block Storage component that keeps everything running behind the scenes

Acronym	Description
On-prem or on-premises	A locally deployed data center set of hardware components
PFMP	APEX Block Storage Manager Platform
PowerFlex Gateway	An APEX Block Storage component that orchestrates deployments and upgrades and acts as the REST API gateway to the APEX Block Storage cluster
PowerFlex GUI	The web user interface used for configuring and managing APEX Block Storage components, often called the Element Manager
Proxy server	A public-facing server used for the "jump server" method of accessing the APEX Block Storage deployment
PV	Persistent Volume (Kubernetes construct)
PVC	Persistent Volume Claim (Kubernetes construct)
RCG	Replication Consistency Group
Region	A collection of cloud data centers within a geographic region, such as North Virginia (us-east-1)
RPO	Recovery Point Objective
RPQ	Request for Product Qualification
RTO	Recovery Time Objective
SDC	Storage Data Client. The Dell APEX Block Storage component installed on a server that requires access to APEX Block Storage volumes
SDR	Storage Data Replicator
SDS	Storage Data Server. The Dell APEX Block Storage component that controls the local disks
SLA	Service Level Agreement
STaaS	Storage as a Service
VPC	Virtual Private Cloud. Network setup without bonding. A logically isolated section of services

Business Case - Dell APEX Block Storage for AWS

For decades, the enterprise infrastructure industry has been optimizing ways for architectural ingredients to interoperate. Storage that meets the needs of discriminating enterprise requirements is a core competency of the Dell Technologies portfolio. Taking advantage of these capabilities, while embracing the agility afforded by cloud infrastructure, brings together the best of both worlds.

As a market leading supplier of software-defined storage, we enable outcomes on the public cloud that take scale, performance, and durability to an outstanding level. This document describes how to bring your enterprise storage experience to the cloud to meet the most demanding needs of your mission critical applications.

To truly embrace multicloud, there is also a fundamental need for STaaS based management and orchestration. The Dell APEX Navigator experience will solve for and manage day zero deployments and day two management of all Dell software-defined cloud storage and data protection endpoints. Optionally, support for Kubernetes storage management can also be delivered as part of this solution.

Native hyperscaler solution offerings can achieve three to four nines of availability. These solutions rely on zone-redundant replication to replicate services around to protect solutions from single points of failure. With Dell APEX Block Storage, we use the native offering instance but take the resiliency many steps forward. The offering can span multiple availability zones with a single cluster. The offering also deploys its own highly available metadata management, as part of the mesh-mirror architecture. This vastly improves resiliency and reliability over and above native hyperscaler offerings.

To make a solid business case for Dell APEX Block storage, we must provide proven on-premises performance, scale, and resiliency, coupled with seamless data mobility and monitoring. These enterprise features can be deployed easily, and with very little friction, to multiple hyperscalers. As a result, we can now place the right data, in the right location, at the right time, for the right consumption service.

The scale out software-defined architecture enables extreme performance by aggregating storage resources and distributing the data access across the cluster to deliver large IOPS with low latency. Aggregating the resources with the solution’s SDS results in performance that is the sum of the component pieces' performance.

One can rapidly expand the cluster to increase both performance and capacity while maintaining low latency. Each additional instance adds additional I/O and capacity to the cluster, with up to 512 nodes in a single deployment. This enables extreme performance and capacity.

You can deploy Dell APEX Block Storage for AWS using either Elastic Block Storage (EBS) for most workloads or the EC2 Instance Store for more performance optimized workloads.

If you choose to deploy Dell APEX Block Storage directly to instances backed with general purpose EBS, Dell APEX Block Storage can aggregate the shared storage resources of the lower performing and cost-effective disks to achieve much higher performance and resilience than native cloud offerings.

How can we achieve this?

The solution aggregates the resources and then shares them across the entire cluster. APEX Block Storage then automatically maintains balance across all resources, to support each application's need. You can also add storage and/or compute on-the-fly with no downtime or impact to applications because APEX Block Storage seamlessly balances the available resources.

The same flexibility exists with Dell APEX Block Storage for other hyperscalers. The software can be deployed on either managed disks for most workloads or on instances with native attached NVMe SSDs for more performance optimized workloads. A deployment of Dell APEX Block can provide these benefits:

- **Flexible deployments:** Provide options to optimize cost and performance based on your workload requirements.
- **Seamless data mobility**: Dell APEX Block Storage offers several ways to move data from on premises to the cloud and back, as well as between regions in the cloud. Features include native asynchronous replication, snapshots, and integrated backup and recovery solutions. Dell APEX Navigator for Multicloud Storage will orchestrate data mobility capabilities, such as snapshot shipping, in a coming release of APEX Navigator.
- **Multi-AZ durability**: Dell APEX Block Storage has a unique ability to spread data protection across multi-availability zones, ensuring data access without having extra copies of data or needing to use replication across AZs.
- **Enterprise data services**: The product provides thin provisioning, snapshots, and replication which can be required for mission critical workloads on the public clouds. Dell APEX Block Storage native replication enables data mobility between on-premises and cloud and/or to have a warm standby copy in another region or on-premises for additional protection against a regional disaster.
- **Performance and linear scalability**: Dell APEX Block Storage is well suited to support very large databases, analytics workloads, and multiple container (Kubernetes) development and production deployments.
- **The cloud experience** – anywhere: Being able to access and interact with cloud services, resources, and applications from anywhere, regardless of the physical location of the data.

Organizations today have data and applications spread across multiple data centers, co-location facilities, public clouds, and increasingly at the edge. Customers enjoy the ease and agility of the cloud experience and are demanding this same experience across all of their deployments. Simply put, this is the business case for Dell APEX Block Storage for Public Cloud—the cloud experience seamlessly delivered wherever organizations have data and applications.

To successfully deliver the business use case to end users, the APEX Block Storage solution provides these benefits.

Best of breed innovation:

- Cloud native storage built for cloud native development
- Storage innovation where it's needed, extending existing IT investments and skillsets.

Data mobility:

- Mobility geared toward both those who already have a PowerFlex on-premises environment and want to migrate to the cloud, and those who have already deployed Dell APEX Block Storage in the cloud and want to move data among various Dell APEX Block Storage deployments.
- Consume the public cloud PaaS services, with only the data attributes that you want, in the cloud.
- Consume the rich hyperscaler PaaS feature set reporting capability by mounting volumes copied from on-premises locations.
- Allowing for seamless data mobility, data teams can now place the right data, at the right time, in the right place.
- Snapshot shipping of volumes to the hyperscaler ecosystem to provide a source for, and to enhance, AI / ML model training.
- Consume just what you need in the hyperscaler of choice, without incurring any egress fees.
- Copy and mount snapshots in the cloud.
- Kubernetes namespace data mobility. Inclusive of all PV and PVCs, with an option to rename a namespace in flight.
- Evolving S3 Object solution sets as a queryable source for a modern database system.
- NOSQL databases – Develop in cloud. Migrate to on-premises or other cloud. Following the *no lock in* pattern of NOSQL distributed systems.

Data storage elasticity on AWS offers on-demand scaling in terms of the additional storage size

Ease of deployment and provisioning:

- Complete automated deployment in less than two hours when deployed with APEX Navigator
- Scaling without physical hardware acquisition channels
- Network automation
- Code upgrades for the APEX Block Storage software defined solution
- Custom deployment guidance:
  - Step-by-step instructions for manually adding to an existing Dell APEX Block Storage infrastructure.
  - Support for use cases not addressed by the orchestrated APEX Navigator deployment option.

Fully automated experience with APEX Navigator:

- Centralize APEX Block Storage deployment and lifecycle management across multiple public clouds, including configuration and automated provisioning and deployment. This process requires only four simple steps.

- Drive data mobility between on-prem and multiple public clouds to meet shifting cloud strategies.
- API-first architecture makes it easy to integrate with popular automation tools such as Ansible and Terraform to accelerate productivity and embrace Dell software defined storage in the public cloud as part of your broader IT strategy.
- Take decisive action with intelligent insights to leverage comprehensive monitoring across your entire Dell storage estate. You can also track license inventory, system health, performance, and capacity with CloudIQ integration.
- Accelerate Zero Trust adoption with role-based access control, single sign-on, and federated identity, all leveraging Zero Trust principles.

In addition to all the benefits that Dell APEX Block Storage provides for on-premises deployments (including extreme performance with sub-millisecond latency, near linear scalability, flexible deployment architecture, and stringent resilience), APEX Block Storage for Public Cloud also offers:

- The ability to achieve extreme performance in the public cloud.
- Linear scalability to 100s of storage instances to address large workloads and to consolidate large sets of applications in one cloud.
- Can be consumed by thousands of compute instances.
- Rapid deployment through an intelligent web-based wizard that removes the complexity of finding which instances best fit the workload and the capacity required. The solution supports the expansion and monitoring of your capacity, as needed.
- A unique use of Availability Zones which, mapped to fault sets, enable extremely high resiliency with a space-efficient and cost-effective use of the distributed mesh data across the data store. This entails enabling ultra-fast reprotection of an entire AZ that could go offline without disrupting access to your Dell APEX Block Storage.
- A choice of running block-based cloud storage on either EBS volumes or an EC2 instance store for performance optimized workloads.

