NetBackup™ for MongoDB Ops Manager Administrator's Guide

Release 11.2

VERITAS
now part of Cohesity

NetBackup™ MongoDB Ops Manager 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

Introduction

This chapter includes the following topics:

About MongoDB Ops Manager

MongoDB Ops Manager is an essential tool for leveraging MongoDB databases. It simplifies database management through real-time monitoring and robust backup solutions. The Ops Manager Automation enables you to configure and maintain MongoDB nodes and clusters.

NetBackup allows you to monitor, back up, and recover the assets from MongoDB Ops Manager portal. The cluster which needs to be protected using NetBackup should be marked as **Third Party Managed** via Ops Manager. The policy in NetBackup supports Full backup, Differential Incremental backups, and Transaction Log (TLOG) backups, as defined in NetBackup policy.

Backup type naming across NetBackup versions

The naming and availability of MongoDB Ops Manager backup types differ across NetBackup releases.

- In NetBackup 11.1 and earlier, MongoDB Ops Manager supported **oplog backups**, which were labeled as **Differential Incremental backups**. Incremental backups were not available.
- In NetBackup 11.2 and later, **incremental backups** are available and are labeled as **Differential Incremental backups**.

Oplog backups are labeled as **Transaction Log (TLOG) backups**.

When reviewing policies, backups, or restore points created across different NetBackup versions, be aware of these naming differences.

Architecture overview

NetBackup supports the following for MongoDB Ops Manager.

- Discovery: NetBackup supports auto-discovery of clusters with protection enabled.

  Note: If the clusters are not displayed, run **Discover Clusters**.

  Navigate to **MongoDB Ops Manager** workload > **Clusters** tab > **Discover Clusters**.

- Backup: NetBackup allows cluster backups, periodically, using policy.
  - Backup types: **FULL, INCR, and TLOG**
  - **FULL** (Full snapshot backup): Creates a complete base snapshot backup of the selected replica set or sharded cluster through MongoDB Ops Manager. A FULL backup is required as the base for subsequent incremental backups.
  - **INCR** (Incremental snapshot backup): Backs up only changed data since the last FULL backup. INCR uses MongoDB Ops Manager APIs to identify which files changed between snapshots and transfers only that changed data.
  - **TLOG** (Transaction log / oplog backup): Backs up the MongoDB oplog (transaction log) only. TLOG is separate from INCR and is used to support transaction-log based recovery objectives (for example, point-in-time recovery).

    Note: Point-in-time (PIT) restore requires Transaction Log (TLOG) backups. FULL and INCR (snapshot-based incremental) backups alone do not support PIT restore.

- Restore: NetBackup allows Full restore and restore to a different cluster. Point-in-time (PIT) restore is supported only when TLOG backups are available for the selected time range.
- NetBackup API and Web GUI

- Asset service
- Discovery API
- Credential management
- Protection through policy
- Recovery point service

