Cohesity Data Cloud for NetBackup™ Administrator's Guide

Release 11.2

VERITAS
now part of Cohesity

Cohesity Data Cloud for NetBackup™ Administrator's Guide

Last updated: 2026-05-28

Legal Notice

Cohesity Support

Reach Cohesity Support

There are several ways to create a Cohesity support case.

- Go to Cohesity Support, to search in our knowledge base; or contact us by phone - United States and Canada: 1-855-9CO-HESI (926-4374), option 2.
- Log in to the Cohesity Support Portal to create a new case.
- Click the (**?**) icon on the Cohesity UI and select Support Portal.

Support/Service Assistance

First, contact the Service Provider that you have contracted for service and support. If you work directly with Cohesity and have a product warranty/entitlement, repair pricing, or technical support-related question, see your options below:

- To find solutions to your product issues or for suggestions or best practices, visit the Cohesity Knowledge Base.
- Log in to the Cohesity Support Portal to create a new case.
- To monitor your open cases, log in to the portal and click the **Cases** tab on the home page. This page should have all the case statuses and updates. You can also view individual case status.

Cohesity Software Running on Partner Hardware

For Cohesity software running on qualified third-party hardware, the following support workflow applies:

1. The customer may contact Cohesity Support first if the issue cannot be determined as a hardware issue.

   Note: Cohesity cannot process hardware replacement requests for partner hardware.

2. Cohesity Support triages the issue. If it is a software issue, Cohesity Support continues to work on it.
3. If it is a hardware/firmware issue or is suspected to be a hardware/firmware issue, Cohesity provides information about the issue to the customer and requests that the customer open a support ticket with the appropriate partner.
4. If needed, Cohesity Support can join a three-way call with the partner and the customer.
5. The customer informs Cohesity Support on the progress of the partner's case.

Contents

Chapter 1

Introducing Cohesity Data Cloud for NetBackup

This chapter includes the following topics:

About this guide

This guide describes how to configure and administer Cohesity Data Cloud for NetBackup. It covers storage provisioning, Direct-IO deployment, data movement and replication, backup policy configuration, and operational tasks.

The guide is intended for NetBackup and storage administrators with working knowledge of NetBackup concepts such as primary servers, media servers, disk pools, storage units, and policies.

About Cohesity Data Cloud for NetBackup

Cohesity Data Cloud for NetBackup integrates NetBackup with Cohesity Data Cloud to provide a scalable and efficient storage platform for backup, duplication, and

replication workflows. The solution uses Cohesity's distributed SpanFS architecture to support high concurrency, optimized data movement, and enterprise scale performance.

In this integration, NetBackup continues to manage backup policies, schedules, catalogs, and job control, while Cohesity Data Cloud provides the underlying storage, data services, and replication capabilities. Optimizations such as Client Direct, optimized duplication, and auto image replication help reduce resource usage and improve backup and recovery operations.

This architecture enables organizations to simplify storage management, support data retention and WORM requirements, and scale backup workloads without changing existing NetBackup workflows.

Key features and capabilities

Cohesity Data Cloud for NetBackup provides the following key capabilities:

- Containerized NetBackup media services on the Cohesity platform
  NetBackup media server functionality runs as containerized services on Cohesity cluster nodes. This design eliminates the need for a dedicated standalone media server infrastructure while allowing an external NetBackup primary server to continue managing policies, scheduling, and catalog operations.
- NetBackup-driven job orchestration and distribution
  NetBackup Job Manager continues to control job scheduling and distributes backup and restore jobs across available media server instances. This ensures load-balanced execution and preserves standard NetBackup job control behavior while leveraging the distributed infrastructure.
- Direct-IO optimized data path to SpanFS storage
  Uses NetBackup Direct-IO to enable a direct, high-performance data path from NetBackup media services and clients to SpanFS storage. This approach bypasses traditional file system mount paths and removes intermediate bottlenecks, improving backup and recovery performance.
- Scale-out, enterprise-grade storage using SpanFS
  Backup data is stored on Cohesity's distributed SpanFS file system, which provides deduplication, compression, snapshots, immutability (WORM), replication, and fault tolerance. This architecture replaces traditional MSDP and array-based deduplication storage.
- High-performance and resilient distributed data plane
  The solution uses distributed services to provide load balancing, secure communication, and parallel data transfer across cluster nodes. This enables scalable, high-throughput backup and recovery operations.

- Support for optimized duplication and Auto Image Replication (AIR)
  Supports the efficient data movement across storage targets using NetBackup storage lifecycle policies, enabling optimized duplication, and cross-domain replication workflows.
- Integration with existing NetBackup workflows
  Preserves an existing NetBackup administration workflows, including policies, schedules, catalogs, and job control, allowing seamless adoption without changes to operational processes.
- Support for data protection and compliance requirements
  Provides WORM storage and retention capabilities to meet enterprise data protection and regulatory requirements.

These capabilities enable efficient, scalable, and secure data protection while preserving familiar NetBackup administration and operational practices.

Architecture and components

Cohesity Data Cloud for NetBackup uses a distributed architecture that integrates NetBackup with Cohesity Data Cloud to support scalable backup, duplication, and replication operations. The solution separates control and data paths while enabling optimized data movement between NetBackup and Cohesity storage.

At a high level, NetBackup continues to manage backup policies, scheduling, catalogs, and job control, while Cohesity Data Cloud provides the underlying storage, data services, and distributed processing capabilities through its SpanFS architecture.

The solution consists of the following main components:

- NetBackup primary server
  Manages backup policies, schedules, catalogs, and overall job coordination. It acts as the control point for configuring and monitoring NetBackup operations.
- NetBackup media server
  Responsible for data movement and metadata processing. The media server participates in optimized data paths and coordinates data transfer between NetBackup and Cohesity storage.
- NetBackup clients
  Host the workloads being protected. Clients can participate directly in data transfer operations when Client Direct is enabled, reducing load on media servers.
- Cohesity Data Cloud cluster

Provides distributed storage, data services, and replication capabilities using the SpanFS file system. The cluster enables scalable storage and optimized data services for NetBackup workloads.

This architecture supports multiple deployment models, including Client Direct, media server–based, and mixed deployments. These models allow customers to scale performance and concurrency while maintaining familiar NetBackup workflows.

Topologies and use cases

Cohesity Data Cloud for NetBackup supports multiple deployment topologies to address different customer requirements such as infrastructure consolidation, phased migration, and storage modernization.

Topology 1: External primary with cluster media servers

External NetBackup primary server + Cohesity Data Cloud cluster (embedded media servers)

In this topology, the NetBackup primary server is deployed externally, while NetBackup media server functionality runs as containerized services on the Cohesity cluster.

Use cases:

- Reduce media server sprawl by adopting a hyperconverged architecture that eliminates external media servers.
- Replace existing media server infrastructure during hardware refresh.

Figure 1-1 External primary with cluster media services deployment

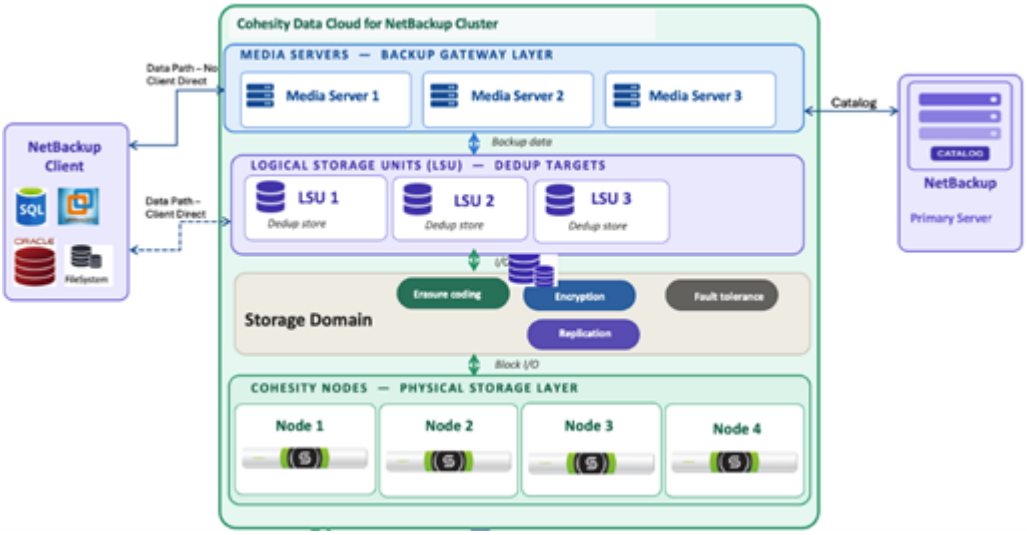

Topology 2: Hybrid deployment

External NetBackup primary server + Cohesity Data Cloud for NetBackup with embedded media servers + External NetBackup media server

In this hybrid model, the NetBackup primary server is deployed externally, while NetBackup media server functionality runs both as containerized services on the Cohesity cluster and on external media servers.