APEX SaaS with Navigator

Imagine a centralized SaaS experience for managing your digital transformation across clouds and infrastructure services. This is our vision for what APEX will deliver – and that experience will be realized through APEX Navigator, which provides simple, secure storage software management spanning multiple storage types deployed in public clouds and on-premises, all centrally accessible within the APEX Console.

As part of the APEX Navigator, there are two sets of enhanced features that customers can select based on their storage management requirements: multicloud storage and Kubernetes.



Figure 1. Dell APEX Navigator

With APEX Navigator for Multicloud Storage you can:

- Centralize Dell Storage software deployment and lifecycle management across multiple public clouds, including easy configuration and automated provisioning and deployment, a process that requires only four simple steps
- Drive **data mobility** between on-prem and multiple public clouds to meet shifting cloud strategies
- **Integrate with popular automation tools** such as Ansible and Terraform, as part of our API-first architecture, so you can accelerate productivity and embrace Dell storage software in the public cloud as part of your broader IT strategy
- Take decisive action with **intelligent insights** with comprehensive monitoring across your entire Dell storage estate to track license inventory, system health, performance, and capacity with CloudIQ integration
- Accelerate **Zero Trust adoption** with role-based access control, single-sign on and federated identity, all incorporating Zero Trust principles for complete control of data.

For public cloud SDS deployment, it all starts with these four simple steps:

1. Select the storage product offering and target hyperscaler.
2. Select and/or add the appropriate cloud access.
3. Select the deployment configuration (storage tier and capacity, network, and so on).
4. Review and deploy.

After the storage software is deployed, APEX Navigator for Multicloud Storage serves as a centralized hub for day-to-day management activities using in-context navigation to access additional element managers. Customers can also monitor the licenses, health, performance, and capacity of their storage deployed in public cloud, or click into CloudIQ for a complete view of information across their Dell storage landscape. APEX Navigator

for Multicloud Storage also drives data mobility between on-premises and public cloud Dell storage.

In conjunction with these multicloud storage capabilities, you can also manage storage requirements for your containerized workloads using the Kubernetes capabilities featured in the APEX Navigator.

APEX Navigator for Kubernetes simplifies Kubernetes persistence management across multicloud, multisite environments.

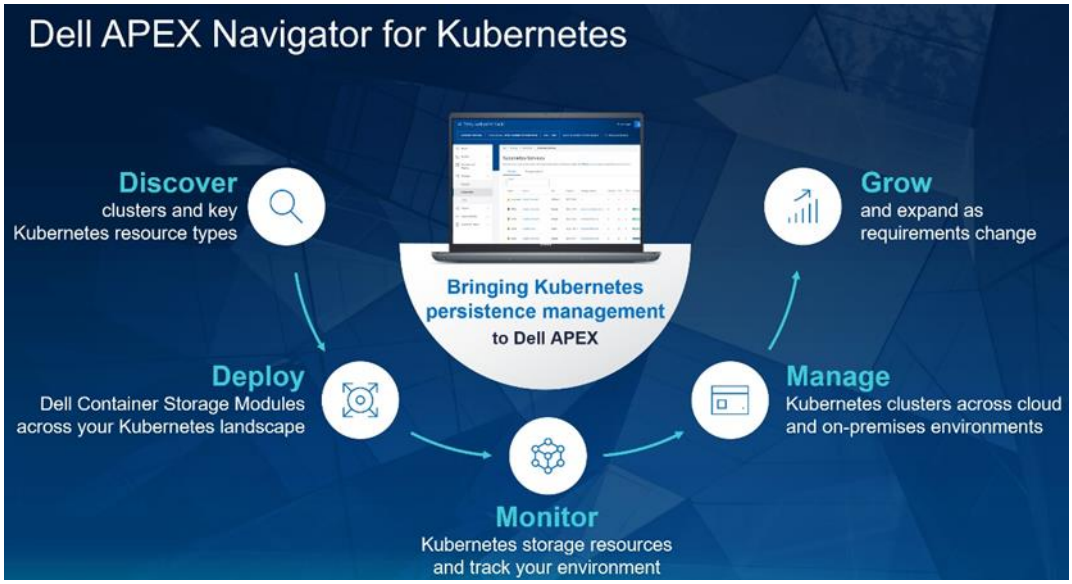


Figure 2. APEX Navigator for Kubernetes

Compared to the APEX Navigator for Multicloud Storage offering (which focuses on provisioning and managing Dell block, file, and data protection across clouds), APEX Navigator for Kubernetes focuses on deploying, managing, and operating **advanced data services** for persistent storage requirements around Kubernetes and containers, such as application mobility, replication, and authorization.

APEX Navigator for Kubernetes streamlines the deployment of Dell's Container Storage Modules at scale across multicloud Kubernetes environments. This improves productivity through automated workflows that allow DevOps team members to dedicate time and effort to innovation and operations.

With the APEX Navigator for Kubernetes, our entire Kubernetes enterprise storage portfolio is purpose-built to integrate with Dell storage. For the release of APEX Navigator for Kubernetes, the offering will support customers running Dell PowerFlex. APEX Navigator for Kubernetes makes the integration process easy for our Container Storage Modules, which expose Dell's storage systems to containerized workloads and allow us to take advantage of all the native array capabilities Dell has to offer:

Integrated zero trust security (SSO, RBAC, and Identity Federation)

Unique performance and deployment delivering extremely high I/O performance is cheaper with APEX Block Storage when compared to public cloud native instances.

A 90-day evaluation license (applies to Dell licensing only)

Federated storage protection with multi-availability zone resiliency and flexibility, protection, and scalability

Data mobility:

- Moving data from on-premises to cloud (ground to cloud)
- Retrieving data from the public cloud (cloud to ground)
- Meet any shifting cloud strategy
- Crash consistent copies
- Efficient – Changed block only snapshot shipping

API-first architecture allowing customers to create their own artifacts for API integration with many of the most popular automation tools, such as Ansible and Terraform.

Current AWS region support for APEX Navigator deployments. (For details, see Regions and availability zones.)

Controlling cost:

- Predictable
- Data services that the cloud alone cannot provide
- Better data services at a reduced cost
  - Thin provisioning providing an efficient storage footprint for volume management.
  - Multiple, efficient, and non-constrained snapshots when compared to native cloud offerings.

Monitoring all endpoints from a single pane of glass:

- Robust monitoring—View of all systems, public or private, to observe and act on the performance, health, inventory, and capacity
- Licensing inventory

Cloud based AIOps for proactive monitoring, machine learning, and predictive analytics so you can take quick action and simplify operations of your on-premises infrastructure and data protection in the cloud.

Custom deployment

For customers who require a custom deployment option, please email APEXBlockStorage.AWS@dell.com.

TCO considerations

The Dell APEX Block Storage solution for public cloud offers an enterprise grade feature set that is not present in EBS. It does so at a price point that over a short amount of time can be significantly lower, dollar-per-I/O, than EBS.

Dell APEX Block Storage can exceed the capacity and performance limitations of any AWS EBS type that will cost less to run overall. The performance of an APEX Block

Storage volume is directly proportional to the scale of the APEX Block Storage cluster created. Simply stated, bigger cluster = faster volumes.

When we look deeper into the total cost of ownership (TCO), note the core tenets of the APEX Block Storage solution: its operational and cost benefits. These core tenets are:

- Data mobility (with Async Replication) to one or many hyperscalers
- The need for a common platform on-premises and in any hyperscaler
- Thin provisioning options that the cloud native offers do not permit.
  - On-premises, a customer can only deploy 2:1 over-provisioning with thin-provisioned volumes.
    - Procurement time for hardware expansion can be long
    - Risk of running out of capacity due to long procurement
    - Customers can overprovision at rates at or exceeding 4:1 without any concern. This would not be an option for on-premises deployments.
- Scale up and scale down the storage cluster to fit current workload needs. Pay for only what you need, when you need it.

Reference to AWS pricing calculator

A good way to understand the list price cost for an APEX Block storage deployment is to map the deployed instances into the AWS pricing calculator. Deployed instances can be found in the Deployment topologies section of this document. As with any pricing guidance for public cloud, the client enterprise agreement will govern the ultimate consumption pricing schedule.

SLA reimagined

On average, cloud providers will deliver anywhere between two to four nines of availability. With each “nine” there quickly adds exponential additional cost. Dell APEX Block Storage for AWS pools together lower cost and lower durability volumes, all while presenting six “nines” of availability[b] to your applications. By aggregating all the lower cost EBS volumes, Dell APEX Block Storage can concentrate the collective performance of those pooled volumes to achieve a IOPs/dollar ratio that is not possible with native hyperscalers.

For a cloud first, or cloud smart, initiative (and there is neither an appetite nor a skills gap that exists to refactor apps), APEX Block Storage for public cloud can deliver without the cost of application re-factoring and redesign, all while maintaining the most stringent on-premises SLAs. Combining the best of Dell software defined storage and cost-effective cloud agility can help satisfy the most mission critical SLA requirements.

