



Distributed Multi-Cloud Data Management Software



Spectra Vail® is a breakthrough cloud data management software solution for the multi-cloud world. Vail unifies all of your data, allowing you to leverage on-premises applications and native cloud services, no matter where your data is created or stored.



Finalist
Cloud Award
2021-2022



DCS AWARDS
FINALIST
2022





Spectra Logic's Integrated Portfolio

Spectra Logic has over 40 years of experience delivering innovative, high-capacity, long-term storage – and with it, a sterling reputation for being known as a trusted storage company. This has given us the foundation to deliver a new generation of data management and storage solutions for a future that is not only cloud, but multi-cloud. And a future where Spectra addresses the challenges many organizations face of increasing internal and external cyberattacks, such as ransomware, that threaten company brands and bottom lines.

Spectra's experience and successful track record in meeting customer needs has led to a portfolio transformation that consists of attack-hardened™ data management and storage solutions for the multi-cloud world.

Whether you are looking to manage data seamlessly across public and private clouds, store data securely for many types of workloads, automate migration to take advantage of the perpetual tier of storage or store data easily on an affordable, scalable platform, or any combination of these, Spectra Logic delivers an innovative portfolio of integrated solutions to help your organization protect, access, secure and harness the power of your data.



What is Vail?



Spectra Vail® is a breakthrough cloud data management software solution for the multi-cloud world. Vail unifies all of your data, allowing you to leverage on-premises applications and native cloud services, no matter where your data is created or stored.

With so many public clouds, so many object storage platforms and so many applications, a solution is needed that makes everything a**VAIL**able in a single managed storage platform without vendor lock-in.

Vail provides a user experience that is built utilizing Amazon® Web Services (AWS®), meaning that users familiar with AWS will have an easy transition to Spectra's Vail software. Vail's multi-synchronization capability manages cloud and on-premises storage repositories so that customers can direct users to either type of storage based on that user's locality and performance requirements, all under a single pane of glass management user interface (UI). Vail enables workflows that leverage the flexibility and agility of cloud services while keeping data locally when needed, allowing egress optimization for lowest cost and fastest access from available storage. For datasets that are on-premises, Vail can synchronize them to public cloud storage enabling public cloud applications to leverage the same data as on-premises applications.

Vail Use Case - University

The Goal:

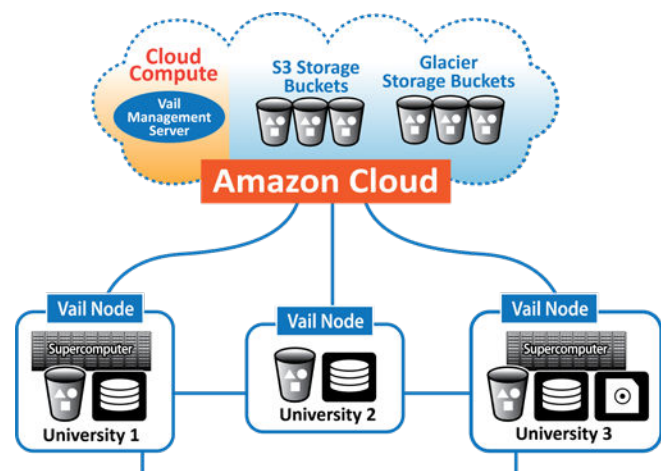
The University has a goal of uniting campuses so they can leverage local storage, but still collaborate and share data with other universities inside their network. Additionally, the local colleges would like to create a disaster recovery (DR) copy of their data in the event that a disaster was to affect their data.

The Challenge:

With their current workflow, sharing and collaborating with other universities was difficult and was only possible through a number of manual data transfers and movements prior to data being accessible. Another challenge was encountered when sharing and distributing data once new findings or results were documented and stored. Lastly there was a concern of becoming locked into any single cloud they chose.

The Solution:

With Vail, this university was able to create site independence with collective data sharing and access across all locations. Through Vail's ability to transfer data directly from site to site without incurring cloud egress fees, the university was able to create a workflow that avoids cloud lock-in by keeping at least one local copy. Additionally, Vail was able to facilitate the process of ingest, research, compute and publication – all integrated into a single system and view of their data



How Does Vail Work?