Use cases:

- Support phased transition, where selected workloads are migrated to cluster media servers while existing workloads continue to use external media servers.
- Support environments that require capabilities like Fiber channel initiator workflows, malware scanning, and other related features.

Figure 1-2 Hybrid deployment with cluster and external media servers

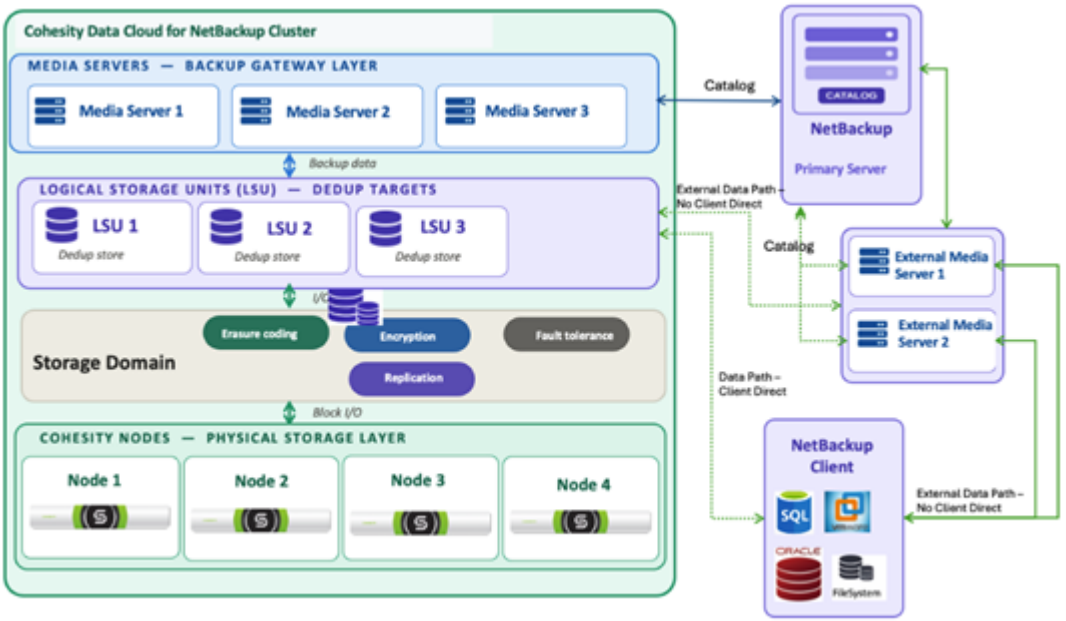

Topology 3: SpanFS as Storage Target

External NetBackup primary server + External media servers + Cohesity Data Cloud for NetBackup cluster

In this topology, the NetBackup primary server and media servers are deployed externally, while Cohesity Data Cloud is used as the storage target using SpanFS.

Use cases:

- Replace or refresh third-party storage with SpanFS without modifying the existing NetBackup environment.
- Support phased transition to a fully integrated, hyperconverged deployment model.

Figure 1-3 External media servers with SpanFS as storage target

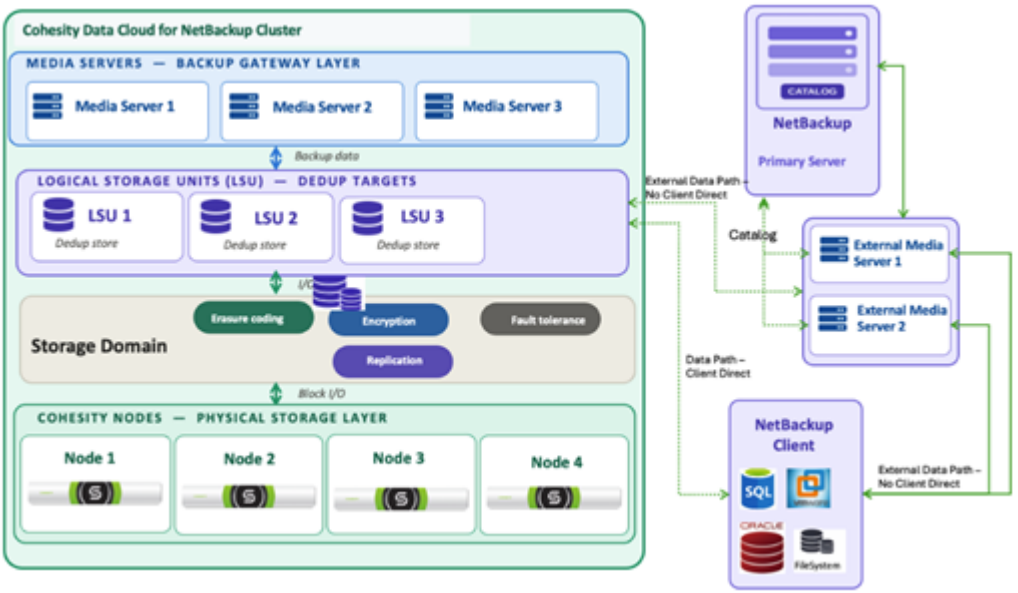

About Client Direct in NetBackup Direct-IO for SpanFS

Client Direct is a key optimization feature in NetBackup Direct-IO for SpanFS that enables backup clients to send data directly to the SpanFS storage cluster, bypassing the media server for the data path. This approach significantly enhances performance, reduces resource consumption on media servers, and improves scalability for backup and recovery operations.

In traditional OST implementations, backup data typically flows from the client to the media server, and then to the storage target. With Client Direct, NetBackup clients equipped with the Direct-IO OST plug-in can deduplicate and stream data directly to SpanFS, while the media server continues to manage job control and metadata.

In the initial release, Client Direct is supported on RHEL and Windows operating systems. Other platforms such as AIX require a media server to act as an intermediary.

Following are the primary advantages of using Client Direct in NetBackup Direct-IO for SpanFS:

- Source-side deduplication: Reduces data volume before transmission, saving bandwidth and storage.
- Improved Performance: Eliminates intermediate hops, resulting in faster backups and restores.
- Reduced Media Server Load: Frees up media server resources for other tasks.
- Scalable Recovery: Leverages SpanFS's distributed architecture for parallel recovery across cluster nodes.

The workflow:

1. The NetBackup media server initiates the backup job.
2. The client, equipped with the Direct-IO plug-in, deduplicates and streams data directly to SpanFS.
3. SpanFS distributes the data across its cluster using its View Domain architecture.
4. The media server manages job control and metadata but does not handle the data payload.

NetBackup Direct-IO feature support matrix

NetBackup Direct-IO integration with Cohesity SpanFS provides optimized data movement for backup and restore operations. However, not all NetBackup features are available when using Direct-IO. The following table lists the capabilities that are supported and not supported.

For detailed information on NetBackup Direct-IO workload support, Client Direct support, and stream handler support, refer to the NetBackup Software Compatibility List (SCL).

Table 1-1 Direct-IO feature support matrix

Feature	Support
Backup to SpanFS storage	Yes
Backup to S3 storage	No
Restore from SpanFS storage	Yes
Restore from S3 storage	No
Accelerator backups	Yes
AIR (Auto Image Replication)	Yes

Table 1-1 Direct-IO feature support matrix *(continued)*

Feature	Support
AIR reverse replication for IRE (Isolated recover environment)	No
Optimized duplication	Yes
Client-side deduplication (Client Direct)	Yes Not supported for policies that require universal share or data mover support.
Media server deduplication	Yes Any workload that is supported or not supported with client direct not configured and sent to the media server for deduplication.
Intelligent catalog archiving	No
Malware scanning	Yes Supported only for MS-Windows File System, Linux File System, and DNAS
WORM	Yes
Job-level deduplication reporting in the Job Details and Activity Monitor	Yes
Transport Layer Security (TLS) protocol	Yes. TLS 1.3 is supported starting NetBackup 11.1.0.2.
VMware Instant Access VM recovery	Yes. Supported only on Linux media servers with MSDP.

Chapter 2

Deploying Cohesity Data Cloud for NetBackup

This chapter includes the following topics:

Pre-deployment requirements

Before you deploy Cohesity Data Cloud for NetBackup, verify that your environment meets the following prerequisites. These requirements ensure proper connectivity, integration, and successful deployment.

Network and DNS requirements

Configure the network, IP addressing, and DNS resolution to support communication between NetBackup components and the Cohesity cluster:

- Assign one physical IP address for each node for node-level communication.
- Configure one Cohesity virtual IP (VIP) per node for cluster services.
- Configure one media server virtual IP (VIP) per node for media services access.

