InfiniStor™ — Sovereign, S3-Compatible Object Storage for Aotearoa New Zealand

White Paper — Technical & Architectural Overview

Version 1.4 • ASI Solutions • New Zealand

Executive summary

InfiniStor is ASI's sovereign, carbon-zero, S3-compatible object storage platform designed for New Zealand organisations that need durable, secure and cost-predictable storage without bill-shock. It delivers two tiers—Hot (instant access) and Cold (archive with recall via S3 Glacier-compatible APIs)—with up to 19×9's durability, tri-site protection across three NZ data centres, strong S3 consistency, and built-in security (immutability, versioning, lifecycle, encryption). Our all-inclusive pricing model eliminates egress and API fees, making regular restore tests and audits practical.

1. Problem landscape

Unpredictable costs in the public cloud, such as egress and API charges, can penalise data restores and limit mobility, discouraging good practices like routine disaster recovery drills. At the same time, cyber-resilience

expectations are rising, with insurers and regulators increasingly requiring immutable backups, enforced retention locks, and clear evidence of recovery testing.

Organisations across councils, education, research, and other regulated industries also face sovereignty and sustainability pressures, demanding data residency within New Zealand and a transparent environmental approach.

InfiniStor meets these challenges by providing sovereign hosting with strong consistency, object-level immutability, and multi-site resilience. Its pricing model is designed to simplify budgeting, helping organizations maintain compliance and resilience without the burden of unpredictable cloud costs.

2. Platform overview

Our storage platform offers two tiers: a Hot tier, which provides always-online S3 access, and a Cold tier, which leverages S3 Glacier-compatible archival storage with convenient recall capabilities. It delivers exceptional durability of up to 19 nines, with geo-dispersed protection across three New Zealand sites, alongside background integrity scrubbing and self-healing to maintain data reliability. Strong security is built in, including Object Lock in Governance and Compliance modes, versioning, flexible lifecycle policies, AES-256 encryption at rest, and TLS for in-flight protection.

Designed with robust S3 semantics, the system ensures read-after-write and list consistency, while offering an intuitive portal for managing tenants, users, buckets, quotas, lifecycle workflows, and monitoring needs. From an economic perspective, it provides predictable per-terabyte pricing with no egress or API request fees, giving customers both transparency and cost efficiency.

3. Architecture deep-dive

3.1 Components and data flow

Access services (stateless) terminate TLS and authenticate S3 requests; operations are distributed across nodes for scale-out throughput. A highly available metadata/control plane coordinates bucket namespace, versioning, Object Lock state and lifecycle events, and enforces strong consistency for PUT/GET/LIST semantics. Storage nodes apply erasure coding and multi-site placement policies. Background services handle integrity scrubbing, dynamic repair and online rebalancing.

3.2 Placement & durability policies (multi-site)

InfiniStor supports policies optimised for capacity, availability, or balance. For large objects in a three-site deployment, examples include:

- 18+7 (balance) — ≈61% usable of raw (≈1.64× overhead); read-available with one site down plus an extra drive failure.
- 18+8 (availability-optimized) — ≈54% usable of raw (≈1.86× overhead); read-available with one site down plus two additional drive failures.

For small objects (≤ ~512 KiB threshold), schemes such as 9+7 or 10+8 are common; pack small files or use multipart to reduce overheads.

3.3 Consistency model

Strong S3 consistency means newly written/overwritten objects are immediately readable and visible in listings—important for backup verification, incremental scans and application correctness.

3.4 Performance model

Throughput scales with client concurrency and multipart uploads. Target ≥1–8 MiB part sizes for sustained ingest, bundle very small files before upload. Endpoints are exposed over HTTPS:443; private connectivity is available for deterministic performance. Capacity and bandwidth scale by adding nodes.

4. Security, compliance, sovereignty

Object Lock (Write Once, Read Many) should be enabled at bucket creation, which automatically turns on versioning, and you can choose between Governance or Compliance modes. Encryption uses AES-256 for data at rest, and TLS secures data in transit. Access control is managed per bucket with credentials supporting least-privilege policies. All data, metadata, and support remain on-shore in New Zealand within carbon-neutral facilities, ensuring sovereign and sustainable operations.