Dell APEX Block Storage provides thin provisioning of storage to a VM consumer and highly effective snapshots without limitations, when compared to native AWS.

[b] Based on Dell analysis of storage software deployable on AWS, Azure, and Google Cloud, May 2023.

Cost considerations

Dell cost

When you deploy Dell APEX Block Storage for the first time, you are automatically provided with a 90-day evaluation period for the Dell APEX Block Storage solution. At the end of 90 days, you can obtain a subscription license, which allows you to use Dell APEX Block Storage for a defined period of time.

AWS cost

The AWS instances that are deployed as part of the Dell APEX Block Storage cluster are aligned to the client subscription. It is the responsibility of the client to maintain the monthly subscription payment, based on the AWS and client agreement, along with the negotiated consumption rates.

Storage cost considerations

Charges are applied to hold the EBS storage in perpetuity when instances backed with EBS storage are deallocated (created but not running), and the APEX Block Storage cluster is gracefully shut down.

Local instance store drives are not perpetually associated with the instance, because an instance can be re-instantiated on a different physical server. Therefore, every time an instance is stopped, AWS wipes the instance store capacity. In this case, Dell PowerProtect with Data Domain Virtual Edition would be required for backup of the data.

Scaling down a cluster deployed into AWS is also allowed. When additional capacity is no longer needed, the cluster nodes can be trimmed back to meet the new requirements.

For more information, see Amazon EBS pricing.

Amazon marketplace

The current Dell APEX Block Storage offering can be purchased from the AWS marketplace.

Unique multi-AZ durability

The Dell APEX Block Storage self-healing architecture simplifies and increases the availability of infrastructure. It is designed to remain running even when it encounters failures. With fast rebuilds, Dell APEX Block Storage recovers rapidly from failures and keeps operating, without degrading performance or compromising availability. You can also extend this availability and configure it to protect your application availability across multiple availability zones.

Durability is defined by:

- Using Dell APEX Block Storage Fault Set architecture
- When spanning a single cluster across multiple availability zones or deploying into a single zone, the solution aggregates the capacity and performance. Having more instances means more performance and more resiliency.
- Retaining critical response time and IOPS performance properties
- Rapid re-protects from an instance failure, or even an AZ failure, with rapid rebuilds

Architecting fault sets, or fault domains, within the public hyperscalers provides unprecedented levels of storage availability that were previously unavailable in any public cloud platform.

The unique multi-AZ architecture is defined by:

- The industry-first feature of a single cluster spanning multiple availability zones allows customers to set up non-replication-based fault tolerance across an entire AWS region
- Uses EC2 Instance Store or EBS
- Scales linearly by federating instances
- No performance limitations
- Chooses your instance configurations based on the use case
- Protection across Availability Zones

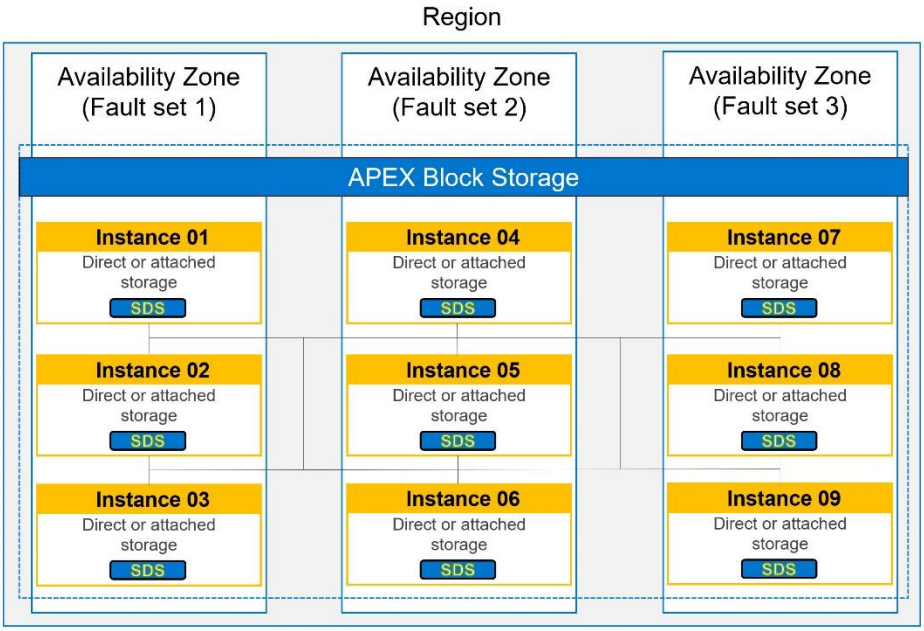


Figure 3. Multiple availability zone architecture

Resourceful workload placement

Knowing where and how to implement efficient workloads is not always a simple task. It takes time, monitoring, and insights. It must also take into consideration big islands of data, either in a database, or loosely stored in file systems. Historically, applications orbited around these large data stores. Fast forward to today, where the right data needs to land in the right place, at the right time. Being able to place data volumes close to the app that consumes them has now become as simple as replicating to the correct location, then mounting and consuming these volumes.

Technical Case - Dell APEX Block Storage AWS-specific architecture

To improve outcomes and repeatability, the number of configurations for running Dell APEX Block Storage has been optimized for the most commonly used configurations. Aligning with dev ops principals, this insures a repeatable and consistent environment that is easier to maintain when compared to providing all possible permutations available in the catalog.

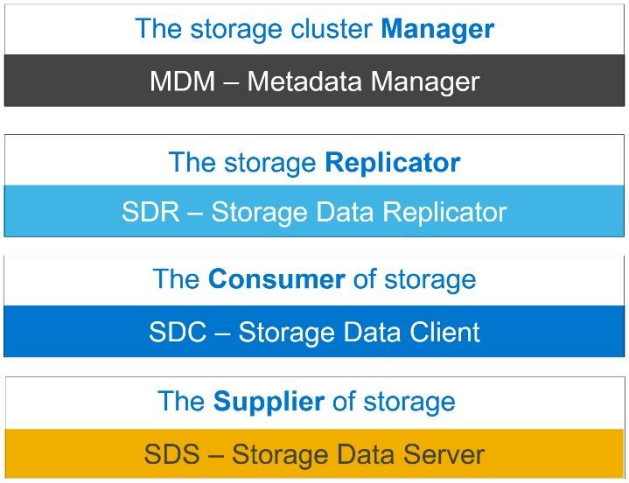


Figure 4. APEX Block Storage building blocks

Dell APEX Block Storage is software defined by design.

To make it run, we need hypervisor virtualized instances, networking, and the Dell APEX Block Storage building block software components.

Dell APEX Block Storage components

The fundamental building blocks of an APEX Block Storage solution are:

MDM – Metadata Manager – The brains of the storage layer

- Oversees storage cluster configuration, monitoring, rebalancing, and rebuilding
- A highly available, independent cluster installed on three or five different nodes
- May reside alongside SDCs and/or SDSs or on separate instances
- Located outside the data path

SDR – Storage Data Replicator

- Handles replication activities between Dell APEX Block Storage systems
- Located alongside the SDS on the same node
- Highly available. If failure occurs, the MDM instructs the SDC to send data to alternative SDRs.
- The Source journals incoming writes and ships them to the target
- The Target receives the journal and applies it to the target volumes

SDC – Storage Data Client (Consumer of storage)

- Provides front-end volume access to applications and file systems
- Installed on instances that are consuming storage
- Maintains peer-to-peer connections to every SDS managing a pool of storage
- Understands how to make read and write requests to different parts of the system

SDS – Storage Data Server (Supplier of storage)

- Abstracts local storage, maintains storage pools, and presents volumes to the SDCs
- Installed on instances that contribute local disks to the storage cluster
- A protection domain is a collection of SDS instances
- Manages the data, that is spread evenly across all of the contributing nodes in the cluster. This provides the opportunity for massively parallel processing.

APEX Block Storage for AWS management instances (MDM and MDM tiebreaker), Secure Connect Gateway (SCG), APEX Block Storage Manager Platform, and proxy server instances use AWS EBS gp3 volumes exclusively.

Architecture

The APEX Block Storage deployment contains instances that may or may not be physical machines, depending on the type of instance selected.

For the deployment, the management stack includes APEX Block Storage Manager, the proxy server, and the temporary management platform installer.

APEX Block Storage does not include SDCs (the consumers of storage). The customer is responsible for deploying the compute instances on AWS instances of their choice. SDC details are listed in the section Downloading and installing the SDC.

Metadata communication

Within a deployed APEX Block Storage solution, only the SDS and the SDC exchange data. The MDM's role is completely outside the data path. Its responsibility is to distribute cluster metadata about the data's layout, about the storage cluster members (the SDS and SDCs), and to make sure that they are still operational.