MongoDB protection workflow - Backup architecture

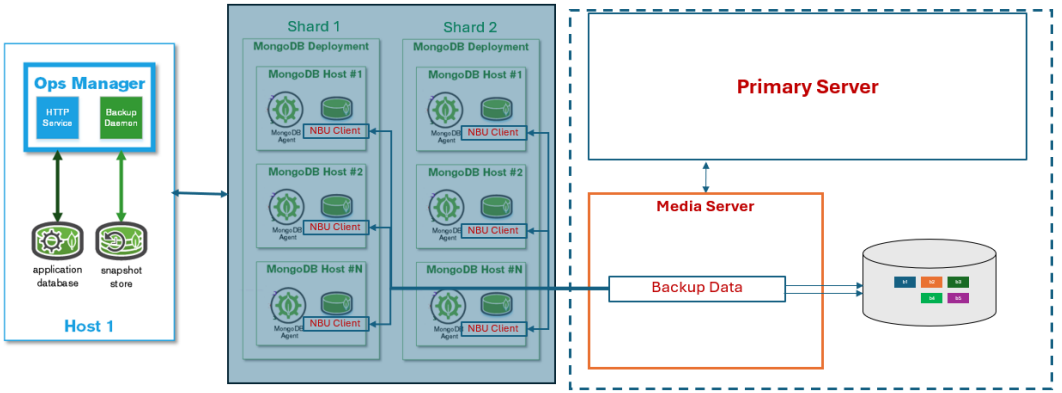

Overview of NetBackup for MongoDB Ops Manager

Table 1-1

Steps	Action	Description
1	Log into NetBackup Web UI as the Default MongoDB Ops Manager	This role has all the permissions that are necessary to manage MongoDB Ops Manager functionalities and assets to back up those assets using policies. **Note:** For more information, see the *Default MongoDB Ops Manager role* in *Default RBAC* roles section of NetBackup web UI Administrator's Guide.
2	Add MongoDB Ops Manager workload	See "Add workload using Central URL" on page 13.
3	Permission	See "Manage credentials" on page 12.
4	Preferred nodes	See "Clusters and preferred node" on page 14.

Table 1-1 *(continued)*

Steps	Action	Description
5	Policies	See "Policies" on page 17.
6	Recovery	See "MongoDB Ops Manager cluster recovery" on page 21.

Prerequisites and best practice

Before proceeding, ensure that the following prerequisites are met.

- In MongoDB Ops Manager, the central URL is a configuration setting that specifies the fully qualified domain name (FQDN) of the Ops Manager application. This URL is provided by MongoDB Ops Manager.

  When configuring credentials

  - Public key: The public key acts as a username during the API requests.
  - Central URL: As mentioned above.
  - Private key: The private key functions as a password and should be kept secure, like any other sensitive credential.
  - Authentication type: When using SSL enabled authentication, ensure that the CA certificate is configured in Manage credentials.
- Third-party APIs allows you to integrate with various external services to enhance monitoring, alerting, and overall management of your MongoDB deployments.

  Note: Contact MongoDB to enable API services.

- Automation, backup, and monitoring are enabled on MongoDB servers.

Prerequisites specific to INCR and TLOG backup types

- **INCR Backup**: A FULL backup must exist before you can run an INCR backup. INCR backups are tied to the base FULL snapshot.
- **TLOG Backup**: Ensure the environment is configured for oplog/transaction log collection in Ops Manager (for example, the oplog/TLOG dump location is configured as part of the Ops Manager cluster configuration).
- **Preferred nodes**: Configure at least one preferred node per replica set (and at least one per shard replica for sharded clusters). Preferred nodes are used during snapshot operations, incremental diff calculation, and transaction log capture workflows.

- **PIT Restore prerequisite**: To perform point-in-time restore, configure TLOG schedules and ensure TLOG backups are retained for the required recovery window. PIT restore is not supported with FULL + INCR only.

Platform and deployment prerequisites

Before integrating NetBackup with MongoDB Ops Manager, ensure that the environment meets the following requirements:

- MongoDB Ops Manager version **8.0.21 or later** is installed in the project that will use third-party backup.
- MongoDB deployments must use the **Enterprise edition**.
- The Feature Compatibility Version (FCV) of the MongoDB deployment must be **4.2 or later**.
- System clocks across all hosts should be synchronized (for example, using NTP) to ensure accurate timestamps and point-in-time recovery behavior.
- The cluster to be protected must be the enterprise version.
- The cluster to be protected must be either a replica set or a sharded cluster.

These prerequisites help ensure stable operation and prevent issues during backup and recovery workflows.

Ops Manager configuration for third-party backup

To use NetBackup as a third-party backup solution, MongoDB Ops Manager must be configured to allow third-party managed backups.

The following configuration steps are required in Ops Manager:

- Enable the third-party backup feature using the supported feature flag.
- Configure a directory that Ops Manager agents can use to store transaction log (TLOG) data.
- Ensure that the MongoDB agents running on all cluster nodes have read and write access to this directory and that sufficient disk space is available.

Operational considerations for third-party backup

When using NetBackup with MongoDB Ops Manager, consider the following operational guidelines:

- Use **only one backup solution per cluster**. Do not run native Ops Manager backup and third-party backup on the same cluster.