5. Integration Partners

5.1 Backup with Veeam® (Hot + Cold + immutability + Direct-to-Archive)

- Add an S3-compatible object repository pointing at the InfiniStor endpoint.
- Select bucket/folder.
- Enable ‘Make backups immutable’ on Object-Lock buckets.
- For archive, use S3-compatible with Data Archiving (S3 Glacier-compatible) and size the archiver appliance per guidance.
- Avoid bucket-side retention/lifecycle when InfiniStor integrates with your backup platform's retention.

5.2 Backup & archive with Commvault®

Add InfiniStor as an S3 Cloud Library; enable WORM/immutability; place short-lived copies on Hot and long-term retention on Cold; let Commvault own retention.

5.3 Other S3-compatible tools

Analytics/data platforms write to Hot; archival tools using S3 Glacier-style recalls integrate with Cold.

6. Operations & monitoring

- Integrity scrubbing and self-healing are automatic; monitor for rebuild events.

- Institute monthly sampled restores and quarterly full-scale DR rehearsals.
- Rotate keys.
- Maintain runbooks for lost credentials, recalls, expansion and incidents.

7. Reference architectures

Cyber-resilient backups (Hot) are maintained by keeping an on-premises primary copy along with an immutable copy on InfiniStor Hot. This setup enables routine restore testing without incurring any egress costs, ensuring that data can be quickly recovered whenever needed.

For long-term archives (Cold), data follows a lifecycle transition from Hot to Cold storage. These archives can be recalled through S3 Glacier-compatible APIs within the configured retrieval window, allowing for efficient and predictable access to historical data.

Research and analytics workflows (Land to Hot) begin by landing data in Hot storage for immediate use and high performance. As datasets age, they transition to Cold storage, preserving capacity and controlling costs. Private connectivity ensures that sovereignty requirements are met while maintaining consistent performance across the data lifecycle.

8. FAQs

Q1: Do listings reflect new objects immediately?

Yes. InfiniStor leverages strong Amazon S3–style consistency, which ensures that any object written to the system is immediately visible in subsequent read and list operations. This means that after you upload a new object, you can reliably perform a read-after-write request or a bucket listing, and the object will appear without any delays. There is no eventual consistency lag for listings or reads.

Q2: Can I disable Object Lock later?

No. Object Lock must be enabled at the time you create the bucket, and it cannot be turned off afterward. This is because Object Lock is designed to enforce data immutability and regulatory compliance. Buckets with Object Lock automatically have versioning enabled, and every new object version will inherit the lock properties. Once enabled, this ensures that objects cannot be unintentionally overwritten or deleted until the retention period expires.

Q3: Which region do I use in tools?

Although S3-compatible clients often require you to specify a region, InfiniStor does not use the region setting internally. Any region value you provide will be accepted by the client SDKs purely for compatibility, but InfiniStor endpoints determine where your data resides. Therefore, always connect to the provided endpoint URL rather than relying on a region for routing or storage placement.

Q4: What about small files?

Small objects can be stored, but they may be less efficient. Each object carries metadata and parity overhead for durability, and with very small files, that overhead is proportionally higher. To optimize storage and performance, it is recommended to either pack multiple small files together into larger objects or use multipart uploads to group data efficiently. This approach minimizes storage inefficiency and can improve both upload and retrieval speeds.

Q5: Do I pay to restore?

No. InfiniStor does not charge egress or API fees for restoration requests. You can perform disaster recovery (DR) testing, restore objects, or validate your backups as frequently as needed without incurring additional costs. This makes routine DR exercises financially feasible and encourages best practices for verifying your stored data.

Appendix A — Installation & Configuration Guide

Veeam InfiniStor Integration Whitepaper

Note on Cold Storage

The InfiniStor Cold Storage can only be utilised through a Scale-out Repository when Archive Tier is configured.

There are the only supported backup types that can be offloaded to the Archive Tier (https://helpcenter.veeam.com/docs/backup/vsphere/archive_tier.html?ver=120):

- Backup files with GFS flags.
- Veeam ZIP backup files.
- Exported backup files.
- Orphaned backups with GFS flags.
- Backups created by Veeam Backup for Nutanix AHV.
- Backups created by Veeam Backup for Oracle Linux Virtualization Manager and Red Hat Virtualization.
- Backups exported by Kasten policies.

Performance Tweaks for InfiniStor

- Configure the SOBR archive tier storage settings to not use 'Store archived backups as standalone fulls' and 'Archive backups only if the remaining retention time is above minimal storage period.