Vail provides an object storage ecosystem with both on-premises and cloud components. It can run, in its simplest form, as a software-only VM-based on-premises solution. Vail includes:

- Any combination of on-premises and cloud storage
 - Online and nearline cloud storage
 - Fast local online flash or disk arrays
 - Deep bulk online disk storage
 - Local nearline tape storage
- Cloud-based command and control enabling:
 - Unlimited capacity, throughput, object count, user count, site count
- Single global namespace with temporal and location-based policies
 - Any user, any data, any location
- Automatic data management
 - Replication, movement, tiering, control, protection, lifecycle management
- Cloud-based portal with simple access, monitoring and control
- Unlimited sharing inside and outside the organization
- Full integration with cloud-based workflows and cloud services, including analysis and artificial intelligence processing
- Public hybrid cloud – a central way to manage everything



There are many ways to extract value from your data. With more public cloud providers delivering overlapping and complementary capabilities, IT consumers want data to be synchronized and available in multiple locations. Vail extends multiple public and private clouds into a single managed structure so that data can be leveraged where it has the most value and at the optimal cost.

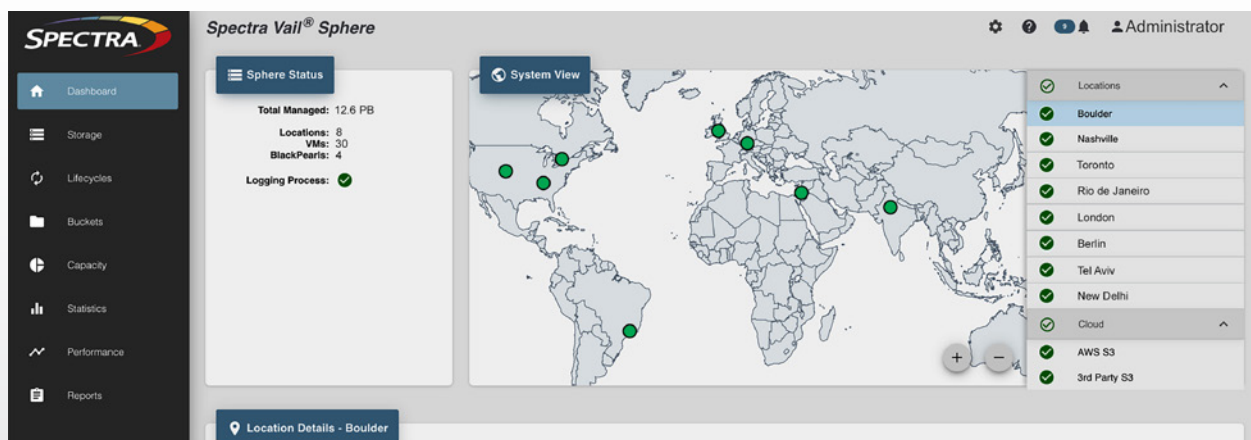
Vail is an object storage solution with a standard cloud S3 interface, which means it will work with existing cloud applications for moving data into and out of Vail's storage. An application that uses cloud or on-premises object storage can be easily redirected to use Vail. After doing so, the customer will experience the benefits of the Vail solution, including higher performance and more control over data, at a lower cost. Vail appliances located at different sites have access to all the same data buckets and underlying objects as they have had previously. Buckets can be mirrored from one site to another for data availability and performance. This ensures that a request for a data object is always fulfilled from the highest performance copy. Unlike traditional storage systems that perform synchronization across a small number of sites, the Vail solution can support an unlimited number of sites.

Central Management and Access

A Vail Sphere has a central point of management provided through a browser-capable UI. This single point of management allows the end user to set up a Vail Sphere with sites, storage targets, and Vail bucket lifecycle policies; observe the state of an entire Vail Sphere; and create and delete users. This will offer a common user experience and familiar access for AWS cloud users.

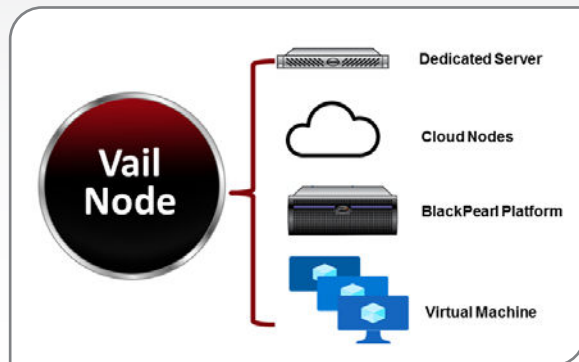
The Vail Sphere runs in its own separate cloud account or sub account such that billing can be isolated. Organizations can choose to have Spectra host and operate the sphere, or they can host and operate it in their own account.