- Do not initiate snapshot backups during a restore operation. Start snapshot backups only after the restore is complete and the cluster is running normally.
- At least one node should be selected as a preferred node within a replica set/shard.
- Do not use multiple third-party backup tools on the same cluster.
- Monitor disk usage for the transaction log directory to ensure that sufficient space is available.
- If the transaction log directory is misconfigured or becomes unavailable, errors such as missing transaction logs may occur, which can reduce or break the point-in-time recovery window.
- Granular recovery is not supported using third party APIs.
- Following the recovery, MongoDB cluster is reverted to the state of the chosen backup images.
- For restore operations, install the NetBackup client on all nodes of the cluster where the restore will be performed. This applies to both the original cluster and any alternate cluster used for restore.
- Ensure that the NetBackup clients are installed on all the nodes of the target cluster.
- If TLOG backups are not configured or have gaps, PIT restore cannot be performed and recovery is limited to the available snapshot restore points (FULL/INCR).

Feature flags for NetBackup integration

The following feature flags are relevant for integrating NetBackup with MongoDB Ops Manager:

- **mms.featureFlag.backup.thirdPartyManaged**
  This is the required and supported feature flag for enabling third-party backup integration.

  Note: The flag value is set to **controlled**.

- **brs.thirdparty.baseOplogFilePath**
  This setting defines the directory used by the MongoDB agent to store transaction log (TLOG) data. This setting is required for point-in-time recovery scenarios.
- **brs.thirdparty.extraLoggingEnabled**
  This is an optional setting used only for troubleshooting and debugging purposes.

Note: The feature flag **mms.featureFlag.backup.thirdPartyWithManagedOplog** is deprecated and has no effect on Ops Manager behavior. This flag is not required for NetBackup integration and should not be used.

For more information, see the MongoDB Ops Manager documentation on enabling third-party backup.

Snapshot and transaction log (TLOG) considerations

For stable backup and recovery behavior, follow these general guidelines:

- Perform a full backup at least once a week.
- Perform incremental backups at least once a day.
- Once a full backup is started, earlier backups cannot be used as the source for subsequent incremental backups, even if the full backup fails.
- Transaction log (TLOG) backups should be taken frequently to maintain a reliable point-in-time recovery window.
- During initial setup or when transaction log gaps are detected, take a transaction log backup followed by a full backup.
- Snapshot creation requires a running cluster with stable data. Snapshots cannot be taken while a restore operation is in progress.

These practices help ensure reliable recovery and minimize the risk of data loss.

Chapter 2

Add MongoDB Ops Manager workload

This chapter includes the following topics:

Manage credentials

Credential management lets you centrally manage the credentials that NetBackup uses to access MongoDB Ops Manager workloads. You can manage NetBackup credentials from Credential management.

Steps to add credentials.

1. Select **Credential management** on the left pane. Credential management page is displayed.
2. Select **Add** . Add credential dialog box is displayed.
3. Select **NetBackup** option as **Credential store** and then select **Start**.
4. Configure the required values for the following tabs:
   - **Basic properties**: Specify the Credential name and other values.
   - **Category**: Configure the appropriate authentication type.

Note: Prerequisites such as a public key, private key, and Central URL are mandatory when specifying the details using *Add credentials* option.

For more information, See "Prerequisites and best practice" on page 8.

- **Permissions**: In this tab, select **Add** and then select the following:
  - Role
  - Permission: The following are the available options:
    - View
    - Create
    - Update
    - Delete
    - Manage Access
    - Assign credentials
- **Review**: Verify the configuration and select **Finish**.

Add workload using Central URL

The central URL in MongoDB Ops Manager is the address you use to access the Ops Manager interface.

This URL is configured during the initial setup and can be updated later, if required. It generally appears as follows:

http://<hostname>:8080

or

https://<hostname>:8443 depending on your configuration.

You can also find the value at the bottom of the Deployment page of the project.

1. Log on to **NetBackup Web UI**.

   Note: If MongoDB Ops Manager workload is not visible, refer to the *Default RBAC roles* section in *NetBackup Web UI Administrators Guide*.

2. Navigate to **MongoDB Ops Manager** under **Workloads** on the left pane.
3. Click **Add**. Add MongoDB Ops Manager page is displayed.