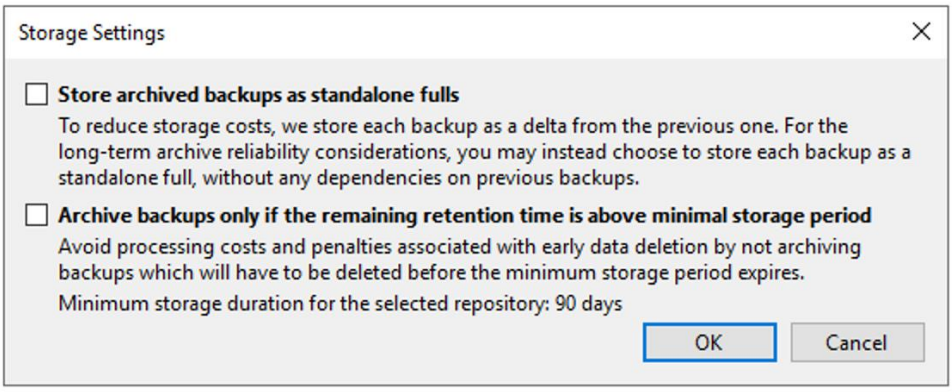

- Configure your backup jobs to use optimal compression and a 4 MB block size.

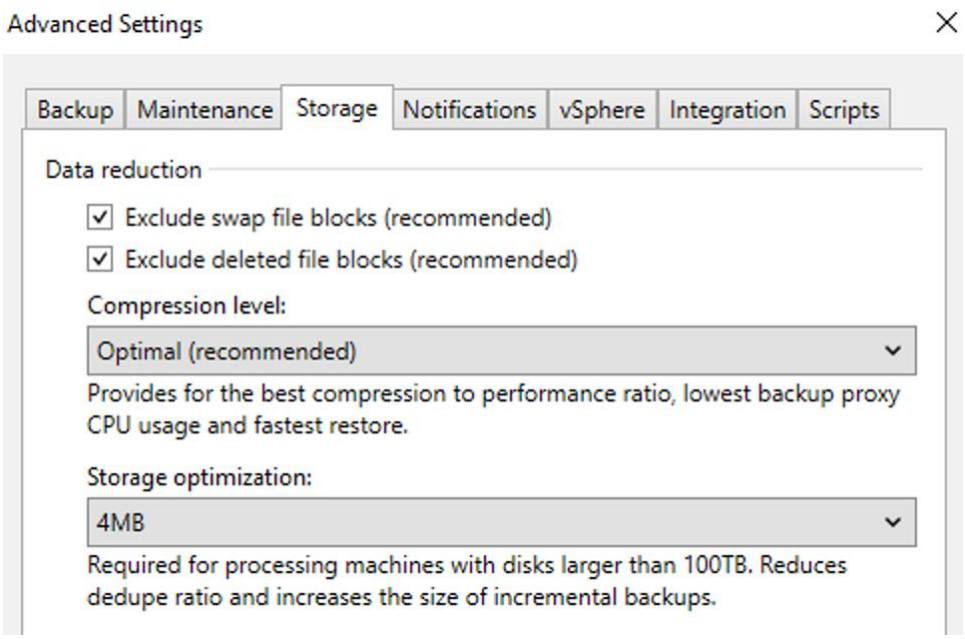

Configuring InfiniStor S3 Buckets

When adding your object storage, you can choose to use separate buckets for your hot and cold storage or use a single bucket with different folders (prefixes) to store your hot and cold data. In this example we will be using a single bucket with different folders.

Adding InfiniStor Hot Storage

- Add a Backup Repository
- Select Object storage > S3 Compatible > S3 Compatible
- Give your repository a name, e.g. InfiniStor-Hot-01
- Limit concurrent tasks should be set once you have completed a sizing exercise