A central user interface provides a real-time overview of a Vail Sphere including system health in each location, overall status of storage and information by site.



Actual Spectra Vail User Interface

Building Blocks of a Vail Solution:



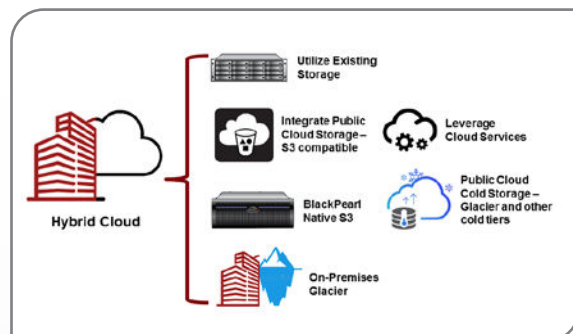
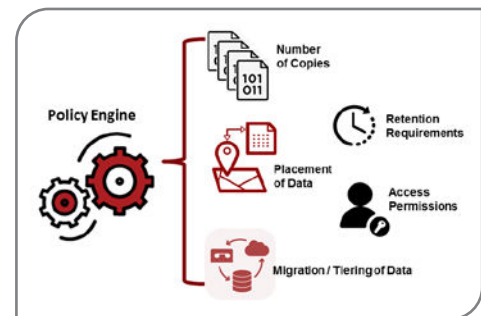
Connection Types (Vail Nodes)

A Vail node is a software package that can be deployed on a number of options:

- A virtual machine or on a Spectra® BlackPearl® storage platform
- Existing storage
- Cloud storage – AWS, **GCP, **Azure
- Spectra storage
- Local glacier* tier (tape behind BlackPearl)

Lifecycle Policies Drive Vail Solutions

A lifecycle policy in a Vail solution consists of rules that dictate where objects are stored and the length of time they are stored in each storage location. Each lifecycle policy is created by the administrator to fit the desired workflow and retention periods. Users can determine the placement of objects in the bucket by specifying the number of copies the rule will keep and the specific storage pools to use. The available storage targets consist of public cloud storage tiers and local storage under control of a Vail node that has been integrated into a Vail Sphere.



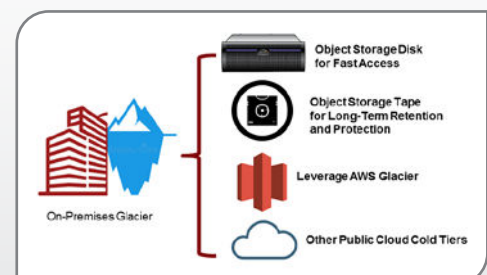
Hybrid Cloud Integrations

Vail presents both on-premises and cloud-based storage options within the lifecycle policies. Copies can be made immediately or after, in any location, to any storage target or medium, and presented to the user as one simple storage system that reacts to their immediate needs.

From storage to compute, Vail integrates with the major cloud providers and any cloud using S3 protocols. This eliminates cloud lock-in and provides the ability to leverage the cloud and cloud services that fit the desired workflow.

Glacier Storage for Long-Term Retention

Creating a local glacier repository is now possible with the help of Vail. Glacier storage delivers substantial cost savings and infinite scalability. The benefits of a Vail solution cannot be overstated. The fact that long-life object storage disk and literally unlimited tape capacity can be realized, makes your own glacier storage tier an easily attainable reality. The seamless combination of affordable disk and ultra-cost-efficient tape, means that any company's growing data universe can be organized and managed in a whole new way – and without any cloud lock-in limitations.