- Define a single fully qualified domain name (FQDN) for Cohesity VIP resolution.
- Define a unique fully qualified domain name (FQDN) for each media server VIP.
- Ensure forward and reverse DNS resolution is configured correctly for all IP addresses and hostnames.
- Verify that firewall rules allow all required network ports.

Proper VIP and DNS configuration is critical for cluster access, load balancing, and high availability in SpanFS environments.

Software requirements

Ensure that all required software components are installed and supported:

- NetBackup version 11.2
- NetBackup Direct-IO plug-in version 1.1.0.2_0014 (bundled with NetBackup primary, media, and client installation package)
- Cohesity Data Cloud for NetBackup version 7.4 or later

Required network ports

Open the following ports to allow communication between NetBackup components and Cohesity Data Cloud for NetBackup:

Table 2-1 Network ports

Source	Target	Ports	Purpose	Direction
NetBackup primary server	Media server (Cohesity Data Cloud for NetBackup)	1556, 13724	Core control communication (PBX, vnetd)	Two-way
Media server	NetBackup primary server	1556, 13724	Job status and catalog updates	Two-way
NetBackup primary server	Client	1556, 13724	Policy push and job initiation	Two-way
Client	NetBackup primary server	1556, 13724	Registration and heartbeat	Two-way
Media server	Client	13701, 13782, 13783	Backup and restore control services	Two-way

Table 2-1 Network ports *(continued)*

Source	Target	Ports	Purpose	Direction
Client	Media server	1556, 13724	Session initiation (Client Direct)	Two-way
Client	Media server	10000–65535 (or custom range)	Backup data transfer	Two-way
Media server	Client	10000–65535 (or custom range)	Data streaming	Two-way
Admin or user	NetBackup primary server	443, 8443	Web UI and REST API access	Two-way

Access key requirements

To enable integration between NetBackup and Cohesity Data Cloud, create an API access key on the NetBackup primary server.

You can create the key using either the NetBackup web UI or REST APIs. The access key is used by Cohesity nodes to authenticate and retrieve NetBackup configuration data.

To validate the access key, run a REST API query from a Cohesity node to confirm connectivity with the NetBackup primary server.

```
{
     "data": {
        "type": "apiKeyCreationRequest",
       "attributes": {
         "expireAfterDays": "P365D",
          "userName": "root",
          "userDomain": "",
         "userDomainType": "unixpwd",
            "description": ""
         }
      }
}
```

Run the following command to validate the access key. Using this access key, the Cohesity node can retrieve NetBackup host information:

```
curl --resolve <<NBU_Primary_Hostname>>:443:<<NBU_Primary_IP>> -s -k
-X GET "https://<<NBU_Primary_Hostname>>/netbackup/config/hosts" \
-H "Authorization: <<Access_Key>>" \ -H "accept:
application/vnd.netbackup+json;version=13.0" | jq
```

Deployment checklist

Use the following checklist to track deployment readiness and progress.

- Confirm that all hardware is provisioned and available.
- Verify that operating systems are installed and updated.
- Configure physical IP addresses on all nodes.
- Validate network connectivity between all components.
- Verify that required firewall ports are open.
- Validate that supported versions of NetBackup, Direct-IO plug-in, and Cohesity Data Cloud for NetBackup.
- Create required NetBackup access keys.
- Validate access key connectivity from all nodes. Verify forward and reverse DNS

Installing the NetBackup primary server and the media server

For detailed instructions to install the NetBackup primary server and the media server, refer to the *NetBackup 11.2 Installation Guide*.

If the deployment involves Cohesity Data Cloud for NetBackup with embedded media servers and you are using the cluster UI to register the storage server, install the Direct-IO plug-in on the NetBackup primary server.

Installing the Direct-IO plug-in on NetBackup media server

The Direct-IO plug-in is included as part of the NetBackup installation package. During the installation or upgrade of NetBackup, ensure that the Direct-IO plug-in is selected when prompted in the installation process.

If the Direct-IO plug-in was not installed during the NetBackup installation, follow the steps below to install the Direct-IO plug-in.

Before beginning the installation, ensure that the following prerequisites are met:

- NetBackup 11.2 is installed and operational.
- The required Emergency Engineering Binaries (EEBs) have been obtained from Cohesity Support.
- You have root or administrative access to the media server.
- Network ports required for SpanFS communication are open.
  See "Required network ports" on page 17.
- The Direct-IO plug-in RPM package is available.

To install the Direct-IO plugin on NetBackup media server

1. Stop the NetBackup services.

   ```
   bp.kill_all
   ```

2. Install the required EEBs.

   For detailed instructions, refer to the following article:

   Using the NetBackup Emergency Engineering Binary (EEB) installer

3. Install the Direct-IO plug-in.

   - RHEL: `rpm -ihv VRTSnbdirectio.rpm.rpm`

4. Start the NetBackup services.

   ```
   bp.start_all
   ```

Installing the Direct-IO plug-in on NetBackup client

The Direct-IO plug-in is included as part of the NetBackup installation package. During the installation or upgrade of NetBackup, ensure that the Direct-IO plug-in is selected when prompted in the installation process.

If the Direct-IO plug-in was not installed during the NetBackup installation, follow the steps below to install the Direct-IO plug-in.

Before beginning the installation, ensure that the following prerequisites are met:

- NetBackup 11.2 client software is installed and operational.
- The required Emergency Engineering Binaries (EEBs) have been obtained from Cohesity Support.
- You have root or administrative access to the client.
- The Direct-IO plug-in installer is available (RPM for RHEL, MSI for Windows).

To install the Direct-IO plug-in on NetBackup client

1. Stop the NetBackup services.

   ```
   bp.kill_all
   ```

2. Install the required EEBs.

   For detailed instructions, refer to the following article:

   Using the NetBackup Emergency Engineering Binary (EEB) installer

3. Install the Direct-IO plug-in.

   - RHEL: `rpm -ihv VRTSnbdirectio.rpm.rpm`
   - Windows: Run the `Cohesity NetBackup Direct-IO.msi` installer and follow the prompts.

4. Start the NetBackup services.

   ```
   bp.start_all
   ```

Direct-IO plug-in configuration values

Create the Direct-IO plug-in configuration file at the following location:

- Windows:
  `<installation_directory>/NetBackup/bin/ost-plugins/spanfs.conf`
- Linux: `/usr/openv/lib/ost-plugins/spanfs.conf`

This configuration file is in JSON format. For example:

```
{ "LOGLEVEL": 5 }
```

If this file does not exist, the plug-in uses the default settings and continues to function normally.

The following configuration values are supported:

Table 2-2 Direct-IO plug-in configuration values

Parameter	Description
`logLevel`	The plug-in logging verbosity. Allowed values: 1 - 6 Default: 3

Table 2-2 Direct-IO plug-in configuration values *(continued)*

Parameter	Description
`logPath`	Specifies the directory where Direct-IO plug-in log files are stored. ■ Linux: `/usr/openv/netbackup/logs/libstspispanfs` ■ Win: `<installation_directory>/NetBackup/logs/libstspispanfs`
`logSize`	The maximum log file size in MB. Default: 500

Cluster virtual IP configuration

Cluster Virtual IP (VIP) configuration in a Cohesity SpanFS environment is designed to provide high availability, load balancing, and simplified cluster access. Even though the NetBackup Direct-IO plug-in provides intelligent data path allocation dynamically for NetBackup jobs to cluster nodes, the cluster virtual IP increases the resiliency and balancing of the control path. By using a single virtual host name, systems like NetBackup can interact with the cluster without needing to manage individual node details.

Besides cluster node IPs, this setup uses floating VIPs that are not tied to specific nodes. These IPs can shift between nodes if any node undergoes an outage and ensures that all control traffic is always serviced.

The following are two primary approaches to configure the cluster VIPs:

- Using external DNS server
  See "Configuration using an external DNS server" on page 22.
- Using DNS delegation
  See "Configuration using the DNS delegation" on page 24.

Configuration using an external DNS server

In this approach, all DNS records including A records for individual nodes and the cluster virtual host name are managed on your organization's external DNS server. Clients and infrastructure components reference this server for name resolution.

When using an external DNS server, all DNS configurations including A records for individual nodes and the cluster virtual host name must be created and managed

on that server. The Cohesity cluster, infrastructure components, and clients will all reference this DNS server for name resolution.