The key constructs of the Metadata Manager are:

- The MDM continuously monitors the SDS and storage
- The MDM provides system metadata and storage configuration information to the SDC
- There are three or five MDMs per system
- There is only one primary MDM at a time
- There are one or two secondary MDMs
- There are one or two tiebreakers

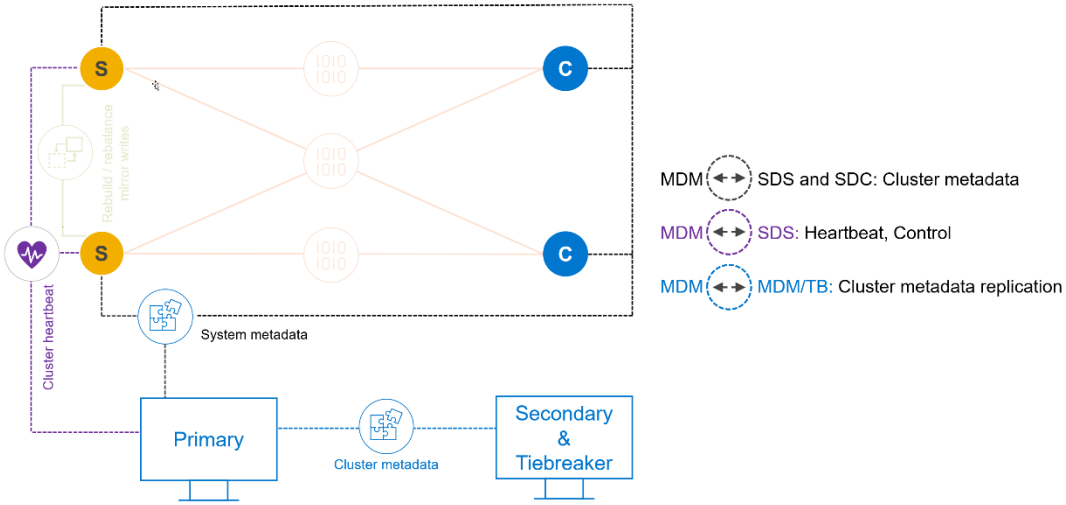


Figure 5. MDM communication

The MDM topology is clustered by design. The primary MDM replication changes to system status synchronously. The MDM monitors the SDS status continuously and informs SDSs of changes to the system and MDM status, regarding instances and data layout. Device awareness can trigger rebuilds ahead of a potential failure.

Note: Temporary failures will not trigger a complete rebuild for a failed instance of storage device if it comes back online.

Regions and availability zones

AWS data centers exist in various regions throughout the globe. APEX Block Storage for AWS can be deployed in any region as a standalone product, where the hyperscaler resources are available. This standalone product is supported in US English.

Consideration should be given to the number of availability zones, if multi-az APEX Block Storage topologies are required.

Availability zone support

Dell APEX Block Storage for AWS supports both single and multiple availability zone deployments. With multiple AZ deployments, a single cluster will span across the availability zones.

Table 2. ECS instances supported

AZ type	Fault set	Instance / EBS volumes
Single AZ	None	• c5n.9xlarge with 10 EBS volumes (gp3) • i3en.metal with 8 x 7.5 TB NVMe devices
Multiple Availability Zones (single region with three AZs)	Three (one per AZ)	• c5n.9xlarge with 10 EBS volumes (gp3) • i3en.12xlarge with 4 x 7.5 TB NVMe devices

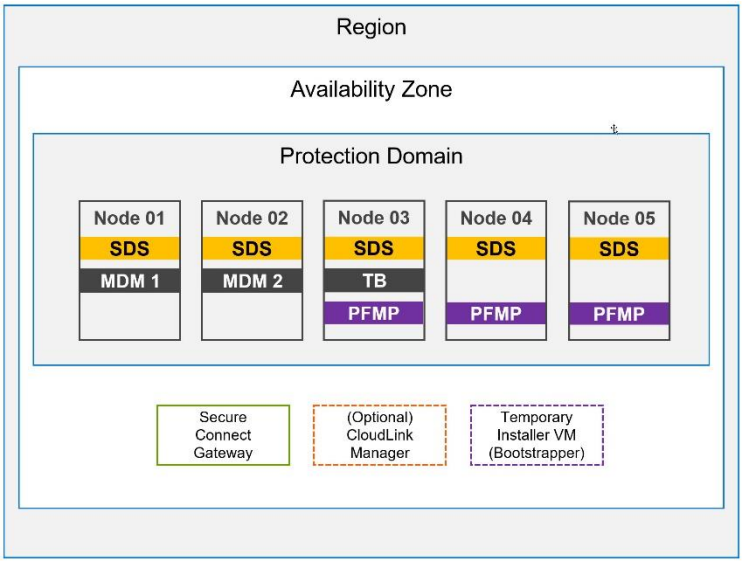


Figure 6. Single AZ logical design

For an explanation of EC2 instance types and drive calculations, see the section Dell APEX Block Storage for AWS limits.

AWS Storage considerations

When selecting the proper AWS storage, the following matrix can help refine your requirements.

You will need access to and permissions from your AWS Infrastructure administrator to access the AWS account for Dell APEX Block Storage.

Because the AWS infrastructure does accrue a cost during the deployment process, it is recommended to have an approved budget for the initial deployment, even if it is only for testing.

Table 3. Balanced use cases

EBS gp3 – Balanced use case
Simple and easy
Inherently persistent
For customers that power on/off regularly
Requires more instances to reach higher performance levels
Software attached and handled by AWS

Table 4. Performance use case

EC2 NVMe – Performance use case
Locally attached at the instance level.
Higher performance with fewer instances. Can be up to 10x.
No additional storage charges but requires more EC2 instances when considering capacity, which do incur an instance cost.

EC2 NVMe – Performance use case
When turning off all nodes of the cluster, persistence requires additional added configuration for resilience.
NVMe drives are ephemeral - Data must be backed up to avoid data loss if Dell SDS cluster is shut off.
Can use DDVE to S3 as a backup solution.

The following table summarizes Dell APEX Block Storage for AWS supported capabilities and system limits.

Table 5. Supported capabilities and system limits

Item	Limit
System raw capacity	16 PB
Device size	Minimum: 240 GB Maximum: 15.36 TB (SSD)
Minimum storage pool size	720 GB
Minimum devices (drives) per storage pool	3, one per fault set
Volume size	Minimum: 8 GB, Maximum: 1 PB
Maximum file system partitions per volume	15
Maximum total number of volumes and snapshots in a system	131,072
Maximum total number of volumes and snapshots in a protection domain	32,768
Maximum total number of volumes and snapshots per storage pool	32,768
Maximum number of snapshots in a single vTree	126
Maximum raw capacity per protection domain	8 PB
Maximum raw capacity per SDS	160 TB (medium granularity)
Maximum SDCs per system	2048
Maximum SDSs per system	512
Maximum SDSs per protection domain	128
Maximum devices (drives) per SDS server	64
Maximum devices per protection domain	8192
Maximum devices per storage pool	300
Total size of all volumes per storage pool	4 PB
Maximum volumes that can be mapped to a single SDC	1024
System Over Provisioning Factor	5X of net capacity per MG layout
Maximum protection domains per system	256
Maximum storage pools per system	1024

Item	Limit
Maximum storage pools per protection domain	64
Maximum fault sets per protection domain	64
Maximum number of snapshots a snapshot policy can be defined to retain per vTree (not including locked snapshots)	60
Maximum snapshot policies per system	1000
Maximum user accounts	256
Maximum number of concurrent logged-in management clients (GUI/REST/CLI)	128
Maximum volumes that can be mapped by API (GUI/REST/CLI) concurrently	1024
Maximum number of configured syslog servers	16 (must be an even number)
Volumes per local consistency group (snapshot)	1024
Maximum number of volume-to-SDC mappings per system	262,143
SCSI-3 reservation type	Write exclusive registrants only
Number of destination systems for replication	4
Maximum number of SDRs per system	128
Maximum number of replication consistency groups (RCGs)	1024
Maximum replication pairs in RCG with initial copy	1024
Maximum number of volume pairs per RCG	1024
Maximum volume pairs per system	32,000
Maximum number of remote protection domains	6
Maximum number of copies per RCG	1
Recovery Point Objective (RPO)	Minimum: 15 seconds Maximum: 1 hour
Maximum replicated volume size	64 TB

Considerations for Availability Zone deployments

Dell APEX Block Storage allows you to configure a deployment across multiple availability zones in a single region. This makes it possible to place SDSs in different AZs. Said differently, this is a single cluster that can span multiple AZs. If one AZ is down, the cluster is still fully operational. Spreading data across multiple AZs in this way creates federated distributed protection. This topology keeps the resiliency at the base storage layer, instead of relying on complicated resiliency architectures that exist higher up in the application stack. This is not something a standard cloud deployment can provide. Dell APEX Block Storage is unique in this space.

Note: There is a data transfer charge between Availability Zones. For pricing details, see Amazon EC2 On-Demand Pricing.

While the multiple-AZ configuration contains the same components as a single-AZ deployment, there are a few noticeable differences to the deployment:

- In a multiple-AZ deployment, fault sets are used to ensure that in the rare event of a total AZ failure, no storage data loss occurs. One fault set is assigned per AZ.
- The placement of MDMs (and the PFMP cluster) is also spread across AZs to minimize the chance of a complete management stack failure.
- Multiple-AZ deployments require a minimum of six SDS instances. Additional SDS instances must be deployed in multiples of three to maintain equal numbers of SDS instances in each fault set. This provides balanced capacity and performance between the SDSs as capacity is expanded.