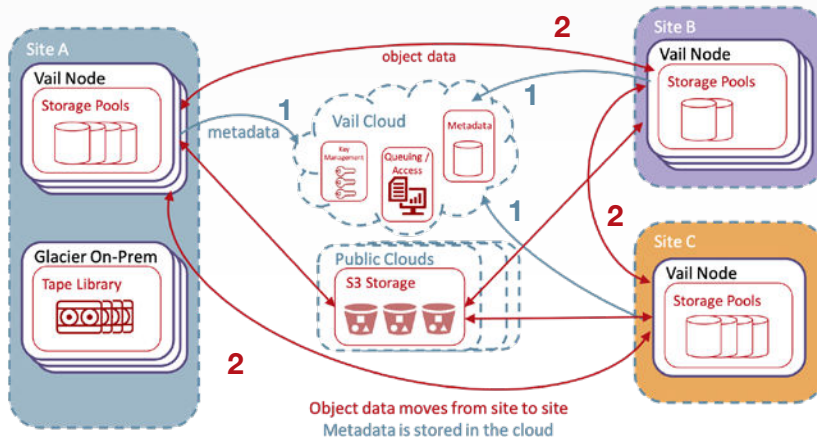


**future release

Data Flow vs. Metadata Flow

One major benefit of Vail is that data does not flow into and out of the public cloud unless a cloud bucket is specified in the policy as a storage target. If a Vail policy dictates that copies are to be stored and accessible at two different physical sites, then the data movers in each Vail node transfer data directly between those sites and not through the public cloud. This prevents excessive egress charges from cloud vendors and removes an additional network hop.

If a policy requires a copy on a cloud bucket as well, data is always copied directly between nodes and to the cloud bucket. A cloud bucket is never used as a source, and that is the only copy of data available.



1. Metadata is always stored and controlled in the cloud. Transactions require multiple interactions with the cloud management server including metadata access, location recording, loading of the job queue, etc. Database transactions, however, are low bandwidth, low cost, and can be fully encrypted. Metadata is also locally cached to maximize performance.

2. Through this architecture, Vail delivers a single management view into all data regardless of its physical location. Data is transferred from site to site, saving on egress fees and cloud storage charges. The transactional, low-capacity metadata is stored and accessed from the public cloud, **saving up to 80%** over a traditional pure cloud workflow.

The Goal:

To find a way to collect globally dispersed data into a central repository where a supercomputer can perform analytics and generate results. Once results are generated, they must be shared with other departments and users.

The Challenge:

Remote offices and remote collection sites were unable to connect to a central location to transfer data, causing delays and other challenges of transferring data. Once results were generated, the public cloud was used to distribute the data. This resulted in unplanned egress and out-charges with the public cloud and alternatives needed to be evaluated.

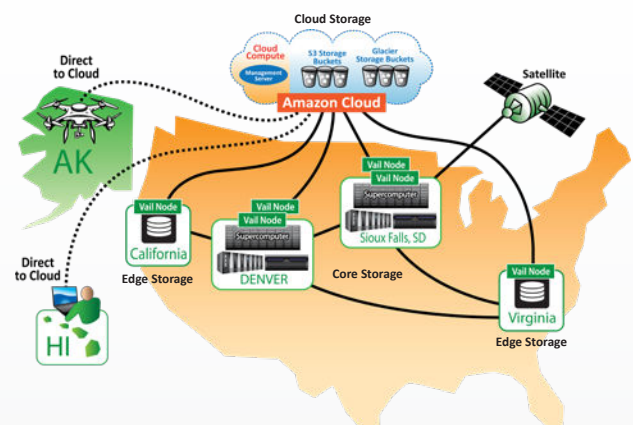
The Solution:

With Vail, this government agency was able to implement an end-to-end data management and delivery service that allows data collection from instruments and contributors from all over the world using a combination of edge, core and cloud computing locations. In addition, they can store data in AWS and on on-premises according to a policy defined to meet service levels that match the value of the data over time, while managing access costs to adhere to the needs of the data consumers and their budget resources.

Lastly, Vail is able to keep historic data in a multi-site, durable, accessible, and affordable datastore. Any dataset can now be accessed in minutes while only paying storage costs.

With Vail, this agency is able to leverage any combination of public clouds and on-premises datastores. They are able to choose the storage tier that matches the value of their data and response time needed for any point in time access to the data. This happens while managing the data based on the elasticity and reliability of AWS cloud services.