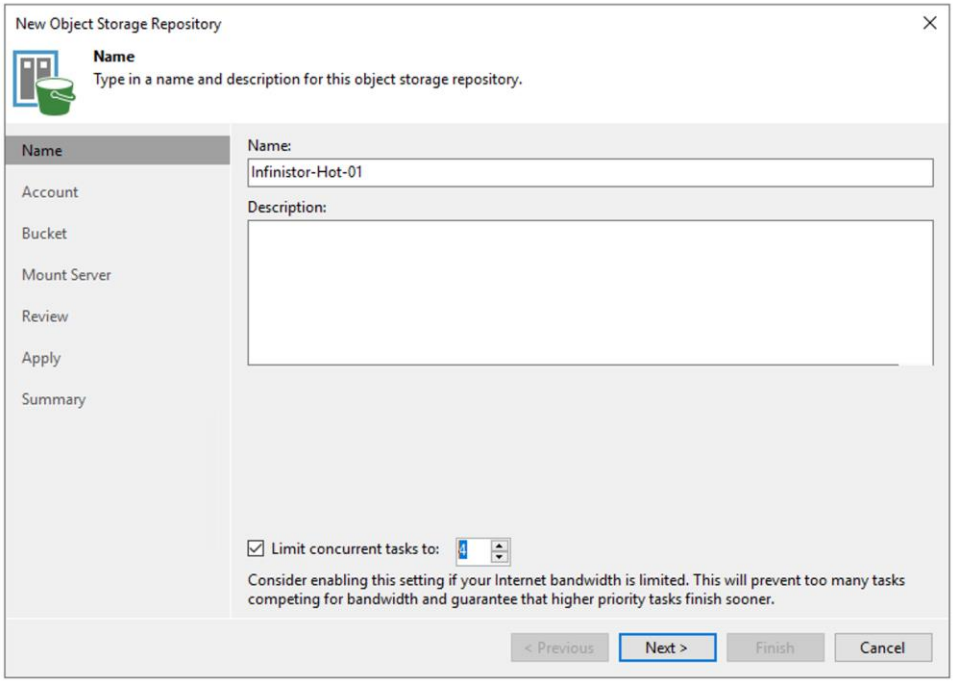

- Fill out the Account details as follows:
    - Service point: https://asi.s3.prod.nz
    - Region: Leave as default us-east-1 (this does not matter with InfiniStor, it is NZ based)
    - Credentials: Add your access/secret key and select the credential.

- Connection mode: Configure as required
  https://helpcenter.veeam.com/docs/backup/vsphere/compatible_repository_account.html?ver=120

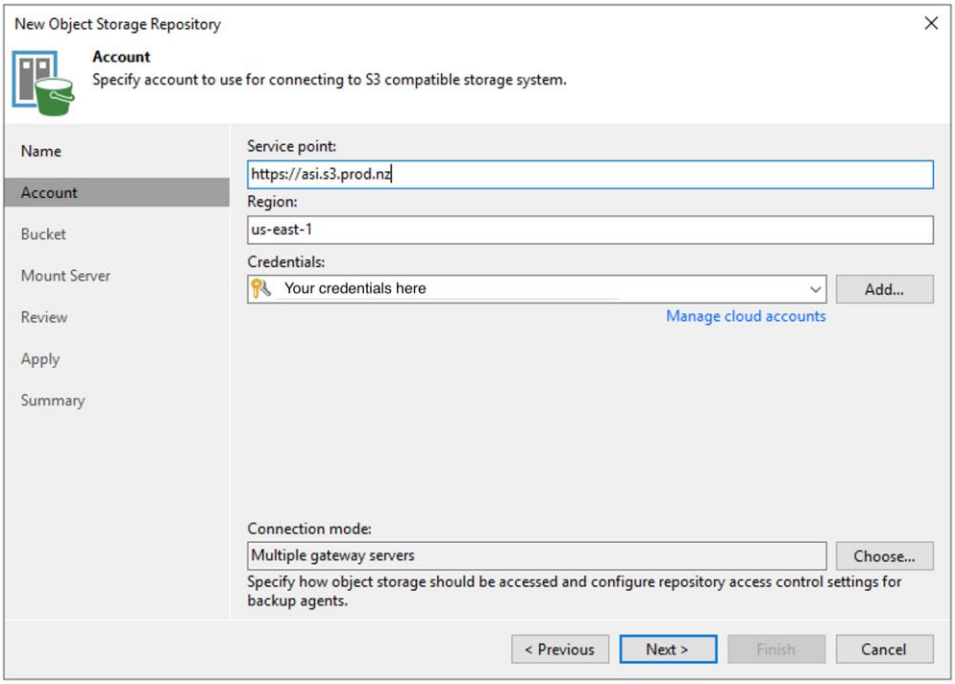

- Browse for your bucket and select it

- Click on ‘Automatic bucket creation enabled’ and uncheck ‘Create new buckets automatically (recommended)’
- Create a folder for Veeam to place the data in, e.g. hot-01
- Configure immutability if required