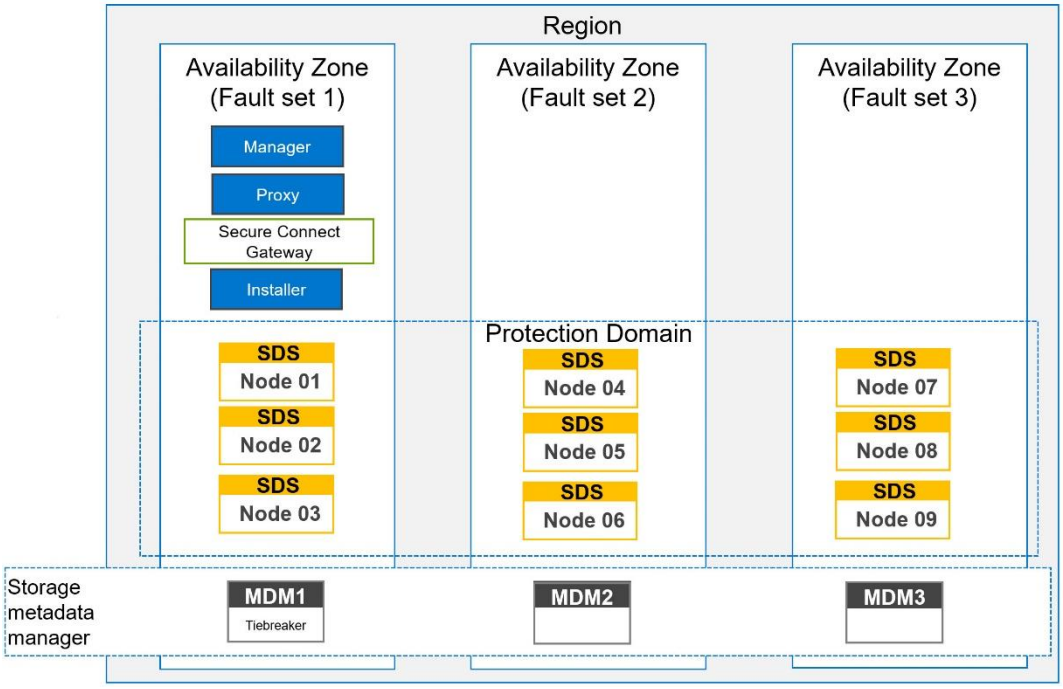


Figure 7. Multiple AZ logical design

Deployment topologies

Instance configuration

Supported AWS instance configurations.

Table 6. AWS instances

Model	vCPU	Memory (GiB)	Instance Storage (GB)	Network Bandwidth (Gbps)
c5n.9xlarge	36	96	EBS-Only	50
i3en.12xlarge	48	384	4 x 7500 NVMe SSD	50
i3en.metal	96	768	8 x 7500 NVMe SSD	100

Storage configuration

The solution pools all the storage resources together, leaving no resources stranded.

Dell APEX Block Storage stores two mirrored copies of each block of data. When a volume is created, that the data for that volume gets spread across all the disks running across all instances. Upon write, the SDS writes a chunk to one disk and another SDS writes the mirrored chunk to another disk. This process automatically protects data, and optimizes and balances resources, based on each application's needs.

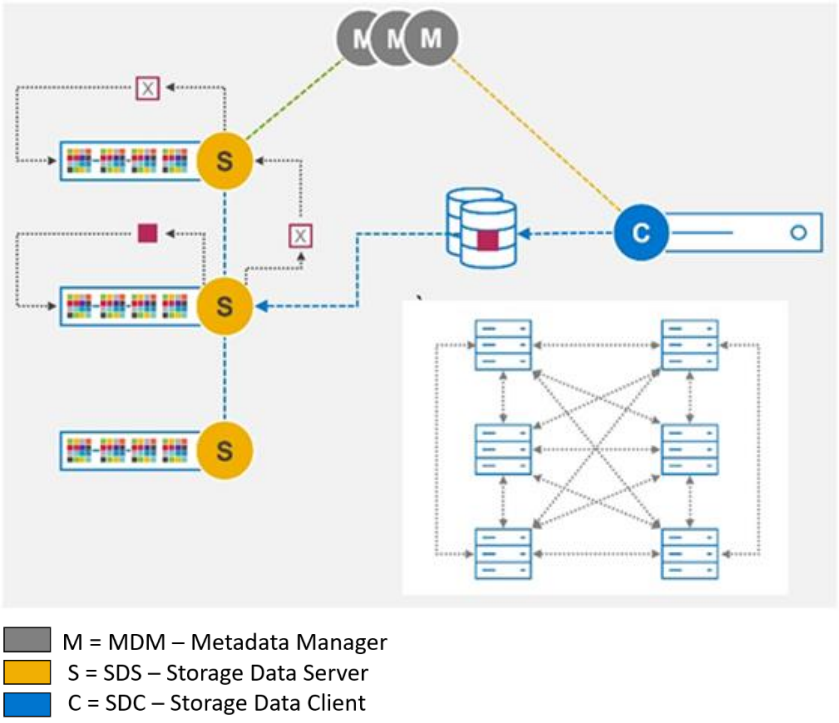


Figure 8. Volume data mirroring and resiliency

- Every write to a volume is written both to the primary SDS and replicated to another SDS
- If a node or disk fails, the primary and secondary copies are quickly rebuilt on the remaining nodes/disks
- The many-to-many relationship among SDSs provides for fast, massively parallel rebuilds and improves availability

There are three basic types of storage capacity:

- **Raw Capacity** – sum of capacity in all disks that are used for storage.
- **Usable Capacity** – sum of capacity available for the client's applications after deducting the mirrored copy and virtual spare space.
- **Spare Capacity** - Additional spare capacity is required to ensure that enough space is available for rebuilding the largest node or fault set in case of failure (for example losing one AZ in a three AZ solution).

Table 7. Storage type raw capacity and use case

Storage	Raw capacity per SDS instance	Use case
EBS gp3	c5n.9large: 5,10,20,40 TB	Data is persistent during planned or unplanned shutdowns.
NVMe SSD	i3en.metal: 8 x 7.5 TB (60TB) i3en.12xlarge: 4 x 7.5 TB (30 TB)	Extreme performance. Data does not persist following either a planned or unplanned shutdown. A backup solution is recommended.

Table 8. Minimum Raw storage capacity for deployment configurations

Raw storage capacity per instance	EC2 instance used	Underlying storage media	Recommended Min raw capacity per cluster	Max IOPs / instance	Supported config (Single vs Multi AZ)
5TB	c5n.9xlarge	gp3 - EBS Volumes	15TB	40K	Single / Multi AZ
10TB	c5n.9xlarge	gp3 - EBS Volumes	30TB	40K	Single / Multi AZ
20TB	c5n.9xlarge	gp3 - EBS Volumes	60TB	40K	Single / Multi AZ
30TB	i3en.12xlarge	EC2 - Local NVMe SSD	90TB	260K	Multi AZ
40TB	c5n.9xlarge	gp3 - EBS Volumes	120TB	40K	Single / Multi AZ
60TB	i3en.metal	EC2 - Local NVMe SSD	180TB	260K	Single AZ

Storage API operations

Dell APEX Block Storage API operations are divided into three categories.

- The legacy API, which is called the block API. It is used to complete tasks related to block, protection, and operations on the core software.
- The Dell APEX Block Storage Manager API, used for lifecycle management of Dell APEX Block Storage deployments, and Dell APEX Block Storage core software.
- The PowerAPI is used to manage SSO, events, and alerts. It is based on a new Dell Technologies standard for product APIs.

The web service for the REST APIs is a single endpoint, for all three categories, hosted on the Dell APEX Block Storage management platform, simplifying programmatic operations, including authentication.

PFMP uses the OAuth 2.0 industry standard protocol for authorization. In OAuth 2.0, a user authenticates with an authorization microservice in PFMP. Upon authentication, the user is passed an access (bearer) token. The bearer token enables the user to access protected resources and to perform protected operations on all areas of the Dell APEX Block Storage system.

For additional information about the REST API, see the white paper PowerFlex REST API Introduction.

Storage Data Client Connectivity

The storage data client is a kernel-level driver. This client installs on any host instances like a virtual HBA. It resides on all nodes that consume storage from the storage pool. The client maintains its peer-to-peer connection to every SDS managing the storage pool, allowing it to map the storage to the OS as a block device.