The following diagram illustrates a typical VIP configuration using an external DNS server:

Figure 2-1 VIP configuration with external DNS server

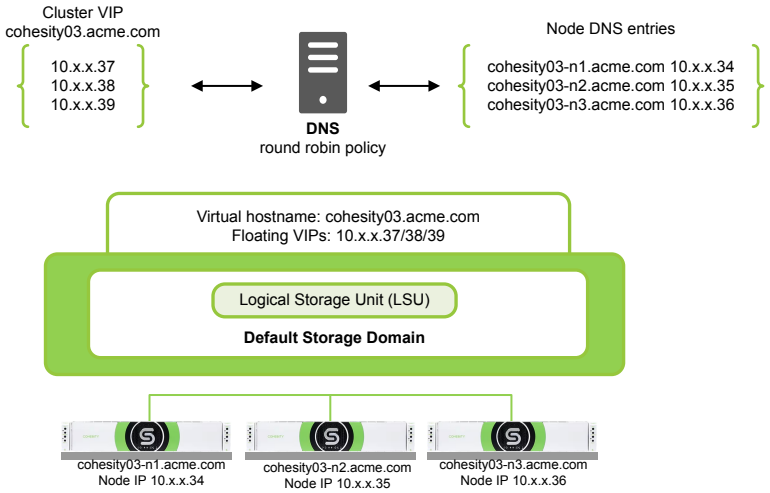

To achieve the above configuration, DNS entries are required as follows:

Table 2-3

Record type	Host name	IP address
A	cohesity-n1.acme.com	10.x.x.50
A	cohesity-n2.acme.com	10.x.x.51
A	cohesity-n3.acme.com	10.x.x.52
A	cohesity-n4.acme.com	10.x.x.53
A	cohesity.acme.com	10.x.x.30
A	cohesity.acme.com	10.x.x.31
A	cohesity.acme.com	10.x.x.32
A	cohesity.acme.com	10.x.x.33

The round-robin policy on the DNS server ensures that requests to cohesity.acme.com are distributed across the available VIPs.

Once the DNS records are in place, the VIPs can be configured during the initial cluster setup. In the Cohesity UI, navigate to: **Settings > Networking > VIPs** tab.

Here, you can view and manage the assigned VIPs, ensuring that they align with the DNS entries.

Configuration using the DNS delegation

This method delegates DNS responsibilities to the Cohesity cluster's internal DNS service, which manages most VIP-related configurations. It simplifies setup and reduces dependency on external DNS infrastructure.

This approach is ideal for environments where internal DNS control is preferred, or external DNS access is limited. For more information about this approach, see the Cohesity documentation.

Benefits of this method:

- Simplified DNS setup and management
- Reduced reliance on external DNS servers
- Integrated control within the Cohesity cluster

Configuring SpanFS subnet whitelists for NetBackup access

SpanFS enforces IP based access control for filesystem RPC operations. To allow NetBackup media servers and clients to communicate with the SpanFS file system during Direct-IO operations, you must configure one or more IP subnet whitelists on the SpanFS cluster. Only the systems within these whitelisted subnets are permitted to access SpanFS services.

You can use the `iris_cli` command-line tool to manage subnet whitelists.

To configure subnet whitelists for NetBackup

1. Add an IP subnet to the whitelist.

   ```
   iris_cli cluster add-subnet-to-whitelist subnet-ip=<IPv4 address>
   subnet-mask-ip4=<IPv4 subnet mask>
   ```

2. View all subnet entries currently allowed to access SpanFS.

   ```
   iris_cli cluster get-subnet-whitelist
   ```

3. Remove a subnet from the whitelist.

   ```
   iris_cli cluster remove-subnet-from-whitelist subnet-ip=<subnet
   IPv4 address>
   ```

Chapter 3

Provisioning the storage

This chapter includes the following topics:

About Direct-IO storage configuration for Cohesity Data Cloud

NetBackup integrates with the SpanFS storage backend using a layered storage configuration model. Each layer serves a distinct purpose, and the components must be created in a specific order before storage can be used by NetBackup policies.

Note: In the NetBackup web UI, you can create storage domains and logical storage units (LSUs) as part of the disk pool creation workflow.

The storage configuration hierarchy is Storage server → Storage domain → Logical storage unit (LSU) → Disk pool → Storage unit.

Each object builds on the previous one to define how data is stored, managed, and accessed by NetBackup.

- Storage server
  A storage server represents the SpanFS cluster in NetBackup. It establishes secure communication between NetBackup and SpanFS and enables NetBackup to discover storage domains and LSUs. A storage server must be created before any disk pools can be configured.
  See "Adding the Direct-IO storage server using the NetBackup web UI" on page 27.
- Storage domain
  A storage domain defines how data is stored on the SpanFS cluster. It controls data layout characteristics such as deduplication, compression, encryption, fault tolerance, quotas, and WORM behavior. A storage domain is required before you can create an LSU.
  See "Creating a storage domain using the NetBackup web UI" on page 28.
- Logical storage unit (LSU)
  An LSU is a logical container created within a storage domain. NetBackup uses LSUs as storage volumes for backup data. Each LSU belongs to a single storage domain and maps one to one with a disk pool in NetBackup.
  See "Creating a logical storage unit (LSU)" on page 30.
- Disk pool
  A disk pool is a NetBackup object that associates a SpanFS storage server with an LSU. It represents NetBackup's view of available storage and defines disk level properties such as high and low water marks. A disk pool is required before you can create a storage unit.
  See "Creating a disk pool for a SpanFS storage server" on page 31.
- Storage unit
  A storage unit is the object that NetBackup policies reference. It is created from a disk pool and controls how backup jobs write data to storage, including concurrency and capacity behavior.
  See "Creating a storage unit" on page 32.

Adding the Direct-IO storage server using the NetBackup web UI

Use this procedure to create a Direct-IO storage server with SpanFS storage server type. You have the option to create a disk pool (local storage or cloud storage) and storage unit after you create a storage server. The recommendation is that you create the disk pool and storage unit if they do not exist in NetBackup.

To add the Direct-IO storage server using the NetBackup web UI

1. On the left, select **Storage > Disk storage**. Select the **Storage servers** tab, then click **Add**.
2. In the **Storage type** list, select **Disk storage servers**.
3. From the **Category** options, select **DirectIO**. Click **Start**.
4. On the **Basic properties** page, select a media server that has the Direct-IO plug-in installed.
5. Under Storage server details enter the following details:
   - Storage server name

     Note: If you want to perform VMware Instant Access VM recovery, you must configure the SpanFS storage server using the hostname. It does not work if the storage server is configured using an IP address.

   - Username and password
6. Click **Next**.
7. When the **Validate Certificate Authority** dialog appears, verify the Certificate Authority details and fingerprint, then click **Yes** to trust the certificate and continue.
8. (Optional) On the **Media servers** page, click **Add** to add any additional media servers that you want to use. Click **Next**.
9. On the **Review** page, confirm that all options are correct and click **Save**.
10. After you click **Save**, the credentials you entered are validated.

Creating a storage domain using the NetBackup web UI

A storage domain defines how data is stored on the SpanFS cluster, including settings for deduplication, compression, encryption, fault tolerance, and immutability. Storage domains are used when creating logical storage units (LSUs), which are then associated with disk pools in NetBackup.

You must create a storage domain before you can create an LSU. You can create a storage domain using the NetBackup web UI during disk pool creation, from the **Add volume** page.

Before you create a storage domain, ensure that:

- The SpanFS storage server is already configured in NetBackup.
- You have the required permissions to create storage domains.
- You have planned the required data storage characteristics, such as deduplication, encryption, and fault tolerance.

To create a storage domain using the NetBackup web UI

1. In the NetBackup web UI, on the left, click **Storage > Disk storage**, select the **Disk pools** tab, and then click **Add**.
2. On the **Disk pool options** page, click **Change** and select **SpanFS** as the storage server type.
3. Click **Select** and enter the following details:
   - Disk pool name
   - (Optional) Description
   - (Optional) Limit I/O streams
   - High water mark and Low water mark
4. Click **Next**.
5. On the **Volumes** page, click **Add** to add a new volume.
6. Under **Storage domain**, click **Add**.
7. In the **Add storage domain** dialog box, enter a name for the storage domain.
8. Configure the required storage domain settings:
   - Deduplication
     Enable deduplication for the storage domain. Select the deduplication mode as Inline (recommended) or Post process.
   - Compression (optional)
     Enable and select Inline or Post process compression.
   - Encryption (optional)
     Enable encryption and select the appropriate key management service (KMS) type.
9. (Optional) Configure **Physical quotas and alerts**:
   - Enable **Storage domain quota**.
   - Specify an **Alert threshold** value. These settings apply to the entire storage domain.
10. Configure **Fault tolerance and redundancy**:
    - Select the required fault tolerance level (for example, 1D:1N)

- Select Erasure Coding settings (for example, EC2:1)
- Review the usable storage percentage displayed for the selected settings

11 (Optional) Enable **Post processing only** if the storage domain is intended exclusively for post processed workloads.

Click **Add**.

After the storage domain is created, it becomes available for use when creating logical storage units (LSUs).

Creating a logical storage unit (LSU)

A logical storage unit (LSU) is a logical container within a storage domain on the SpanFS cluster that NetBackup uses to store and retrieve backup data. Each LSU represents a logical grouping of storage that NetBackup consumes when creating disk pools, which are then used to create storage units.

You must create an LSU before it can be selected when creating a disk pool.