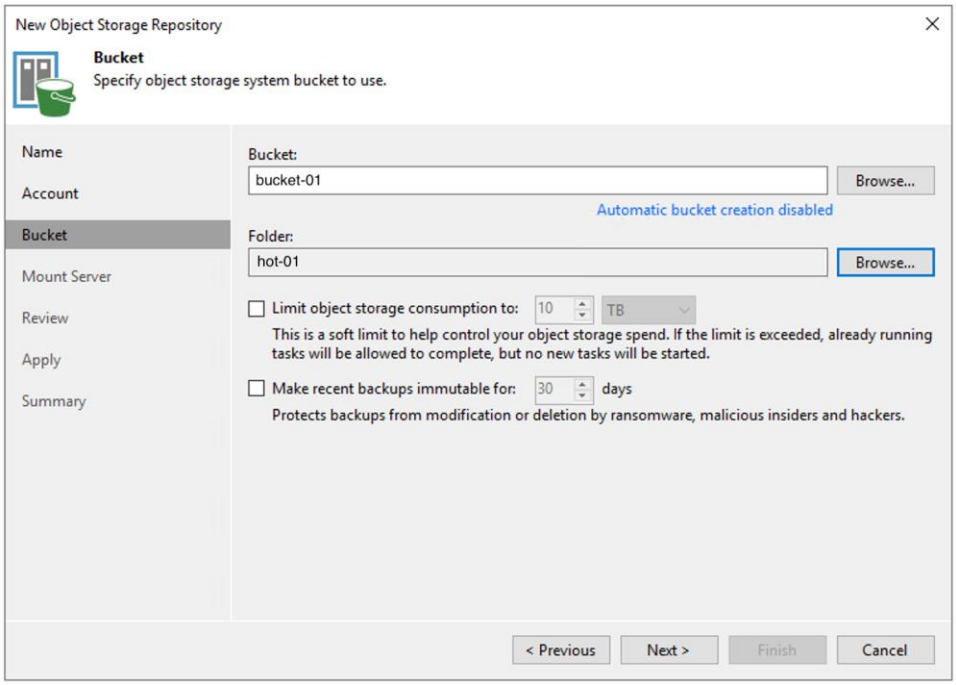

- Configure Mount Server as required
- Review and apply the settings to create your S3 Compatible backup repository

Adding InfiniStor Cold Storage

- Add a Backup Repository
- Select Object storage > S3 Compatible > S3 Compatible with Data Archiving
- Give your repository a name
- Fill out the Account details as follows:

    - Service point: https://asi.s3.prod.nz
    - Region: Leave as default us-east-1 (this does not matter with InfiniStor, it is NZ based)
    - Credentials: Add your access/secret key and select the credential or re-use the credential you added in the previous InfiniStor Hot Storage section Archiver appliance: Configure as required (sizing guide https://bp.veeam.com/vbr/2_Design_Structures/D_Veeam_Components/D_backup_repositories/gateway.html)

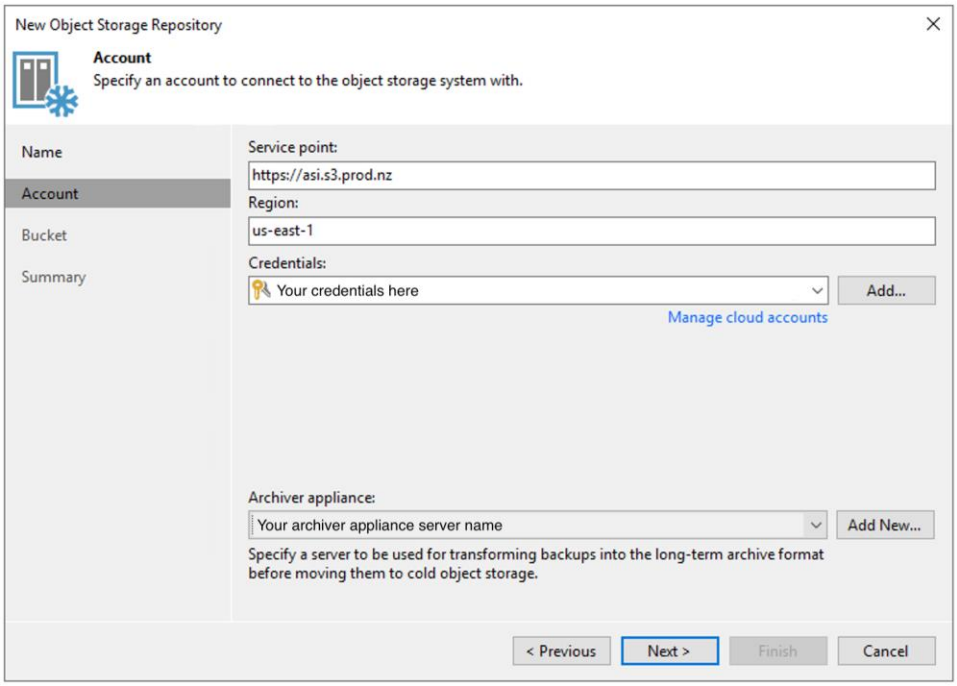

- Browse for your bucket and select it
- Create a folder for Veeam to place the data in, e.g. cold-01
- Configure immutability if required