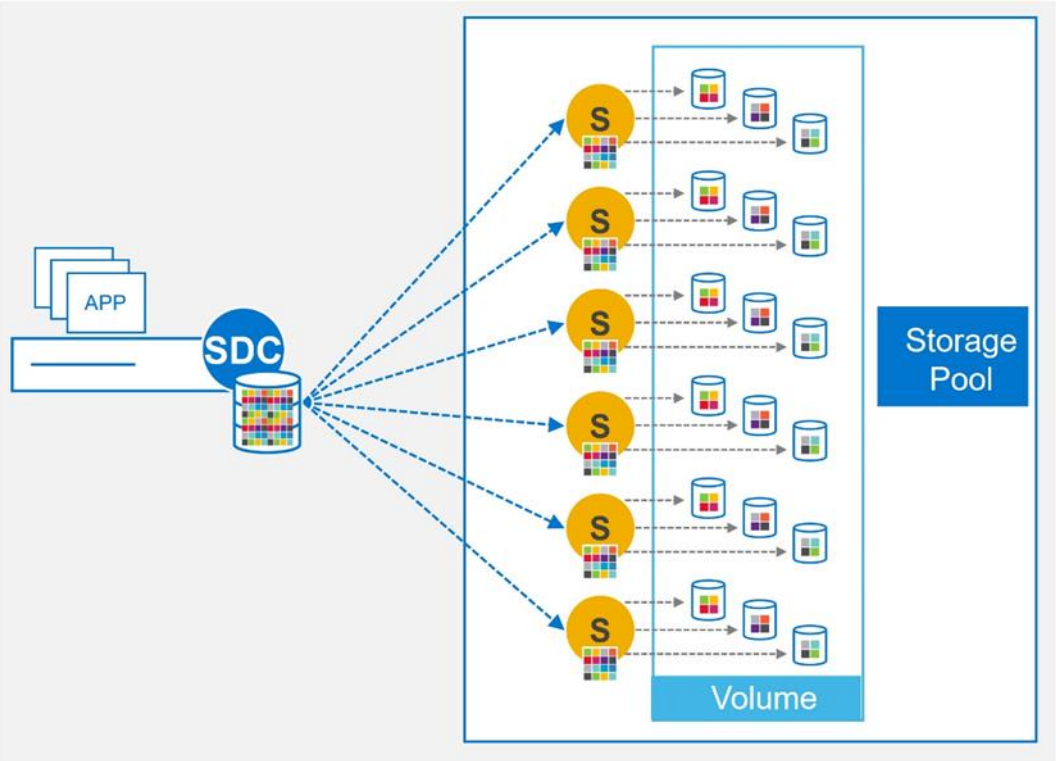


Figure 9. Storage data client connectivity

Downloading and installing the SDC

Confirm that you have the required installation software package for the operating system of the SDC instance that will be consuming the APEX Block Storage. The installation files are available for download as the “PowerFlex Complete Software Download” on the Dell Technologies support site.

The SDC can be downloaded from the Dell support website here.



Figure 10. SDC Download details page

Note: The build numbers in this figure are for example only, and do not reflect those in the current build.

Create your AWS virtual instance according to your defined processes and procedures. Ensure that the cloud instance that will utilize the SDC, and the one that will consume the APEX Block Storage, exist in the same AWS region. While it is possible to connect SDCs to APEX Block storage instances across different regions, the resulting data transfer cost would be substantial, and performance would not be optimal.

The second zip file in the package contains all the SDC files that will need to be manually installed on the machine instances within the public cloud environment that will be consuming the APEX Block Storage system.

There is no official limitation regarding the minimum instance size. However, the instances must be able to supply enough network bandwidth to support the anticipated workload speed. Because the storage traffic traverses the same connection as the rest of the network traffic, Dell Technologies recommends that you select instances with a minimum 10 Gb uplink speed.

The SDC must have proper access to the APEX Block Storage environment:

- Connectivity: The network must be reachable. The SDC can use the same subnet as the APEX Block Storage deployment. If you use a different subnet, configure the routing to ensure proper connectivity.
- Security groups: Adjust the security groups for both the SDC and the APEX Block Storage side to allow connectivity.

For any OS instance, ensure that you have the following:

- The IP address of the proxy server
- The IP address of each SDC Windows or Linux instance
- Login credentials for the Windows or Linux instance
- The IP addresses of the MDM management instances

Operating system requirements

The following is a list of operating system components supported by APEX Block Storage.

Operating System	Requirement
Linux	Packages required for all components, all Linux flavors: • numactl • libaio1 • wget • apr • libapr1 • python2-rpm or python3-rpm • which
Windows	• Only SDC, and LIA are supported components on Windows • To install SDC, ensure that Microsoft Security Update KB3033929 is installed. • To use the secure authentication mode, ensure that these are installed on all servers in the system: • OpenSSL 64-bit v1.0.1 or later (v1.1, however, is not supported) • Visual C++ redistributable 2010 package, 64-bit

Installation process for Windows instances:

Steps:

1. Copy the following files to the proxy server:
   - EMC-ScaleIO-sdc*.msi
   - EMC-ScaleIO-lia*.msi
2. From the proxy instance, open an RDP connection to the Windows instance and log in using the username and password of the instance.
3. Copy the SDC and LIA MSI installation files (EMC-ScaleIO-sdc*.msi, EMC-ScaleIO-lia*.msi) to C:/Users/ec2-user on the instance.
4. Run the following commands to install the software:

   ```
   msiexec /i EMC-ScaleIO-sdc*.msi MDM_IP=<list_mdm_ips>
   msiexec /i EMC-ScaleIO-lia*.msi
   ```

   The PowerFlex Installer opens.
5. At the prompt, accept the license agreement.
6. When the installation is complete, close the PowerFlex Installer status message, and click **Mark operation completed**.
7. Repeat Step 2 through Step 6 for any additional Windows instances.

Installation process for Linux instances (Ubuntu):

Steps:

1. Copy the following files to the proxy server:

- scini.ko
- EMC-ScaleIO-sdc*.deb
- EMC-ScaleIO-lia*.deb

2. Log in to the proxy server and open a terminal window.
3. Install the generic kernel (for example, 5.4.0-113-generic):

   ```
   sudo apt-get update && sudo apt install linux-generic -y
   ```

4. Open /etc/default/grub.cfg in a text editor and edit the default boot option to change the kernel version to the respective one on the system:

```
set GRUB_DEFAULT="Advanced options for Ubuntu>Ubuntu, with Linux 5.4.0-113-generic"
```

5. Run the following command:

   ```
   sudo update-grub && sudo update-grub2
   ```

6. For Ubuntu Server 20 only, open the file `/etc/default/grub.d/40-force-partuuid.cfg` for editing and comment out the following line:

   ```
   GRUB_FORCE_PARTUUID
   ```

   For Ubuntu Server 18.04, skip the next step.

7. Reboot the OS.
8. Run `uname -r` to verify that the system has switched to the generic kernel.

   Output similar to the following is displayed:

   ```
   5..4.0-113-generic
   ```

9. Upload EMC-ScaleIO-sdc*.Ubuntu.*.x86_64.deb and scini.ko to the Ubuntu instance. The scini.ko file should match the generic kernel version (5.4.0-113-generic).
10. Upload EMC-ScaleIO-lia*.Ubuntu.*.x86_64.deb to the instance.
11. Copy scini.ko to the driver cache:

```
sudo mkdir -p /bin/emc/scaleio/scini_sync/driver_cache/Ubuntu/3.6.400.107/5.4.0-113-generic/
sudo cp scini.ko/bin/emc/scaleio/scini_sync/driver_cache/Ubuntu/3.6.400.107/5.4.0-113-generic/
```

12. Install the SDC.

    ```
    sudo MDM_IP=<list_mdm_ips> dpkg -i <sdc_path>.deb
    ```

    where

    - <*list_mdm_ips*> is a comma-separated list of all primary, and secondary MDM IP addresses.
    - <*sdc_path*> is the path where the SDC installation package is located.

13. Install LIA:

```
sudo TOKEN=My_PowerFlex_Password dpkg -i <sdc_path>.deb
```

Installation process for Linux instances (RHEL, SUSE, CentOS):

Steps:

1. Copy the following files to the proxy server:
   - RPM-GPG-KEY-ScaleIO_*
   - EMC-ScaleIO-lia*.rpm
   - EMC-ScaleIO-sdc*.rpm
2. Log in to the proxy server and open a terminal window.
3. Copy the SDC installation file (EMC-ScaleIO-sdc*.rpm) from the proxy server to the first SDC instance:

```
scp -i keyfile.pem EMC-ScaleIO-sdc-* ec2-user@<sdc_ip>:/home/ec2-user/
```

4. Copy the PowerFlex LIA installation file (EMC-ScaleIO-lia*.rpm) from the proxy server to the SDC instance:

```
scp -i keyfile.pem EMC-ScaleIO-lia* ec2-user@<sdc_ip>:/home/ec2-user/
```

5. Log in to the SDC instance using SSH and install the GPG key. From the install folder, run:

```
rpm --import RPM-GPG-KEY-ScaleIO
```

6. Run the following commands to install the software:

```
ssh -i keyfile.pem ec2-user@<sdc_ip>
sudo TOKEN=My_PowerFlex_Password rpm -ivh EMC-ScaleIO-
lia*.rpm
sudo MDM_IP=<list_mdm_ips> rpm -ivh EMC-ScaleIO-sdc*.rpm
```

7. Confirm that the driver is loaded:
   - For SLES 15:

     ```
     systemctl status scini
     ```

   - For all other Linux types:

     ```
     service scini status
     ```

Volume management

The smallest unit of allocation within an APEX Block Storage environment is a volume. Volumes start at 8GB and can be as large as 1PB. They are the logical container of blocks, physically located on many independent disks that a host will format with a file system. All volumes are separate and independent systems, with Dell APEX Block system coordinating their storage using standard TCP/IP networking.