4. Specify the **Central URL** and then click **Next**.
5. Post URL validation **Permissions** page is displayed.
6. Click **Manage credentials**. Credentials page is displayed with the following two options:
   - **Select from existing credentials**. If you have already defined credentials, select this option. NetBackup displays the list of existing credentials that you may define using Credential management. Select appropriate credentials and click **Next**.
   - **Add credentials**. Click Add credentials if the credentials are not define and follow the instructions. For more information, See "Manage credentials" on page 12.
7. Click **Finish**. NetBackup validates the specified configuration and displays success message.

Note: If cluster discovery fails after the validation step is completed, the discovery status and failure details are not displayed on the discovery page. This information is recorded in the Events section.

Check the **Security** >> **Security Events** section to view the discovery outcome and failure details.

Clusters and preferred node

On completion of Credentials validation, auto discovery process is initiated. Once the auto discovery process is complete, all the organization, projects and clusters within Ops Manager are discovered and added as Assets. When you click on a particular asset, you will find all the details about Ops Manager, Organization, Project, and replica sets/shards within the cluster.

Steps to configure preferred node.

1. Click **MongoDB Ops Manager** workload on the left pane.
2. Click **Clusters** tab. A list of auto discovered clusters is displayed.

   Note: The clusters and nodes which are displayed as **Inactive** are either deleted or removed from Third Party managed at Ops Manager.

   ---

3. Click the required cluster for which you want to configure the preferred node. The discovered node page is displayed.

Note: Auto discovery is an asynchronous process that takes time to retrieve and push clusters and nodes details into the assets. This process may cause delay in displaying the details.

4. Click **Actions**, displayed on the right side for the respective node.
5. Click **Select preferred nodes**. A Select preferred node dialog box is displayed.

Note: For replica set, select at least one node as the preferred node for the cluster to be protection ready.

For Sharded, select at least one preferred node from each shard replica.

Arbiter node will not be displayed.

To ensure successful backups, update the policy that contains the cluster after updating the preferred nodes.

6. Select the required preferred node or nodes. Click **Next**.
7. Further, you can also arrange them as per your requirements. Click **Save**.
8. A success message is displayed.

Chapter 3

Backup

This chapter includes the following topics:

MongoDB Ops Manager configuration file settings

You can create a configuration file in NetBackup at the specified location on the NetBackup primary server. These settings will be applicable to all the MongoDB Ops Manager backups/restores.

Backup type naming across NetBackup versions

The naming and availability of MongoDB Ops Manager backup types differ across NetBackup releases.

- In NetBackup 11.1 and earlier, MongoDB Ops Manager supported **oplog backups**, which were labeled as **Differential Incremental backups**. Incremental backups were not available.
- In NetBackup 11.2 and later, **incremental backups** are available and are labeled as **Differential Incremental backups**.
  Oplog backups are labeled as **Transaction Log (TLOG) backups**.

When reviewing policies, backups, or restore points created across different NetBackup versions, be aware of these naming differences.

File name and location

File name: `mongodbopsmanager.conf`

File location:

- Windows: `Install_Dir\NetBackup\var\global\`
- Unix: `/Install_Dir/var/global/`

Parameters

Table 3-1 Configurable parameters in mongodbopsmanager.conf file

Configurable parameter	Description	Note
NO_OF_BACKUP_STREAMS_PER_NODE	Number of backup streams per node. Maximum 32 streams per node can be triggered. Default value: 1, which is single stream per node.	A new stream will be created for meta data backup which will be an additional stream despite this value.
SNAPSHOT_STATE_CHECK_TIMEOUT_IN_SEC	The timeout value during snapshot state check operations in seconds. Default value: 120 seconds.	
SNAPSHOT_TIMEOUT_IN_MINUTES	The timeout value during which complete snapshot backup will be performed in minutes. Default value: 360 minutes.	
RESTORE_TIMEOUT_IN_MINUTES	The timeout value during which complete restore and recovery will be performed in minutes. Default value: 720 minutes. The value can be re-configure to accumulate larger data set to restore and recover.	

Example

```
NO_OF_BACKUP_STREAMS_PER_NODE=3
SNAPSHOT_TIMEOUT_IN_MINUTES=120
SNAPSHOT_STATE_CHECK_TIMEOUT_IN_SEC=300
RESTORE_TIMEOUT_IN_MINUTES=300
```

Policies

NetBackup allows you to configure **FULL**, **INCR**, and **TLOG** backups for MongoDB Ops Manager assets using policies and schedules.