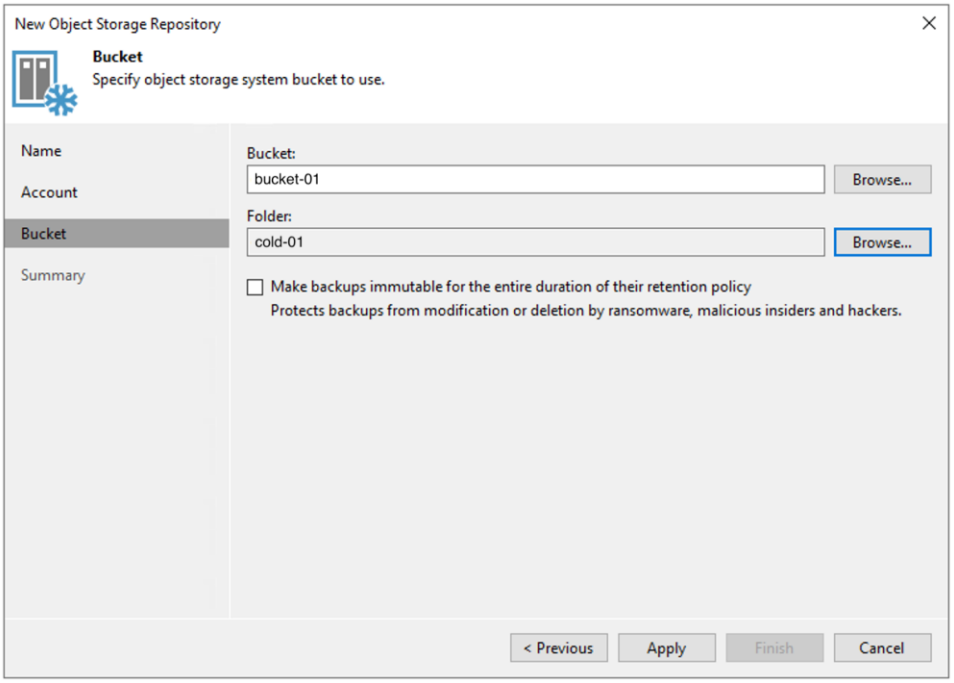

- Apply the settings to create your S3 Compatible with Data Archiving backup repository

Configuring Scale-Out Backup Repository for InfiniStor Direct to Archive

- Add a Scale-out Repository
- Give your scale-out repository a name
- On the Performance Tier settings, add the hot storage backup repository as a performance tier extent