After installing SDCs, volumes can now be added using the APEX Block Storage Manager.

When defining a volume name, be sure that it:

- Contains less than 32 characters
- Contains only alphanumeric and punctuation characters
- Is unique within the object type

Dell recommends that you give each volume a meaningful name that reflects its operational role.

Using the APEX Block Storage Manager, you are now ready to map the volumes to the SDC. Mapping exposes a volume to the specified SDC, effectively creating a block device on the SDC.

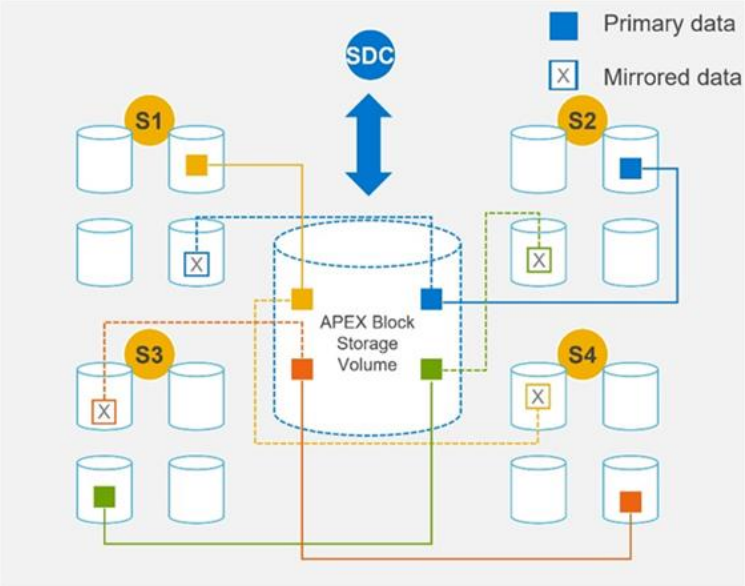


Figure 11. Volume management and mirroring architecture

- A volume appears as a single object to the application / file system
- The SDC is always accessing data from multiple devices on multiple nodes
- Logical collection of mirrored, distributed chunks in a Storage Pool
- SDC only accesses primary chunks

Data mobility

There are a couple of options to address data mobility with APEX Block Storage. The product provides native replication and snapshot shipping. With native replication you can asynchronously transfer data between APEX Block Storage systems. With snapshot shipping, the product can transfer entire volumes. This process will only be available and orchestrated with the APEX Navigator

Snapshot considerations

APEX Block Storage enables you to create instantaneous copies of volumes, that is, snapshots. A snapshot can be manipulated as any volume exposed by the storage system. Snapshots taken at the same time form a consistency group can be manipulated as a group.

The APEX Block Storage system enables you to take snapshots of existing volumes, up to 128 per volume. The snapshots are thin provisioned and are extremely quick to create. When a snapshot is generated, it becomes a new, unmapped volume in the system. You can manipulate it in the same manner as any other volume exposed by the APEX Block Storage system.

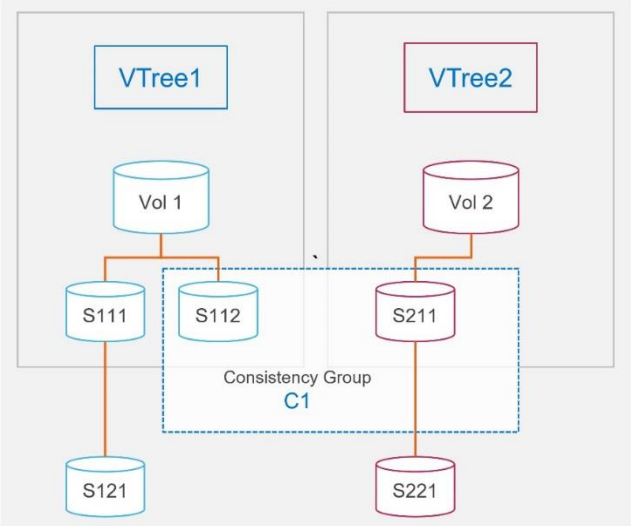


Figure 12. Snapshot operations and consistency groups

The structure related to all snapshots that result from one volume is referred to as a V-Tree. When taking a snapshot in the system, you can specify more than one volume. They are consistent in the sense that they were all taken at the same time. If there is a contextual relationship between the data contained by all the snapshot members, then that set is meaningful. The consistency group allows manipulation of the entire set.

If you remove an entire consistency group, all the snapshots that were taken together are removed.

Snapshot policies

Snapshot policies enable you to define policies where you can configure the number of snapshots to take at a given time for one or more volumes.

Snapshots are taken according to the rules defined. You can define the time interval in-between two rounds of snapshots as well as the number of snapshots to retain, in a multi-level structure. For example, take snapshots every x minutes/hours/days/weeks. There are one to six levels; the first level has the most frequent snapshots.

Secure snapshots

Snapshots can be created with a “Secure Snapshot” option. These snapshots are protected until the expiration time is reached. During this time, these rules apply:

- No data change
- No snapshot deletion
- No refresh

- No overwrite
- Are read only

Volume trees

A V-Tree consists of the root volume and all its descendant volumes and snapshots.

A V-Tree is the structure that consists of a volume and the snapshots resulting from that volume. It is a tree that spans from the source volume at its root, whose descendants are either snapshots of the volume itself or snapshots of a snapshot. In Figure 12, S_{111} and S_{112} are snapshots of V_1. S_{121} is a snapshot of snapshot S_{111}. Together, V_1 and S_{1xy} are the V-Tree of V_1.

Native replication

APEX Block Storage enables native asynchronous journal-based replication, with a defined RPO. With async replication you are also replicating entire volumes or sets of volumes (Replication Consistency Groups).

This crash-consistent replication can be used to recover from a physical or logical disaster quickly and easily. Additional use cases can align with data migration, test data at a remote site, or to offload a volume backup.

This replication can also be from an existing PowerFlex on-premises deployment to APEX Block Storage for AWS.

The SDR is the APEX Block Storage component that enables asynchronous journal based replication. The SDR is located alongside the SDS on the same nodes. The same scalability as part of the storage topology also exists for native replication. Only I/Os of replicated volumes will flow through SDR. This is governed by the metadata cache that the MDM manages on the SDC and SDS. All IO traffic for replicated volumes goes through the SDR (both reads and writes).

The RPO is down to 15 seconds which reduces the data-loss to a minimal time in case of disaster recovery.

The SDR orchestrates the following functions:

At the source:

- Forwards incoming writes to target volume SDSs
- Journals incoming writes
- Continuously sends journal data to the target system

At the target:

- Receives, stores, and applies journal data to target volumes in a consistent manner

Thin provisioning

Capacity is “on reserve,” and not allocated until used. This policy enables more flexibility in provisioning. Whereas thick capacity is limited to available capacity, thin capacity provisioning can be oversubscribed, as follows:

Maximum thin capacity provisioning = 5 * (gross capacity - used capacity)

When capacity usage reaches the level where it can cause IO errors, alerts are generated. At certain higher capacity levels, volumes (even thin volumes) can no longer be created.

Consequently, thick provisioning capacity is allocated immediately, even if not actually used. Thick capacity is limited to available capacity.

Network configuration

The basic network infrastructure for AWS is a network environment that includes routing, subnets, security groups, network gateways, IP address allocation, and other network requirements.

The following rules apply when using an AWS VPC:

- Regions and availability zones (AZs) define the fault domain for the network infrastructure.
- Network Interface Cards (NICs) are assigned during instance creation. Network performance (IOPS) is determined by the EC2 instance type.

Public-facing Dell APEX Block Storage components are accessible from the external network if the AWS firewall allows it.

The proxy server is the only component with an external IP address and is used for the "jump server" method of accessing the Dell APEX Block Storage deployment. Public-facing APEX Block Storage components are accessible only through the proxy server, provided that this is allowed by the AWS firewall.