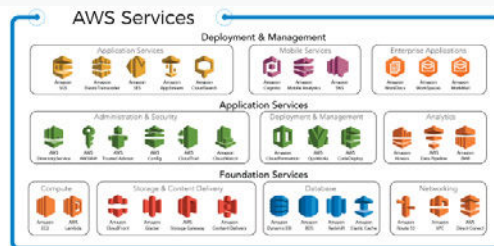
Vail Use Case - Government Agency



Leveraging Cloud Services with Vail

Modern data centers are moving toward a hybrid model where some processes are to be performed on-premises and others in the public cloud. Previously, storage was tightly coupled with where the processing was performed, meaning that where the compute goes, the storage will follow. Through its patent-pending technology called ‘bucket synchronization’, Spectra Vail has broken the bond that has traditionally existed between data processing and data locality. This frees users to store data where it makes the most sense, independent of where processing is performed. In this manner the customer may utilize cloud processing services, such as AWS media services or Google AI services, while not being required to store the data in those clouds long term.

To demonstrate this capability, consider a customer that wants to utilize AWS media services for transforming raw high-resolution video images into their lower resolution counterparts. Furthermore, they would like to archive all high-resolution and related low-resolution files to an on-premises tape system. Vail’s lifecycle policy manager allows for setting up multi-directional synchronization between a Vail bucket to a public cloud bucket, for example, AWS S3 storage.



One cloud workflow example is where all high-resolution images are ingested into a Vail bucket. These raw images are then synched to an AWS S3 bucket that AWS media services are set up to use as its input. When objects are downloaded into this bucket, AWS media services automatically create low-resolution versions of the content and write it out to the same bucket. The Vail bucket multi-directional synchronization feature becomes aware of these new objects and copies them back into a pre-determined Vail bucket. Furthermore, the Vail lifecycle policy will delete the high-resolution objects in the AWS S3 bucket. This can be done immediately or as a time-based policy to reduce any early deletion charges incurred from some cloud providers. At a different point in time, typically at a project’s completion, all objects associated with the project can be transferred to tape for long-term digital preservation.

Vail Use Case - Post Production Studio

The Goal:

This post production organization was looking to leverage the public cloud in their workflow. The desire was to utilize cloud services, such as transcoding, as well as store a long-term preservation copy of their finished work.

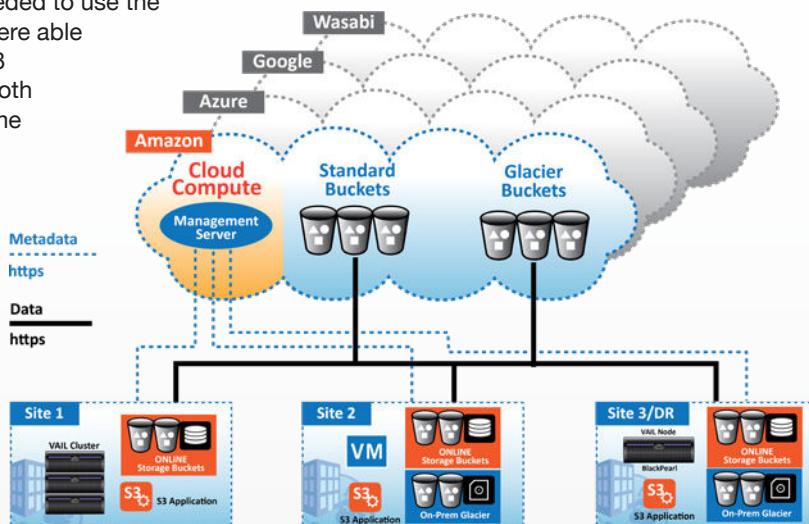
The Challenge:

The primary challenge was how to efficiently and affordably get their content into the cloud to begin utilizing cloud services. Additionally, cloud storage and egress fees became a big concern.

The Solution:

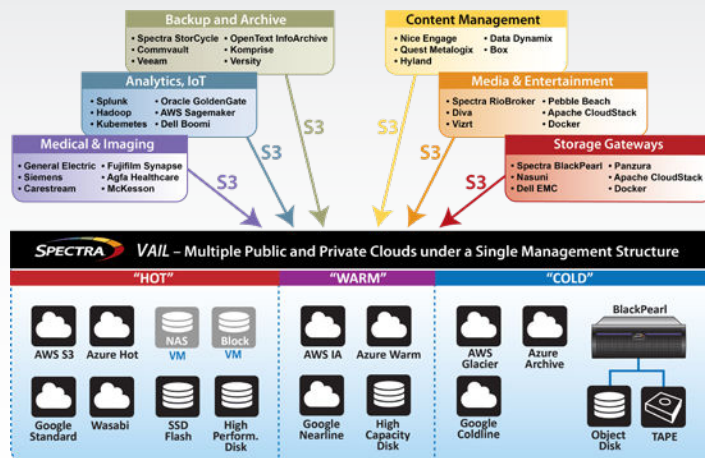
Vail enabled this organization to extend their data center into the public cloud by managing object storage in multiple locations under a single, policy-based management structure. Customers wanted to run applications on one or more public clouds with production data, but still needed to use the production data on-premises. With Vail, they were able to synchronize object data to any number of S3 storage pools. This allowed them to leverage both on-premises and cloud applications for the same data without negative impacts to their SLAs.