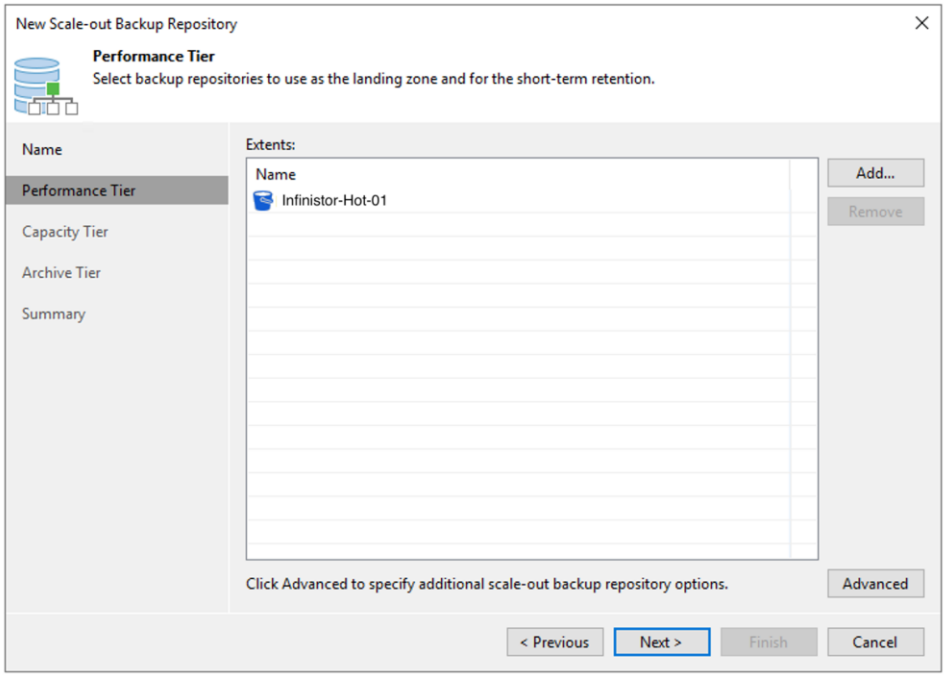

- Skip Capacity Tier
- On the Archive Tier settings, check ‘Archive GFS full backups to object storage’ and select your cold storage backup repository.
- Configure the ‘Archive GFS backups older than X days’ to 0 days if you want GFS restore points to be offloaded to cold storage as soon as possible, otherwise configure to a suitable amount of time for offloading e.g. 7 days will mean the GFS restore point needs to age 7 days before it will be offload to cold storage, 30 days needs to age 30 days before it will be offloaded and so on.

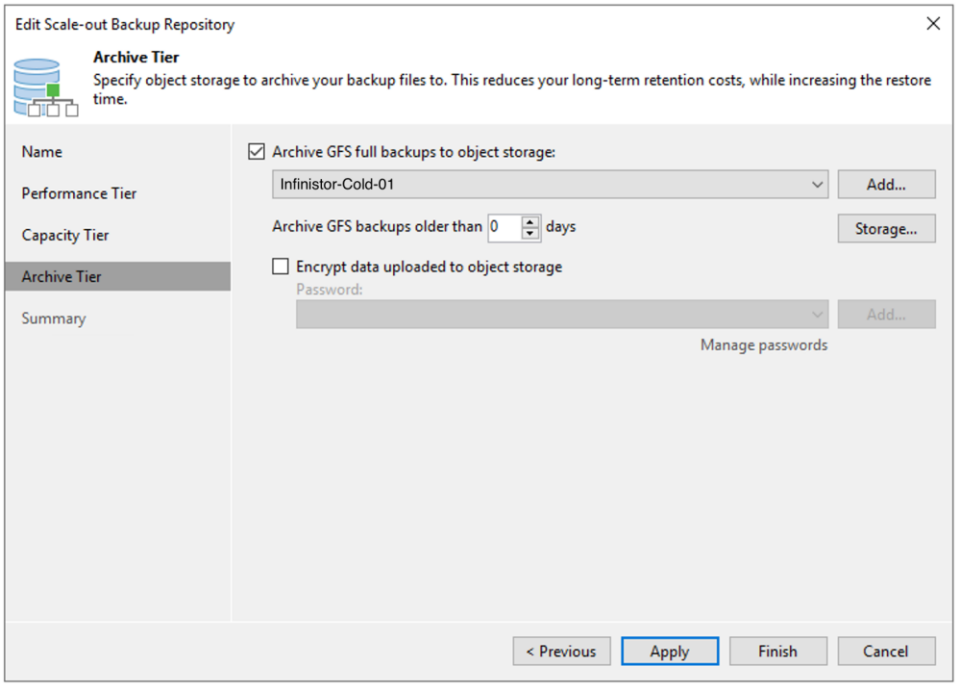

- Configure ‘Storage’ settings as defined above in Performance Tweaks for InfiniStor**.**