Before you create an LSU, ensure that:

- A storage domain already exists on the SpanFS cluster.
- The SpanFS storage server is configured in NetBackup.
- You have planned whether the LSU requires immutability (WORM).

Note: Storage domains define the data layout characteristics such as deduplication, compression, encryption, and fault tolerance.

See “Creating a storage domain using the NetBackup web UI” on page 28.

To create a logical storage unit (LSU)

1 In the NetBackup web UI, on the left, click **Storage > Disk storage**, select the **Disk pools** tab, and then click **Add**.

2 On the **Disk pool options** page, click **Change** and select **SpanFS** as the storage server type.

3 Click **Select** and then click **Next**.

4 On the **Volumes** page, click **Add**.

5 Configure the LSU details:

- Volume name: Enter a name for the LSU.

- Storage domain: Select the storage domain in which the LSU must be created.
  Click **Add** to add a new storage domain.
  See "Creating a storage domain using the NetBackup web UI" on page 28.
- WORM (optional): Enable WORM if the LSU requires immutable storage. Select one of the following:
  - Enterprise: Allows retention settings to be modified by authorized administrators.
  - Compliance: Enforces strict, non modifiable retention to meet regulatory requirements.

6 Click **Add**.

The logical storage unit is created and becomes available for selection during disk pool creation.

Creating a disk pool for a SpanFS storage server

When you create a disk pool, you associate a SpanFS storage server with one or more logical storage units (LSUs). A disk pool is required before you can create a storage unit in NetBackup.

When creating a disk pool, you specify:

- The storage server (SpanFS storage server)
- The logical storage unit (LSU) to include in the disk pool
- Disk pool properties

Cohesity recommends that disk pool names and disk volume names be unique across your enterprise.

Before you create a disk pool, ensure that:

- The SpanFS storage server is already configured in NetBackup.
- You have the required permissions to create disk pools.
- At least one storage domain exists, or you are prepared to create one during the workflow.

Note: A storage domain must exist before you can create a volume. You can create a storage domain in advance, or directly from the Add volume page during disk pool creation.

To create a disk pool for a SpanFS storage server

1. In the NetBackup web UI, on the left, click **Storage > Disk storage**, select the **Disk pools** tab, and then click **Add**.

   You can also start disk pool creation by clicking Create disk pool after creating a storage server.

2. On the **Disk pool options** page, click **Change** and select **SpanFS** as the storage server type.

3. Click **Select** and enter the following details:

   - Disk pool name
   - (Optional) Description
   - (Optional) Limit I/O streams
   - High water mark and Low water mark

4. Click **Next**.

5. On the **Volumes** page, select a volume from the volume list.

   The list displays the logical storage units (LSUs) created on the SpanFS cluster.

   Click **Add** to add a new volume.

   See "Creating a logical storage unit (LSU)" on page 30.

6. Click **Next**.

7. On the **Replication** page, click **Add** to add the replication target.

   See "Pairing clusters, storage domains, and LSUs" on page 33.

8. Click **Next**.

9. On the **Review** page, verify that all configuration details are correct, and then click **Finish**.

After the disk pool is created successfully, you can create a storage unit using the disk pool.

Creating a storage unit

A storage unit defines how NetBackup writes backup images to storage and is the final object required before a disk pool can be used by a NetBackup policy. A storage unit references a disk pool, which in turn is backed by a logical storage unit (LSU) on a SpanFS storage server.

You create storage units after creating a disk pool.

Before you create a storage unit, ensure that:

- A disk pool is already created and in an UP state.
- The disk pool is associated with the correct SpanFS storage server and LSU.
- You have NetBackup administrator privileges.

Note: Disk pool properties such as high water mark and low water mark are inherited by the storage unit.

To create a storage unit

1. In the NetBackup web UI, on the left, click **Storage > Storage units**.
2. Click **Add**.
3. On the **Basic properties** page, enter the unique storage unit name and configure the following storage unit properties as required:
   - **Maximum concurrent jobs**: Specify the maximum number of jobs that can write to this storage unit at the same time.
   - **Maximum fragment size**: Specify the maximum fragment size (in MB).
4. Click **Next**.
5. On the **Disk pool** page, elect the disk pool to associate with the storage unit, and then click **Next**.
6. On the **Media server** page, choose one of the following options:
   - **Allow NetBackup to automatically select**: Allows NetBackup to select a media server based on load balancing.
   - **Manually select**: Select one or more specific media servers from the list.
7. Click **Next**.
8. On the **Review** page, verify the configuration details, and then click Save.

The storage unit is created and becomes available for selection in NetBackup policies and storage lifecycle policies.

Pairing clusters, storage domains, and LSUs

To support optimized duplication and Auto Image Replication (AIR), NetBackup must establish a replication relationship between the source and target SpanFS clusters, along with their associated storage domains and logical storage units (LSUs).

In NetBackup, the pairing for Auto Image Replication (AIR), is established when you add a replication target to a disk pool during disk pool creation or by updating an existing disk pool. For optimized duplication, it is done during SLP creation of optimized duplication.

Before you add the replication target, ensure that:

- Source and target SpanFS clusters are reachable over the network.
- Storage domains and LSUs exist on both the source and target SpanFS clusters.
- Both SpanFS clusters are added as storage servers in NetBackup.
- You have NetBackup administrator privileges.

When you add a replication target to a disk pool, NetBackup performs the following actions:

1. Associates the source disk pool with a target storage server.
2. Identifies the corresponding target storage domain and LSU.
3. Establishes trust and communication between the source and target SpanFS clusters.
4. Enables inter cluster data transfer required for replication.

To pair clusters, storage domains, and LSUs by updating an existing disk pool

1 In the NetBackup web UI, click **Storage > Disk storage**, and then select the Disk pools tab.

2 Select the disk pool to update.

3 Click **Edit**.

4 Navigate to the **Replication** page.

5 Click **Add** to add a replication target.

6 In the Add replication targets window:

- Under **Select trusted primary server**, select the trusted primary server. Under **Select target storage server**, select the target SpanFS storage server and the appropriate target volume (LSU).

These settings apply only to Auto Image Replication (AIR) between NetBackup domains.

7 Click **Add**.

8 Save the disk pool configuration.

Chapter 4

Configuring the optimized duplication and replication

This chapter includes the following topics:

Configuring optimized duplication

Optimized duplication allows NetBackup to efficiently copy backup images between two SpanFS storage servers by transferring only the required data blocks. It uses the OpenStorage (OST) interface to enable direct communication between the source and target SpanFS clusters, resulting in faster duplication and reduced network usage.

Optimized duplication is configured using a Storage Lifecycle Policy (SLP) after the required storage objects are created on both the source and target NetBackup domains.

Before you configure optimized duplication, ensure that:

- The source and target SpanFS clusters are paired.
- Storage domains and LSUs are created and paired between the clusters.
- Both SpanFS clusters are added as storage servers in NetBackup.
- Disk pools and storage units are created on both the source and target NetBackup domains.

Note: If an LSU is created or updated after the disk pool is created, update the disk pool to discover the changes before configuring optimized duplication.

When the storage lifecycle policy is created, NetBackup establishes the pairing between the source and target SpanFS clusters, along with their associated storage domains and LSUs, if the pairing does not already exist.

To configure optimized duplication

1. In the NetBackup web UI, on the left, click **Storage > Storage lifecycle policies**.
2. Click **Add** to create a new storage lifecycle policy.
3. On the **Storage lifecycle policy** page, enter the storage lifecycle policy name.

   (Optional) Configure Data classification and Priority for secondary operations.
4. Click **Add** to add an operation.
5. In the New operation window:
   - From the **Operation** list, select **Backup**.
   - Under **Destination storage** attributes, select the source storage unit backed by the source SpanFS cluster.
   - Configure the required **Retention** settings.
   - Click **Create** to add the backup operation.
6. Select the Backup operation in the operations table, and then click **Add child**.
7. In the **New operation** window:
   - From the **Operation** list, select **Duplication**.
   - Under **Destination storage attributes**, select the target storage unit associated with the target SpanFS cluster.
   - Configure the required **Retention** settings.
   - (Optional) Configure duplication options such as **Alternate read server**.
8. (Optional) Configure additional settings, such as postponing the duplication until the source copy is about to expire.
9. Click **Create** to add the duplication operation.
10. Review the storage lifecycle policy configuration, and then click **Create**.

The storage lifecycle policy is created and can be used by backup policies to perform optimized duplication.

Configuring the auto image replication

Auto Image Replication (AIR) enables automatic replication of backup images from a source NetBackup domain to a target NetBackup domain for disaster recovery. Auto Image Replication is configured using a Storage Lifecycle Policy (SLP) that includes backup and replication operations.