Note: For policies created before NetBackup 11.2, a **Differential Incremental** schedule runs as an Incremental backup after upgrade.

To continue oplog backups (TLOG), update the policy and select **Transaction log (TLOG)** as the backup type.

Warning: When defining policies for the same cluster, ensure that backup jobs of the same type do not run at the same time.

- Use FULL schedules to create the base snapshot.
- Use INCR schedules between FULL backups to transfer only changed data since the last FULL backup.
- Use TLOG schedules to back up transaction logs (oplog) independently from snapshot backups.

Warning: FULL and INCR schedules do not support point-in-time (PIT) restore. To meet PIT recovery objectives, configure and retain TLOG schedules.

Here are the key components of a NetBackup policy:

- Policy attributes
- Schedules
- MongoDB Ops Manager Assets

By configuring these components, NetBackup ensures that your data is backed up according to your organization's requirements and recovery objectives.

How INCR incremental backups works

INCR backups are snapshot-based incremental backups that avoid re-backing up unchanged data. NetBackup uses the following Ops Manager APIs during INCR processing:

- **File list API**: Retrieves the list of files (and sizes) contained in the incremental snapshot for a node. NetBackup uses this to understand what is present in the snapshot before backing it up.
- **Incremental file diff API**: Identifies only the files that changed between the previous snapshot and the current incremental snapshot. NetBackup uses this "diff" to back up only changed data instead of the full snapshot contents.

Result: INCR runs are typically faster and smaller than FULL backups because unchanged data is not re-read and re-transferred.

Limitation: Incremental (INCR) backups under specific load conditions

Under certain timing and workload conditions, incremental (INCR) backups taken using MongoDB Ops Manager with NetBackup 11.2 and later may be incomplete.

This behavior can occur when new data is created very shortly after a snapshot is taken (for example, within a few seconds). The likelihood increases in multi-node replica set environments that are under heavy load or have a high rate of data change. In such cases, one or more collections may not be included in the incremental backup.

Although the backup job may complete successfully, restores that rely on the affected incremental backups can result in missing or inconsistent data.

If an incremental backup is affected, all subsequent incremental backups that depend on the same backup chain may also be unreliable. There is currently no mechanism to detect this condition during backup execution. The issue can only be identified during restore validation.

Based on current analysis:

- This issue applies only to block-level incremental (INCR) backups introduced in NetBackup 11.2 and later
- Full (FULL) backups are not affected Transaction log (TLOG) backups are not affected
- The issue is less likely in single-node deployments

This issue is currently under investigation with MongoDB.

For workaround and guidance , See the section called “Limitation: Incremental (INCR) backups under specific load conditions” on page 25.

Add a policy

This section provides information about how to configure a policy to backup the MongoDB Ops Manager assets.

Upgrade behavior for MongoDB Ops Manager backups

If you upgrade to NetBackup 11.2 or later, existing MongoDB Ops Manager policies remain available.

However, backup behavior changes for policies that were created before the upgrade.

- Policies that used Differential Incremental backups now run as Incremental (INCR) backups.
- Transaction log (oplog) backups do not continue automatically after the upgrade.

Required action:

To continue transaction log–based backups and point-in-time recovery, you must manually edit existing MongoDB Ops Manager policies and change the backup type to **Transaction Log (TLOG)**.

If this change is not made, backups may run successfully, but transaction log recovery will not be available.

1. On the left pane, click or expand **Protection** and then click **Policies**.
2. Click **Add**. Create policy page is displayed.
3. Specify the details for the following tabs:
   - **Attributes**: Select **BigData** from the Policy dropdown and then select **MongoDB Ops Manager** from the workload type dropdown.
   - **Schedule**: On this tab, configure all the settings required to schedule backup. You can also select **Full backup, Differential incremental backup, or Transaction log backup**.
     - **Full backup**: Creates a complete copy of all data for the selected cluster at a specific point in time.
     - **Differential incremental backup**: Backs up only the data that has changed since the last full backup.
     - **Transaction log backup**: Backs up transaction log data so that changes made after a backup can be captured.