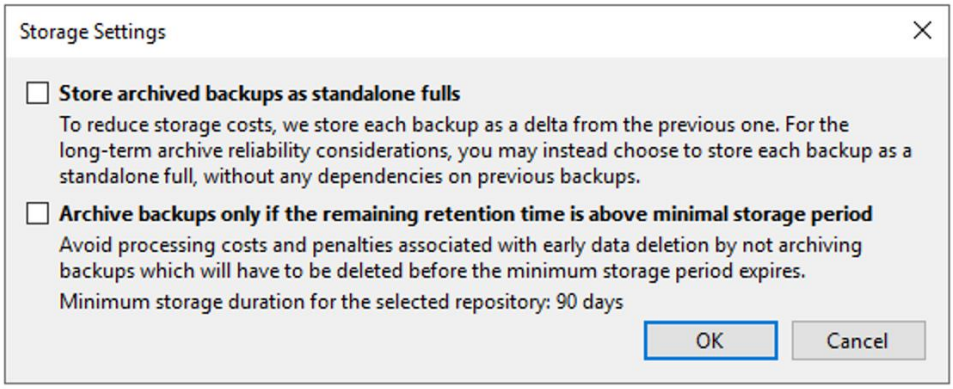

- Configure encryption if required (Note: This will impact how data is stored on cold storage as it will treat all GFS restore points as standalone fulls and not leverage previous files for deltas)
- Apply the settings to create your InfiniStor Scale-out Repository

Configuring Backup Jobs

This will be a simple overview on the settings you need to use with InfiniStor.

- Create a new backup job
- On the Storage settings, select the InfiniStor Scale-out Repository that you setup at the Backup repository.
  - Configure ‘Keep certain full backups longer for archival purposes’ for GFS retention/Archive Tier retention policy requirements as needed.
  - Click on Advanced > Storage and configure the Performance Tweaks for InfiniStor.

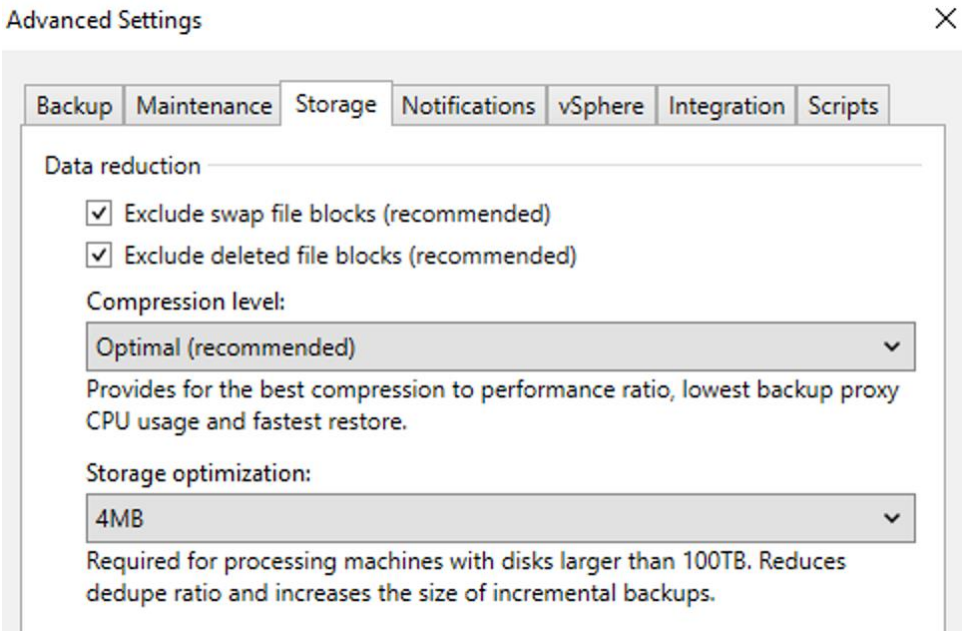

- Review and apply the settings to create/update your backup job

Notes on Archiving

- By default, there is an offloading job on the Scale-out Repository that occurs every 4 hours. This can be adjusted if you have the Capacity Tier configured but isn't used with the Direct to Archive configuration here.
- When an archiving job occurs, data is read from the performance tier via a gateway server, then transferred to the archive appliance for processing into the archive tier. The archiving job is limited to the concurrent tasks on the performance tier extent. If there is no limit and you have many workloads being processed for offloading, you may encounter a memory exhaustion issue. Please use the Veeam Sizing Guide for Gateways in the References section below for the archiver appliance sizing.

References

- Veeam Best Practices for Object Storage: https://bp.veeam.com/vbr/2_Design_Structures/D_Veeam_Components/D_backup_repositories/object.html
- Veeam Sizing Guide for Gateways: https://bp.veeam.com/vbr/2_Design_Structures/D_Veeam_Components/D_backup_repositories/gateway.html
- Veeam Adding S3 Compatible Object Storage Repositories: https://helpcenter.veeam.com/docs/backup/vsphere/adding_s3c_object_storage.html?ver=120
- Veeam Immutability for Object Storage: https://helpcenter.veeam.com/docs/backup/vsphere/immutability_object_storage_repositories.html?ver=120