Before you create an SLP, ensure that:

- Source and target SpanFS clusters are paired.
- Storage domains and LSUs are created and paired between the clusters.
- Both SpanFS clusters are added as storage servers in their respective NetBackup domains.
- Disk pools and storage units are created on both the source and target NetBackup domains.
- An import type storage lifecycle policy exists on the target NetBackup domain

---

Note: AIR requires two NetBackup domains: a source domain and a target domain. Replicated images are cataloged and managed independently on the target domain.

To configure Auto Image Replication (AIR)

1. In the NetBackup web UI, on the left, click **Storage > Storage lifecycle policies**.
2. Click **Add** to create a new storage lifecycle policy.
3. On the **Storage lifecycle policy** page, enter the storage lifecycle policy name.

   (Optional) Configure Data classification and Priority for secondary operations.
4. Click **Add** to add an operation.
5. In the New operation window:
   - From the **Operation** list, select **Backup**.
   - Under **Destination storage** attributes, select the source storage unit backed by the source SpanFS cluster.
   - Configure the required **Retention** settings.
   - Click **Create** to add the backup operation.
6. Select the Backup operation in the operations table, and then click **Add child**.
7. In the **New operation** window:
   - From the **Operation** list, select **Replication**.

- Under **Destination storage attributes**, select replication target storage servers (across different NetBackup domains) for the source storage server.
- Under **Target import SLP**, select the import storage lifecycle policy created on the target domain.
- Configure the required **Retention** settings.

8 (Optional) Configure additional settings such as postponing replication until the source copy is about to expire.

9 Click **Create** to add the replication operation.

10 Review the storage lifecycle policy configuration, and then click **Create**.

Chapter 5

Configuring the backup policy

This chapter includes the following topics:

Creating a backup policy

A backup policy defines what data NetBackup protects, when backups run, and which storage unit is used as the backup target. After you create a storage unit backed by SpanFS storage, you must create or update a backup policy to direct backups to that storage.

Before you create a backup policy, ensure that:

- A storage unit is already created and available.
- The storage unit is backed by the correct disk pool and LSU.
- You know the clients or workloads to include in the policy

Note: Backup policies reference storage units. You cannot select SpanFS storage directly in a policy.

To create a backup policy

1. In the NetBackup web UI, on the left, click **Protection > Policies**.
2. Click **Add**. On the **Attributes** tab, enter the following basic details:
   - **Policy name**:
     Enter a unique and descriptive name.
   - **Policy type**:

Select the appropriate policy type for the workload (for example, Standard, VMware, or Database).

Direct-IO supports the following primary NetBackup policy types:

- Standard
- Windows
- VMware
- DNAS
- Oracle

For detailed list of the supported workloads, refer to the NetBackup Software Compatibility List (SCL).

- **Policy storage**:
  Select the storage unit created for SpanFS storage.

3. On the **Schedules** tab, configure all the necessary schedules. For example, Full and incremental schedules.
4. On the **Clients** tab, depending on the policy type that you selected, specify the client systems or workloads to be protected.
5. On the **Backup selections** tab, depending on the policy type that you selected, add the files, database instances, or other objects that you want to protect.
6. For the policy types that have additional tabs, review and select the other policy options that are needed to complete the setup.
7. Review the policy configuration and click **Create**.

For more information about the policies, see *NetBackup Web UI Administrator's Guide*.

Chapter 6

Managing WORM storage and retention

This chapter includes the following topics:

About NetBackup SpanFS cluster storage support for immutable and indelible data

NetBackup SpanFS cluster storage server supports immutable and indelible data storage.

For more information, refer to the *Configuring immutability and indelibility of data in NetBackup* chapter in the *Veritas NetBackup Administrator's Guide, Volume I*.

A retention period lets you define a time for protecting the backup image. Once you define a retention period, SpanFS cluster stores a timestamp along with the image metadata to indicate when the retention period expires. After the retention period expires, the image data can be deleted.

You can get the following parameters for the retention period:

- Lock Minimum Duration
- Lock Maximum Duration

For more information refer to the workflow to configure immutable and indelible data topic in the *Veritas NetBackup Administrator's Guide, Volume I*.

SpanFS cluster storage supports the following retention period modes:

- Compliance mode
  Any type of user cannot overwrite or delete the data that is protected using the compliance mode for the defined retention period. Once you set a retention period for the data storage, you cannot shorten it and can only extend it.
- Enterprise mode
  Users require special permissions to disable the retention lock and then delete the image. Only the SpanFS cluster security administrator user can disable the retention lock and then the delete the image if required. You can use the enterprise mode to test the retention period behavior before you create a compliance mode retention period.

Updating LSU WORM attributes

You can update LSU WORM attributes like WORM mode, minimum retention, and maximum retention after LSU creation.

To update LSU WORM attributes

1. Log in to cluster node where the WORM LSU is configured.
2. Run the following command to update the LSU WORM attributes of the LSU:

   ```
   iris_cli lsu update
   ```

   For example, to update maximum retention period for WORM, run the following command:

   ```
   iris_cli lsu update name=wormlsu worm-max-retention-secs=1728000
   ```

   Sample output:

   ```
   LSU ID                  : 13516
   LSU NAME                : wormlsu
   NBU DOMAIN NAME         : <domain-name>
   STORAGE DOMAIN ID       : 2
   WORM CONFIG:
     MODE                  : Enterprise
     MAX RETENTION IN SEC  : 1728000
     MIN RETENTION IN SEC  : 60
   ```

3 On NetBackup primary server, run the following command to update the configuration:

```
nbdevconfig -updatedv -stype SpanFS -dp spanfs-worm-dp -dv wormlsu
```

4 Run the following command to verify that maximum retention period for WORM is updated:

```
nbdevconfig -previewdv -storage_server 10.84.87.59 -stype SpanFS
-U
```

Deleting WORM-locked images

You can shorten the WORM lock for a view to a near-future time. Once the lock expires, the view is unlocked, and you can delete the corresponding backup image in NetBackup.

Note: You can only shorten the WORM lock if the LSU is in Enterprise mode. Shortening the lock is not allowed in Compliance mode.

To delete WORM-locked images

1 Log in as a Data Security user.

2 Run the following command to update the view:

```
iris_cli view update include-internal=true name=<name>
datalock-expiry-usecs=<timestamp>
```

- Set <timestamp> to a time in the near future in microseconds.
- The value must be greater than the configured worm-min-retention-secs.

3 Verify that the DATALOCK STATE of the view is now UNLOCKED by listing the view.

4 On NetBackup primary server, run the following command to delete the WORM-locked backup:

```
bpexpdate -backupid <backup_image> -try_expire_worm_copy -d 0
-copy 1
```

5 Verify that the backup is now deleted in NetBackup catalog or run the following command:

```
bpimagelist
```

6. Verify on SpanFS by listing the view that the view no longer exists
7. Verify on SpanFS that the view does not exist any longer by listing available views.

Chapter 7

Cohesity Data Cloud for NetBackup operations

This chapter includes the following topics:

Managing the certificates for cluster communication

Secure communication between NetBackup and the Cohesity Data Cloud (SpanFS cluster) requires proper management and maintenance of SSL certificates.

This section describes how to renew the cluster certificate and update it in NetBackup.

See "Renewing the Iris certificate on the SpanFS cluster" on page 45.

See "Updating the certificate in NetBackup" on page 46.

Renewing the Iris certificate on the SpanFS cluster

The Iris service certificate is used for secure communication between NetBackup and the SpanFS cluster. If the certificate is expired, perform the following steps.

When the certificate on the cluster expires, it renews automatically. If it is an external certificate, it needs to be renewed manually.

To renew the Iris certificate:

1. Log in to the SpanFS cluster.
2. Run the appropriate CLI command to update the SSL certificate.

For detailed instructions about updating the certificate, see the Cohesity documentation.

After the certificate is renewed, you must update the certificate in NetBackup to ensure uninterrupted communication.

Updating the certificate in NetBackup

After you verify that the Iris service certificate is valid on the cluster, update the certificate configuration in NetBackup.

To update the certificate using the web UI

1. In the NetBackup web UI, navigate to**Storage > Disk storage**, and select the **Storage servers** tab.
2. From the list of configured storage servers, select the SpanFS storage server for which you want to update the certificate.
3. Click the action menu (three dots) for the selected storage server, and then select **Edit credentials**.
4. When the **Validate Certificate Authority** dialog box appears, review the certificate details.
5. Click **Yes** to trust the certificate and continue.
6. In the **Edit credentials** dialog box, enter the required credentials (username and password) for the storage server.

   Click **Update**.