     Note: Operational warning: Point-in-time (PIT) restore is supported only with Transaction log (TLOG) backups. PIT restore is not supported with Differential incremental backups (INCR) alone.

   - **MongoDB Ops Manager Assets**: On this tab, a list of clusters which are protection ready is displayed. Select the required cluster.
4. Click **Create**.

Warning: If the preferred nodes are shared with multiple clusters, adding multiple clusters in the policy may lead to backup failures.

Chapter 4

Recovery

This chapter includes the following topics:

MongoDB Ops Manager cluster recovery

Recovery typically refers to the process of restoring data and service after a failure or disaster. Ensure that the Configuration parameters are configured for recovery. See “MongoDB Ops Manager configuration file settings” on page 16. for more details.

Note: Taking a snapshot while a restore operation is in progress is not supported.

During a restore, the database process is stopped and existing data on the cluster is removed before data from the selected snapshot is copied back. At this time, there is no stable data available on disk to create a snapshot.

Any snapshot request started during a restore will fail or may remain stuck and will not generate a Snapshot ID. Snapshot requests are not queued or retried automatically.

Start snapshot operations only after the restore has completed successfully and the cluster is fully running and healthy.

PIT restore preconditions

TLOG backups must be configured and available for the target time.

If TLOG backups are missing or contain gaps for the selected interval, PIT restore cannot be performed and recovery is limited to snapshot restore points.

1. Log on to NetBackup Web UI.

2. On the left pane, select **Recovery**.
3. Click **Start recovery** from the Regular recovery option.

   Specify the Basic properties details as follows:

   - Select MongoDB Ops Manager as the **Policy** type.
   - Select Normal backups as **Restore** type.
   - Click **Select cluster**. Source clusters list is displayed. The list displays the clusters which are discovered and added as asset. Select the required source cluster to recover and then click **Next**.

     Note: If a cluster has been renamed after taking a backup, new name of the cluster is displayed.

4. Specify the **Recovery details** such as date range, files, and folders.
5. Specify the **Recovery options** such as Restore target, MongoDB options, and number of streams. The default stream is configured to One.

   Note: Restore to an alternate cluster is supported only when both the source and target clusters are managed by the same MongoDB Ops Manager. Restoring data across different Ops Manager hosts is NOT supported.

6. **Review** the recovery options provided and then click **Start recovery**.

   Note: Post recovery, the files and folders permission will be 777.

   Note: The '**max stream count per node**' provided by the user for recovery is the maximum number of streams that NetBackup is allowed to create per node for data transfer. The actual number of streams to create and amount of data allocated to each stream will be decided at runtime based on the total data size and the number of files. The actual number of streams will always be less than or equal to the max number specified by user. NetBackup has enforced a max stream count limit capped at 32 which is user can specify any value between 1 and 32.

Post-recovery considerations

During MongoDB restore operations, NetBackup creates additional directories under the `NetBackupCopiedOplogs` path. These directories are created for each restore job and are named using the restore job ID.

After the restore operation completes successfully, these directories are not removed automatically. If the restore job is completed and the data is no longer required, you can safely delete the directories related to those completed restore jobs.

Removing these directories helps free up disk space on the system.

Error codes and messages

For more information, refer NetBackup Status Code Reference Guide.

Chapter 5

Troubleshooting

This chapter includes the following topics:

Troubleshooting MongoDB Ops Manager issues

For more information about MongoDB Ops Manager troubleshooting, verify the following details:

- Verify the Job Details section of the job on the Activity monitor page for the failures.
- Verify the following logs:

  Auto-discovery log files (/usr/openv/logs)

  - ncfnbcs

  Backup log files (/usr/openv/netbackup/logs)

  - bpbrm
  - nbaapidiscv
  - bpbkar

  Recovery log files (/usr/openv/netbackup/logs)

  - bprd
  - nbaapireq_handler
  - tar

  Asset/Backup/Recovery API (/usr/openv/logs)

  - nbwebservice

Note: For API call failures, also check MongoDB Ops Manager logs such as `mms0.log` and `mms0-access.log`.