The private subnet components are isolated to ensure security and speed within the Dell APEX Block Storage environment. Private subnet components are as follows:

- SDS (storage)
- SDC (client)
- APEX Block Storage management platform
- SDR (when asynchronous replication is used)
- MDMs and tiebreaker (management)
- CloudLink (if used)

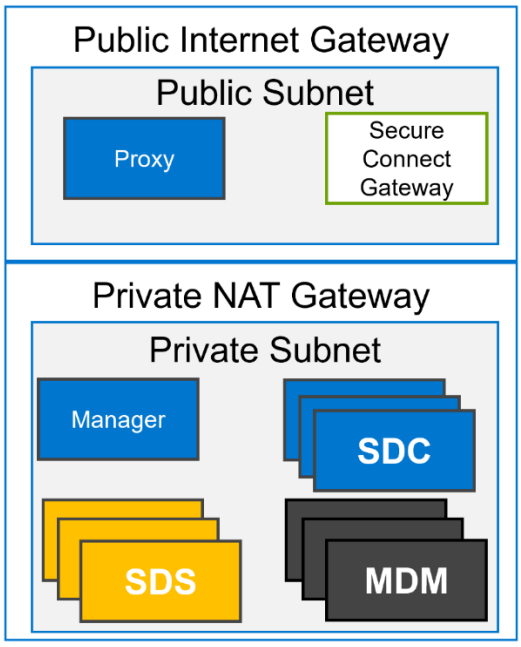


Figure 13. Logical network design

Security groups

Security groups are the core of AWS security. Each interface on an EC2 instance is managed by a security group. You can apply more than one security group to a single instance interface.

APEX Block Storage for AWS supports two types of security groups:

- Public security group - externally available. The proxy server and SCG require external access and are the only components to use the public security group. Instances assigned to the public security group are also assigned external IP addresses.
- Private security group - internal subnets only, with no external IPv4

You must create a security group and policies before creating the ASW instances. AWS firewall controls are available for every security group to further protect any public-facing components from unauthorized access. After the wizard has completed the deployment, the customer is responsible for adjusting the firewall controls for the public security group.

Required ports

The following tables list the ports required for the private ports' inbound and outbound rules that will need to be part of your VPC subnet.

Inbound rules

Table 9. Private subnet inbound rules

Rule Number	Type	Protocol	Port range	Description
1	Custom TCP	TCP	8080	APEX Block Storage - gateway
2	Custom TCP	TCP	8083	SSO internal pod listener

Rule Number	Type	Protocol	Port range	Description
3	Custom TCP	TCP	7072	APEX Block Storage – SDS listener
4	Custom UDP	UDP	51821	RKE2 - Canal CNI with WireGuard IPv6/ dual stack
5	Custom TCP	TCP	7611	APEX Block Storage – MDM peer connection
6	Custom TCP	TCP	61714	ActiveMQ
7	Custom TCP	TCP	8161	ActiveMQ
8	Custom TCP	TCP	9433	Thin deployer
9	Custom UDP	UDP	53	APEX Block Storage – pods connectivity
10	Custom UDP	UDP	4420	APEX Block Storage - core SDT
11	Custom UDP	UDP	9098	APEX Block Storage – SDS listener
12	All ICMP - IPv4	ICMP	All	RKE2 - Cilium CNI health checks
13	Custom TCP	TCP	30000-32767	RKE2 - NodePort port range
14	Custom TCP	TCP	8611	Default mTLS port for MDM
15	Custom TCP	TCP	4789	RKE2 - Calico CNI with VXLAN
17	Custom TCP	TCP	10250	RKE2 - kubelet
18	Custom TCP	TCP	9345	RKE2 agent nodes - Kubernetes API
19	Custom UDP	UDP	5000	RKE2 – Docker registry
20	Custom TCP	TCP	80	APEX Block Storage - gateway
21	Custom TCP	TCP	8009	APEX Block Storage - core SDT
22	Custom TCP	TCP	30400	NodePort access from proxy
23	Custom TCP	TCP	6443	RKE2 – Kubernetes API
24	Custom TCP	TCP	28765	APEX Block Storage
25	Custom TCP	TCP	9402	RKE2 - cert-manager
26	Custom UDP	UDP	51820	RKE2 - Canal CNI with WireGuard IPv4
27	Custom TCP	TCP	5000	RKE2 – Docker registry
28	Custom UDP	UDP	9099	APEX Block Storage – SDS listener
29	Custom UDP	UDP	8472	RKE2 - Cilium CNI VXLAN
30	Custom TCP	TCP	8443	APEX Block Storage - gateway
31	Custom TCP	TCP	6611	APEX Block Storage – MDM listener
32	Custom TCP	TCP	2380	RKE2 - etcd peer port
33	Custom TCP	TCP	443	APEX Block Storage – Allow SSH for EC2

Rule Number	Type	Protocol	Port range	Description
34	Custom UDP	UDP	12200	APEX Block Storage - core SDT
35	Custom TCP	TCP	9099	APEX Block Storage – LIA listener
36	Custom TCP	TCP	179	RKE2 - Calico CNI with BGP
37	Custom TCP	TCP	9098	RKE2 - Calico Typha health checks
38	Custom TCP	TCP	61613	ActiveMQ
39	Custom TCP	TCP	22	Allow SSD for installer
40	Custom UDP	UDP	12200	APEX Block Storage - core SDT
41	Custom TCP	TCP	5473	APEX Block Storage – Calico CNI with Typha
42	Custom TCP	TCP	9011	APEX Block Storage – MDM cluster member
43	Custom TCP	TCP	2379	RKE2 - Cilium CNI health checks
44	Custom TCP	TCP	4240	APEX Block Storage – SDS listener
45	Custom UDP	UDP	7072	APEX Block Storage – SDS listener
46	Custom UDP	UDP	8009	APEX Block Storage - core SDT
47	Custom TCP	TCP	4420	APEX Block Storage - core SDT

Outbound rules

Table 10. Private subnet outbound rules

Rule Number	Type	Protocol	Port range	Description
1	All traffic	All traffic	0.0.0.0/0	Allow external connection

Permissions for managed policies

This section describes the AWS permissions required for a Dell APEX Block Storage deployment into AWS. To deploy, you need the following AWS-managed policy permissions completed as a pre-requisite activity.

The AWS account administrator for your organization will need to make these changes using the AWS Management Console in the **Identity and Access Management (IAM) > Policies** dashboard for each AWS Region in which Dell APEX Block Storage will be deployed.

- AmazonEC2ContainerRegistryFullAccess
- AWSCloudFormationFullAccess
- IAMReadOnlyAccess
- AWSSupportAccess
- AWSLambda_FullAccess
- AmazonSSMReadOnlyAccess

- AWSLambdaRole

Required policy actions

The following is a list of the required policy actions. Ensure that AWS account permissions are adjusted to include these capabilities:

Allow the following actions:

- elasticfilesystem:DeleteTags
- elasticfilesystem:DeleteAccessPoint
- elasticfilesystem:DeleteFileSystemPolicy
- elasticfilesystem:DescribeBackupPolicy
- elasticfilesystem:DescribeFileSystems
- elasticfilesystem:DescribeFileSystemPolicy
- elasticfilesystem:DescribeLifecycleConfiguration
- elasticfilesystem:DescribeMountTargets
- elasticfilesystem:DescribeMountTargetSecurityGroups
- elasticfilesystem:DescribeTags
- elasticfilesystem:DescribeAccessPoint
- elasticfilesystem:ModifyMountTargetSecurityGroups
- elasticfilesystem:PutBackupPolicy
- elasticfilesystem:PutLifecycleConfiguration
- elasticfilesystem:PutFileSystemPolicy
- elasticfilesystem:UpdateFileSystem
- elasticfilesystem:TagResource
- elasticfilesystem:UntagResource
- elasticfilesystem:ListTagsForResource
- elasticfilesystem:Backup
- elasticfilesystem:Restore
- elasticloadbalancing:*
- events:DeleteRule
- events:EnableRule
- events:List*
- events:Put*
- events:RemoveTarget
- events:TestEventPattern
- kms:DescribeKey
- kms:ListAliases

- logs:*
- purchase-orders:*PurchaseOrders
- route53 resolver:*
- network-firewall:List*
- s3:*
- servicelens:*
- ssm:StartSession
- ssm:TerminateSession
- "sts:DecodeAuthorizationMessage
- synthetics:*
- xray:*
- iam:PassRole
- Iam:CreateServiceLinkedRole
  - Iam:AWSServiceName:
    - Autoscaling.amazonaws.com
    - Ec2scheduled.amazonaws.com
    - Elasticfilesystem.amazonaws.com
    - Elasticloadbalancing.amazonaws.com
    - Spot.amazonaws.com
    - Spotfleet.amazonaws.com
    - Transitgateway.amazonaws.com

Conclusion

From both the business and technical uses cases conveyed in detail in this paper, the opportunity for real multicloud innovation is now clear. Dell APEX Block Storage for Public Cloud delivers on combining enterprise-class storage services across the evolving multicloud landscape.

Software-defined, cloud-native storage has fundamentally changed the landscape regarding how storage is managed and consumed. Dell Technologies brings together the best of breed Dell storage with the best of the hyperscalers. Combined now as a single resilient multicloud platform to address any storage need requirement, this solution helps place the correct data in the correct location, to be consumed at the correct time.

References

Dell Technologies documentation

The following Dell Technologies documentation provides other information related to this document. Access to these documents depends on your login credentials. If you do not have access to a document, contact your Dell Technologies representative.

- Dell APEX Storage for Public Cloud

APEX Navigator is aligned with a client beta program. Which also includes a 90-day free trial period. Contact your Dell representative for details.

Additional resources

For more information about our APEX Storage for Public Cloud, visit this page.

- Dell APEX and APEX Navigator playlist