To update the certificate using the command-line

- Run the following command on the NetBackup primary or media server:

  ```
  tpconfig -update -storage_server <server_name> -stype
  <server_type> \ -ca_file_path <SpanFS_CA_certificate_path> \
  -sts_user_id <user_ID> -password <password>
  ```

  Where:

  - `<server_name>` is the SpanFS storage server configured in NetBackup.
  - `<server_type>` is the storage server type (SpanFS).
  - `<SpanFS_CA_certificate_path>` is the path to the retrieved certificate file (for example, `/tmp/ca.pem`).
  - `<user_ID>` and `<password>` are the credentials used for the storage server

Chapter 8

Troubleshooting

This chapter includes the following topics:

Direct-IO plug-in is not installed

NetBackup does not include the Direct-IO plug-in by default. Therefore, the plug-in must be installed separately. If the plug-in is not installed, you may encounter an error when attempting to configure a SpanFS storage server.

Example error:

```
# nbdevconfig -creatests -storage_server <storage_server_IP> -stype
 SpanFS -media_server `hostname`
```

```
Could not connect to storage server. Please verify that
(<storage_server_IP>) is a valid storage server of the required type
 and is reachable on the network. If an OpenStorage plug-in is
required for this storage server, verify it is installed.
Failed to create storage server <storage_server_IP>, host is
unreachable
```

To resolve this issue, install the Direct-IO plug-in on both the NetBackup client and the NetBackup media servers.

See "Installing the Direct-IO plug-in on NetBackup media server" on page 19.

See "Installing the Direct-IO plug-in on NetBackup client" on page 20.

The `tpconfig` utility cannot find SpanFS cluster

This issue occurs when the required storage server has not been created before running the `tpconfig` command.

The `tpconfig` is a NetBackup utility used to configure and manage tape devices and OpenStorage (OST) storage servers. When working with the storage, `tpconfig` helps register the storage server with NetBackup by specifying credentials, server type, and certificate paths.

Cause:

The most common reason for this error is that the storage server was not created using the `nbdevconfig` command. Without this step, `tpconfig` cannot retrieve the necessary storage server information.

Error message:

```
# tpconfig -add -storage_server <storage_server_IP> -stype SpanFS
-sts_user_id <user> -password <password> -ca_file_path /tmp/ca.pem
Failed to obtain storage server information to <storage_server_IP>:
 Error = 2050200
Authorization failed for OpenStorage server <storage_server_IP>
```

Resolution:

Ensure that the storage server is created using the following command before running `tpconfig`:

```
nbdevconfig -creatests -storage_server <storage_server_IP> -stype
SpanFS -media_server `hostname`
```

Once the storage server is properly configured, `tpconfig` should be able to connect and complete the setup successfully.

AIR import fails with status code 191 on target SpanFS cluster

After configuring Auto Image Replication (AIR) between two SpanFS clusters, you may encounter a status code 191 during import operations at the target domain. It indicates that the disk pool and disk volume are not properly configured to recognize the replication source.

Cause:

The target domain may be missing replication metadata. If the following command shows disk volumes without a **Replication Source** field, it means the disk pool and disk volume need to be updated:

```
nbdevquery -listdv -stype SpanFS -U
```

Example output with the missing replication source:

```
Disk Pool Name              : dp
Disk Type                   : SpanFS
Disk Volume Name            : lsu4
...
Num Repl Sources            : 0
Replication Source          : [missing]
```

Resolution:

To resolve the issue, update both the disk volume and disk pool using the following steps:

1. Identify the disk pool and disk volume names:

   ```
   nbdevquery -listdv -stype SpanFS -U
   ```

2. Update the disk volume:

   ```
   nbdevconfig -updatedv -dv <disk_volume_name> -dp <disk_pool_name>
   -stype SpanFS
   ```

3. Update the disk pool:

   ```
   nbdevconfig -updatedp -dp <disk_pool_name> -stype SpanFS
   ```

After these steps, run the `nbdevquery` command again to verify the update. You should now see the **Replication Source** field populated, indicating successful configuration.

Example updated output:

```
Disk Pool Name              : dp
Disk Volume Name            : lsu4
```

```
...
Num Repl Sources           : 1
Replication Source         : <source_cluster>:<volume_name>
```

The client direct backup fails with status code 50 due to client memory exhaustion

When running a backup job, you may observe a failure with status code 50: client process aborted in the Activity Monitor. This issue is often caused by insufficient memory on the client system, leading to the termination of the `nbostpxy` process by the operating system.

Check for the following on the Activity Monitor:

```
Info nbjm: started backup job for client <client_name>, policy
<policy_name>, schedule Full on storage unit <storage_unit>
Info bpbrm: started process
Info bpbrm: connecting
Info bpbrm: connected
Info bpbrm: begin writing
Critical bpbrm: unexpected termination of client <backup_id>
Error bptm: media manager terminated by parent process
Info dbclient: done. status: 50: client process aborted
```

Check `/var/log/messages` for out of memory events. You may see entries like:

```
Out of memory: Killed process <PID> (nbostpxy) total-vm:<value>kB,
anon-rss:<value>kB, ...
```

Run free -g on the client. If the output shows zero free memory, it confirms that the system is under memory pressure:

```
              total        used        free      shared  buff/cache
   available
Mem:             15          12           0           0           2
           1
Swap:             7           6           1
```

Cause:

The `nbostpxy` process allocates a large write buffer (default: 512 MB per instance). On systems with limited memory, this allocation can lead to excessive memory usage and OOM conditions, causing intermittent backup failures.

Resolution:

To resolve this issue, reduce the size of the write buffer that is allocated by `nbostpxy`. By default, each `nbostpxy` instance allocates 512 MB for writing. You can lower this value to prevent excessive thrashing by setting the `maxWriteBufferSizeMB` and `maxPendingWriteDataSizeMB` parameters in the `spanfs.conf` file.

For example:

```
{
    "logLevel": 5,
    "logSize": 30,
    "logPath": "/var/log/spanfs",
    "writeParams": {
        "maxWriteBufferSizeMB": 16,
        "maxPendingWriteDataSizeMB": 8
    }
```

The target LSU is not listed during duplication configuration

When configuring duplication between SpanFS Logical Storage Units (LSUs), you may encounter a situation where the target LSU is not listed in the configuration interface or command output.

Cause:

This issue occurs when the target LSU resides in a different NetBackup domain than the source LSU. For duplication to work between SpanFS LSUs, both LSUs and by extension, their associated DataProtect clusters must be part of the same NetBackup domain.

Resolution:

Ensure that:

- Both the source and the target LSUs are configured within the same NetBackup domain.
- The DataProtect clusters hosting these LSUs are properly registered and recognized by the NetBackup master server.

Once the LSUs are aligned within the same domain, the target LSU should appear during duplication setup.

The files and folders restore fails for original location

When attempting to restore files and folders to their original location in a guest VM, the operation may fail due to an unsupported root file system on the client.

Cause:

The root file system on the client is **Btrfs**, which is not supported for mount point mapping during restore operations. As a result, NetBackup maps files using the device path (for example, /dev/sda3/...) instead of the actual mount point. Since the restore process cannot create files directly on device paths, the original location restore fails.

In the Activity Monitor, you may see:

```
Error bpVMutil: Failed to restore the selected files and folders.:
SYM_VMC_CONNECT_ERROR
Info bpVMutil: Finished restore operation in guest VM.
Info bpVMutil: Successfully deleted VM <temporary_VM_ID>
VMware policy restore error(2820)
In the nbtar_rt log (if AGENTLESS_KEEP_STAGING_LOCATION = 1 is set
in bp.conf), you may find:
ERR - Failed to create parent directory /dev/sda3/sanyukta. [Not a
directory]
```

Resolution:

Perform an alternate location restore instead of restoring to the original location. This bypasses the unsupported device path issue and allows the restore to complete successfully.

Duplication fails with "operation not supported" error

While configuring duplication Storage Lifecycle Policy (SLP) for VMware backup policies, the job may fail with an error indicating that the operation is not supported.

Cause:

This issue occurs when the source and target are located in separate SpanFS storage domains. Duplication between LSUs is only supported when both are within the same NetBackup domain.

Job details:

```
begin writing
Critical bpdm: sts_copy_extent failed: error 2060016 operation not
supported
Critical bpdm: image copy failed: error 2060016: operation not
supported
Error bpdm: cannot copy image from disk, bytesCopied = 0
requesting resource <storage_unit_name>
```

Resolution:

Ensure that both the source and the target LSUs are:

- Part of the same NetBackup domain
- Hosted within the SpanFS clusters that are properly registered and recognized by the NetBackup master server.

If the LSUs are in separate domains, duplication using SLP is not supported and results in failure.

The client direct backup fails due to client data streaming patterns

In some cases, BMR and file system backup failures may occur due to the data streaming behavior of the client's file system.

During an initial backup, all data is considered new and is written to storage. However, as the backup policy continues to run, the combination of schedule settings and acceleration attributes determines how much data is new and how much is referenced.

If most of the data is referenced, or if the access pattern is unevenly distributed, it may appear that no new data is being written. This perceived lack of progress can lead to time-outs, causing the backup to terminate prematurely.