At the same time, they kept their cloud costs under control by populating data to the public cloud when it was needed and deleting it when it was not, without incurring egress charges. The migration to the cloud didn’t occur in one day. During the migration to the cloud, Vail kept the on-premises and cloud storage synchronized so that the content was usable in both locations simultaneously.



Vail Ecosystem and Application Connections

Vail enables IT organizations to store object data in any number of cloud locations – whether in a public cloud, on-premises location or both. Vail can join public cloud storage and on-premises storage to become a single namespace of object storage. Data is managed by user-defined policies to optimize service level agreements and costs to match the value and location of the data. Customers can synchronize any S3 source with any number of S3 targets to make data available to any application in any cloud. Vail delivers the flexibility to balance data location, performance and cost in a single managed, unified cloud operating environment.



The Goal:

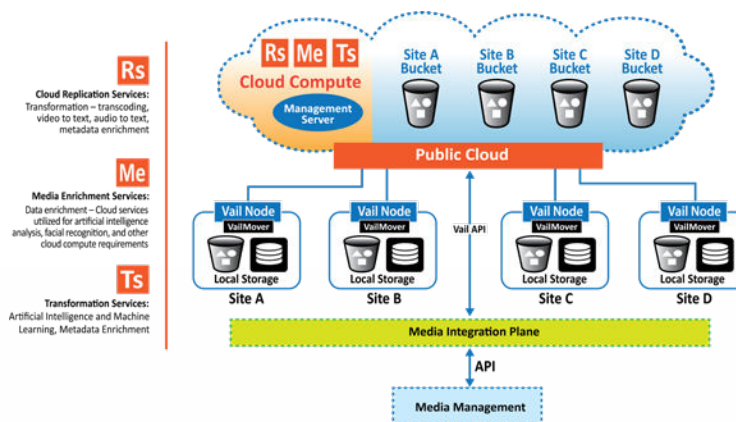
To create an automated workflow that seamlessly utilizes cloud services such as transcoding, edit, playout and metadata stripping. Additionally, to create a common platform for all assets, making them available from anywhere, to anyone.

The Challenge:

With over 100 sites spread across the nation, and with every station able to choose tools that work best for their station, a challenge of how to store and manage content across sites was discovered. Cloud services and storage costs both began to rise, becoming cost multipliers in the overall workflow.

The Solution:

With Vail, this broadcast group was able to deliver a single media manager across all sites, accessed by thousands of users, in a future-proof S3 interface. Cloud-based workflow automation was achieved by storing high res and raw video content locally, reducing costs. Vail was able to manage media and metadata (sidecar) as objects within a single namespace. Vail also automated site placement of data based on tags with staging, auto caching, network optimization – and direct site-to-site transfers. By only using AWS for cloud services and keeping HD and raw content locally, the cloud bill was reduced by over 70%.



Enterprise-Class Security



Data security is one of Vail's core principles and has been designed to ensure that data is protected, regardless of its physical location. By leveraging the inherent air gap present in a tape system, hackers cannot gain physical access to local glacier data copies and tape can prevent overwrites and changes typical of ransomware attacks. Versioning also provides an added and instantly retrievable copy option. Vail protects data when it is stored and transferred, as well as the database with all the valuable metadata. With a Vail solution, peace of mind is at hand, with users knowing data will be safe and protected. Vail protects data through the following methods:

- Data can be encrypted, single key per object, locked to your account
- Database encrypted, metadata, object names and locations
- Endpoints hidden behind firewalls
- Object locking** for ransomware protection
- HTTPS, V4 authentication for in-flight protection
- Self-healing and versioning for data corruption/malicious protection
- Multi-site for physical security
- Cloud/storage mobility to protect against cloud lock-in
- Multi-level CRC hashing to detect and correct errors



Differences Between Vail and BlackPearl S3

In addition to the Vail software solutions, Spectra has an appliance-based version called BlackPearl S3. BlackPearl is a multi-purpose hybrid storage platform that easily and cost effectively scales up as object and file volumes grow; adapts as operational requirements change; and enables easy synchronization of data between on-premises and cloud storage. With its S3 extension, the BlackPearl appliance can be a local S3 storage system with the ability to integrate with the public cloud. For single cloud and low number of sites, BlackPearl S3 may be the best option, whereas organizations looking to leverage multiple clouds and multiple sites would benefit from the Vail software.

	 Vail	 BLACKPEARL S3 Storage
Deployment	Software system that combines on-premises and multi-cloud storage into a single S3 namespace	Appliance-based (BlackPearl) local S3 storage with cloud/tape tiering options.
Storage	Broad spectrum of on-premises storage	Local S3 standard as BlackPearl/ZFS pool – DR or gateway to AWS and/or tape
	Any - VM, BlackPearl, cloud node	No VMs - on BlackPearl only – Appliance
	BlackPearl, VM, block, local S3, public cloud	No 3rd party storage (no block/LUN, other local S3 etc.)
	Multi-Cloud / Multi-Site	Single Cloud / Single Site
Failover	Cloud or local nodes for failover	Optional enhancement for Cloud DR "node" for failover
Notifications	Notifications allowed – active S3 content and sync plus Multi-directional sync	No notifications (linked buckets) Bi-directional synch Optional enhancement to enable linked buckets for BlackPearl S3
Multi-Site Options	Multiple nodes of any type, VM, BlackPearl, other	Multiple BlackPearls allowed in a BlackPearl S3 solution (only BlackPearls)
Description	This software is primarily used to orchestrate data, manage data placement and access, and acts as a data mover with intelligent policy management built in. Vail provides a single global namespace to manage all clouds, on-premises storage, and application data into a single, managed namespace that can be accessed from anywhere.	Storage device that creates a synchronous link to cloud/tape. Optionally, customer can set up a lifecycle rule using move commands that allows immediate or delayed move (copy/delete) and allows BlackPearl S3 to work as a gateway enabling more data in cloud/tape than on local disk.

Spectra World-Class Support

Industry-Leading Support to Keep You Safe



The SpectraGuard Difference – Customer Focused. Results Driven.

Spectra Logic's SpectraGuard support takes a personal and attentive approach to every support interaction regardless of customer size or severity level. It's what sets Spectra apart and will earn your highest satisfaction. We provide the correct tools, expertise and resources to drive your issues to resolution. We offer a comprehensive set of support options spanning a wide range of business requirements, from simple phone support and parts replacement, to custom add-ons such as stocked parts at customer sites and focused service account management.

A Better Approach to Support

24x7 Global Support – Our technical experts and management teams are available around the clock to ensure you are receiving the support you can rely on. With an industry-leading time-to-answer average of 30 seconds, and an unprecedented 98% incident resolution rate at Tier 1, you can count on Spectra support to assist you. Our centralized location at our corporate headquarters in Boulder, Colorado is complemented by our European and Asian support teams to provide global time-zone coverage.



Unparalleled Customer Satisfaction – As customer satisfaction is one of our core values, we are dedicated to ensuring you receive the best support out there. Our customers validate their positive experience by consistently providing us with near-perfect ratings on routine satisfaction surveys.

Professional Services – Quickly implemented and according to specification, these services allow a rapid return on your storage investment from the time of first contact to the end of the product lifecycle. Services include: Installation and Integration, Media Migration, Preventative Maintenance, Customized Training, Data Center Services, and many more.

Assisted Self Maintenance – For those organizations that require immediate parts availability, this industry-first support option provides faster resolution time with onsite, customer replaceable parts. Secure sites, remote locations, and companies with hands-on oriented technical staff will find this supplemental support ideal to meet their unique requirements.

AutoSupport Phone Home – When a critical error on the library occurs, BlueScale® AutoSupport notifies our support team, sends error logs via email, and automatically opens an incident on your behalf. The support process is initiated without you even making a phone call, and because we understand the importance of response and recovery times, this benefit is part of every support contract – at no additional cost.