Error message:

```
Info bpbkar (pid=11936) done. status: 14: file write failed
Error nbpem (pid=921441) backup of client <host name> exited with
status 14 (file write failed)
```

Example scenario:

In a Windows BMR backup using ALL_LOCAL_DRIVES (for example, C and E), a large portion of data may be written in the first few seconds. After that, the remaining data is mostly referenced. When drive E is snapshotted, another burst of writes occurs. Over time, these "includes" accumulate, contributing to the stalled state.

Workaround:

To resolve this issue, increase the retry limit for stalled data streaming by modifying the `bp.conf` configuration file. Specifically, adjust the `OST_CD_BUSY_RETRY_LIMIT` parameter to allow more retries before the process fails.

This overrides the default limit of 500 retries, giving the backup process more time to recover from stalled states.

The client direct backup fails with status code 83: media open error

When running a backup job, the process may fail with status code 83: media open error in the Activity Monitor.

In the Activity Monitor, you may observe the following error messages:

```
Error  nbpem (pid=1710263) backup of client <client_name> exited with
 status 83 (media open error)
end writing
Critical bpbrm Client Direct plugin not available (2060009)
Info bpbkar (pid=0) done. status: 83: media open error
```

To resolve this issue, install Direct-IO plug-in on the client and run the backup job again.

Troubleshooting missing subnet whitelist for SpanFS access

NetBackup needs network access to the SpanFS file system for Direct-IO operations. If the required IP subnet isn't added to the SpanFS subnet whitelist, SpanFS rejects the RPC requests, causing backups to fail.

Issue:

When the subnet whitelist is missing or incorrect, backups fail with status 83 and logs show errors such as:

```
image open failed: error 2060037: access not allowed
```

Resolution:

To resolve this issue, add the required subnet to the SpanFS whitelist.

```
iris_cli cluster add-subnet-to-whitelist subnet-ip=<IPv4 address>
subnet-mask-ip4=<IPv4 subnet mask>
```

See “Configuring SpanFS subnet whitelists for NetBackup access” on page 24.

Duplicate replication target volumes displayed in multi-VLAN Direct-IO configurations

In a multi-VLAN Direct-IO setup, a target volume (LSU) may appear multiple times while adding replication targets, with each entry associated with different network interfaces of the same storage server. Although each LSU is associated with only one storage server, the UI may display multiple entries, and the displayed storage server association is not always correct.

Before selecting a replication target, verify that the volume is associated with the correct storage server. Avoid selecting duplicate entries that are not mapped to the intended storage server.

Appendix A

Performance tuning and optimization

This appendix includes the following topics:

NetBackup concurrency control

To achieve optimal performance from a NetBackup storage server, it's essential to manage how many jobs can run concurrently. NetBackup provides several configurable settings that control job concurrency across disk pools, storage units, policies, and clients. Proper tuning of these parameters ensures balanced resource usage, prevents bottlenecks, and improves overall backup and restore efficiency.

- Disk Pool I/O Stream Limit
  Controls the maximum number of concurrent I/O operations (streams) allowed on a disk pool. Each backup/restore job counts as 1 stream.
  Each optimized duplication (opt-dup) or AIR job counts as 2 streams.
- Storage unit (STU) maximum concurrent jobs
  Limits how many jobs can run concurrently on a specific storage unit. Ensure the total concurrent jobs across all STUs using the same disk pool do not exceed the disk pool's I/O stream limit.
- Limit jobs per policy
  Restricts the number of concurrent jobs that can run from a single backup policy. It prevents a single policy with many clients from consuming all available resources.

- Maximum jobs per client
  Limits how many concurrent jobs can run for a single client.

Recommendations

SpanFS permits, also known as Nebula permits, are used to evenly distribute jobs across SpanFS nodes and prevent resource oversubscription, ensuring optimal performance. By default, each NetBackup data-related job consumes a fraction of a data path permit. Upon successful job completion, the acquired permit is explicitly released. However, if a related process such as **bptm**, **bpdm**, or **nbostpxy** crashes before the release call is made, the permit may not be freed immediately. In such cases, garbage collection of the permit takes longer, potentially impacting resource availability.

Permit information

Permits are calculated based on available resources and allocated resources for a stream. The following are some available permits:

- Each SpanFS physical node, including C5K, C6K and C8K, has 24 permits.
- VMRobo has 12 permits.
- Each VM edition node has 9 permits.

For any backup or recovery job, when a permit is granted, the SpanFS cluster returns a virtual IP from its available resources to the NetBackup job. This virtual IP, which represents a specific node in the cluster at that moment, is used by the NetBackup job to read from or write data to the cluster, depending on the operation.

Supported Direct-IO configuration

Configuration rules:

- One LSU can only be used by one NetBackup domain.
- One or more than one LSU can be created from the same SpanFS storage domain that is also known as a view box.
- One view box is a single, independent deduplication pool (deduplication domain).
- NetBackup LSU and NetBackup disk pool are 1:1 mapping.
- One or more than one LSU may be created from the same view box.

Some typical configurations:

- Two NetBackup domains use two LSUs created from the same view box of a SpanFS cluster. This can achieve data deduplication for backups from two different NetBackup domains.
- One NetBackup domain uses more than one LSU created from the same view box of a SpanFS cluster.
- One NetBackup domain uses more than one LSU created from more than one view box's of different SpanFS clusters.

View box and job concurrency in SpanFS clusters

It is recommended to configure one View box per cluster. For illustration purposes, we consider a 4-node SpanFS cluster.

When the number of concurrent jobs is less than 96, each job can use the maximum allowed memory, ensuring throughput-optimal performance for each job. Once the number of concurrent jobs exceeds 96, each job will consume less than the maximum memory allowed to accommodate the increased load.

Each node can handle up to 480 concurrent jobs, allowing a 4-node cluster to support up to 1,920 concurrent NetBackup jobs. If additional nodes are added to the cluster, job concurrency scales linearly with the number of nodes.

Managing disk pool limits for high concurrency

If a user anticipates higher NetBackup job concurrency than the cluster can support, it's essential to configure NetBackup disk pool limits appropriately for all disk pools created from the same SpanFS cluster.

Conservative approach:

- The sum of disk pool limits should be less than or equal to the cluster’s maximum job concurrency.
- To account for rolling upgrades or temporary node failures, disk pool limits should be adjusted downward accordingly.

Aggressive approach:

- If the user does not expect all disk pools to reach their maximum job limits simultaneously, the combined disk pool limits may exceed the cluster’s maximum concurrency.
- However, if actual job concurrency exceeds the cluster’s capacity, some jobs may fail due to lack of permits.

Refer to the *Best practices: Disk pool configuration - setting concurrent jobs and maximum I/O streams* topic of the *NetBackup Backup Planning and Performance Tuning Guide* for the general principles that are applicable to NetBackup Direct-IO.

Optimizing CPU and memory usage in NetBackup Direct-IO

The Direct-IO plugin allows high-bandwidth backups while offering control over system resource usage. By tuning specific configuration parameters, administrators can manage CPU and memory consumption per backup job, especially useful in resource-constrained environments.

To control how much CPU each backup job consumes, adjust the `maxConcurrentWrites` parameter. The following table summarizes how this parameter affects CPU usage:

Table A-1

Parameter	Description
`maxConcurrentWrites`	Controls the number of concurrent write operations per backup job. ■ Lower the value to reduce CPU usage per job. ■ Set to 1 to limit usage to less than one CPU core. ■ Actual usage depends on deduplication rate, network latency, and SpanFS responsiveness.

The following parameters control memory allocation during the write operations:

Table A-2

Parameter	Description
`maxWriteBufferSizeMB`	Specifies the maximum size of the write buffer. This buffer temporarily holds data before it is written to the storage backend. Allowed values: 1 – 2000 MB Default: 512 MB

Table A-2 *(continued)*

Parameter	Description
`maxPendingWriteDataSizeMB`	Defines the maximum amount of data that can be queued in memory before being added to the write buffer. Allowed values: 1 – 1000 MB Default: 128 MB

Configuration example:

```
/usr/openv/lib/ost-plugins/spanfs.conf
{
  "writeParams": {
    "maxWriteBufferSizeMB": 512,
    "maxPendingWriteDataSizeMB": 128
  }
}
```

To estimate the total memory usage for write operations, use the following formula:

Total memory usage (MB) = `maxWriteBufferSizeMB` + `maxPendingWriteDataSizeMB`

For example, with the above configuration, total memory usage = 512 + 128 = 640 MB

This value represents the peak memory that may be consumed during intensive write operations. It is recommended to adjust these values based on available system memory and workload characteristics.

While these configuration changes help conserve system resources, they may also affect backup performance. Consider the following tradeoffs when applying these settings:

- Lower resource usage typically results in longer backup times.
- These settings are best suited for environments where minimizing CPU and memory usage is more critical than backup speed.

Index