Technical Account Management – This personalized support service provides a single point of contact at Spectra. Familiar with the intricacies of your data storage environment, your personal advocate is accountable for managing and overseeing all technical support activities for your organization. This premium service saves you effort, involvement and time, allowing you to focus on achieving your core business objectives.

Storage Crisis Lifeline – A disaster recovery program, Storage Crisis Lifeline provides customers with a complimentary loaner library and installation services in the event a disaster situation damages or destroys the original library to the point that data cannot be restored.



SPECTRA'S DATA STORAGE STORY

Spectra Logic develops a full range of Attack-Hardened™ data management and data storage solutions for a multi-cloud world. Dedicated solely to data storage innovation for more than 40 years, Spectra Logic helps organizations modernize their IT infrastructures and protect and preserve their data with a broad portfolio of solutions that enable them to manage, migrate, store and preserve business data long-term, along with features to make them ransomware resilient – whether on-premises, in a single cloud, across multiple clouds, or in all locations at once.

IT Challenges That Spectra's Vail Software Can Solve

Vail enables organizations to store object data in any number of locations – whether in a public cloud, on-premises or both. Vail extends your public cloud storage and on-premises storage to become a single platform of object storage locations managed by user-defined policies to optimize service level agreements and costs to match the value and location of the data. Organizations can synchronize any S3 source with any number of S3 targets to make data available to any application in any cloud. Vail delivers the flexibility to balance data location, performance and cost in a single managed, unified cloud operating environment.

Challenge	How Vail can solve the problem
How do I transition to the public cloud while still utilizing my on-premises storage during the transition?	Vail provides a single namespace to all storage it manages regardless of physical location, including: flash, disk, tape, and public clouds, creating a connected infrastructure that is able to store data locally or in the cloud.
How do I minimize cloud storage retention and egress fees?	With Vail you can leverage what the cloud does best – automated services and disaster recovery – while keeping data onsite for lower cost. You can even use Vail to create your own on-premises glacier storage tier for greatest cost savings.
How do I connect or transition my data between multiple cloud storage vendors?	With Vail, cloud lock-in is no longer an issue when adopting a cloud workflow. By leveraging multiple public clouds, Vail delivers freedom to choose the cloud (and cloud services) that fits your needs. Choose the clouds that work for you. Painlessly move data between cloud vendors without your applications even being aware that the data has moved.
How can I enjoy public cloud benefits if my organization cannot store data in the public cloud?	For organizations that are unable to utilize the public cloud for any number of reasons, Vail delivers a cloud interface and management experience, making all data accessible from anywhere in the world, all while using on-premises storage and maintaining control over your data.
How can I leverage my existing investment in storage technology when moving to the cloud?	With Vail's storage-agnostic approach, existing hardware and software can be utilized in a Vail workflow providing investment protection and utilizing existing equipment, and in turn offsetting cloud costs with local storage.
How do I connect my multi-site environment to enable sharing and collaboration easily?	Vail provides the ability to manage data regardless of physical location and bucket-level policies can be applied to data to determine where the content should reside, where it should be transferred and with whom it should be shared.
How can I remain flexible in a world moving to the cloud?	Vail provides value-add to solutions with a cloud-like workflow by providing a hybrid system that highlights on-premises value while enabling cloud services and storage.
How can I manage opex, capex, and surprise costs associated with storing and accessing all of my data?	With Vail's licensing flexibility, options are available to control costs and allocate budgets to balance the upfront expense with ongoing expenses, reducing surprise costs traditionally associated with cloud workflows.
How can I use cloud services without keeping all my data, long-term in the cloud?	Vail allows for raw content and high-resolution content to be kept locally, while cloud service data output and lower-resolution proxy files are used in the cloud. Once the final cut list is created, high-res and raw content can be accessed from local storage, saving time and money.



Spectra World Headquarters

Toll Free: 800-833-1132 • 303-449-6400 • 6285 Lookout Road • Boulder, CO 80301 USA
International offices in Bracknell, United Kingdom and Sydney, Australia

SPECTRALOGIC.COM

V1-091322

© 2022 Spectra Logic. All rights reserved worldwide. All other trademarks and registered trademarks are the property of their respective owners.

*Amazon Glacier® is a registered trademark of Amazon Technologies, Inc.