In earlier releases, transaction log backups were referred to as `oplog` backups. In this release, they are referred to as `TLOG (transaction log)` backups.

Limitation: Incremental (INCR) backups under specific load conditions

Workaround and guidance

- Avoid running incremental backups immediately after rapid data or schema changes
- Allow sufficient time for data to stabilize and synchronize across replica set nodes before starting incremental backups
- When possible, use FULL backups or Transaction log (TLOG) backups instead of relying solely on incremental backups until this issue is resolved

Errors and recommended actions

The following table describes the problem that might occur.

Table 5-1 Error and recommended actions

Error message or cause	Explanation and recommended action
Post changing the preferred nodes, running incremental backup process fails. Error: Backup of client exited with status 112.	This is a known issue with MongoDB, in case of any topology change. Take a incremental backup of <PrefNode1> node. Change the preferred node from <PrefNode1> to <PrefNode2> and again take a incremental backup. **Note:** The Ops Manager point-in-time restore might have oplog gaps in some situations. oplog is the MongoDB operation log. The backup type previously referred to as “oplog backup” is now explicitly named TLOG (transaction log). The backup fails with an error. Recommended action: ■ Take a full backup post topology changes.

Table 5-1 Error and recommended actions *(continued)*

Error message or cause	Explanation and recommended action
Running backup on old media server. Error: The backup selection fails with an error.	Cohesity recommends upgrading the media server to latest.
Run full and incremental backup on the RECOVERING cluster. Then recovery the cluster and run incremental backup. Warning: The executed incremental backup, post cluster recovery, is not accurate.	Cohesity recommends to run full backup post cluster recovery.

Chapter 6

NetBackup using API

This chapter includes the following topics:

- APIs to manage, protect, or recover MongoDB Ops Manager assets

APIs to manage, protect, or recover MongoDB Ops Manager assets

This section contains information of MongoDB Ops Manager APIs.

NetBackup uses third-party MongoDB Ops Manager APIs to support FULL, INCR, TLOG, and restore workflows. These include:

- File list API (incremental snapshot contents) and Incremental file diff API (changed files between snapshots) for INCR backups.
- Cluster and snapshot status APIs (for polling and lifecycle state checks).
- FULL snapshot lifecycle APIs (create/start/poll/finish/fail).
- INCR snapshot lifecycle APIs (create with base snapshot reference/start/poll).
- Oplog snapshot lifecycle APIs (create/start/poll/finish/fail/details).
- Restore lifecycle APIs (create/start/poll/fail).
- Preferred node discovery APIs and group/cluster discovery APIs.

The following sections are the part of this topic:

- Add a MongoDB Ops Manager
- MongoDB Ops Manager Discovery API
- Preferred nodes of replica sets/shard in a cluster
- Create a MongoDB Ops Manager Policy
- MongoDB Ops Manager Recovery

Add a MongoDB Ops Manager

Table 6-1 Add a MongoDB Ops Manager

API	Important variables and options

`POST /netbackup/asset-service/queries`
 	
`GET /netbackup/asset-service/queries/{aqcId}`
 	
`GET /netbackup/asset-service/workloads /mongodb-ops-manager/assets`
	■ `centralURL`: The central URL is a fully qualified URL with the port number of the Ops Manager Application.
	■ `credentialName`: credentials associated with MongoDB Ops Manager.
	Note: Ensure that the credentials with `credentialName` exist.

MongoDB Ops Manager Discovery API

Table 6-2 Discover the MongoDB Ops Manager assets for given central URL of Ops Manager

API	Important variables and options

`POST /netbackup/admin/discovery /workloads/mongodb-ops-manager/start`
 	
`POST /netbackup/admin/discovery/workloads /mongodb-ops-manager/stop`
 	
`GET /netbackup/admin/discovery/workloads /mongodb-ops-manager/status`
	■ `serverName`: central URL to initiate the discovery process.

After discovery, you need to configure the preferred nodes for the replica sets in the cluster.

Preferred nodes of replica sets/shard in a cluster

Table 6-3 Select preferred nodes of replica sets/shard in a cluster

API	Important variables and options
`POST /netbackup/asset-service/queries`	■ `assetType`: is cluster. ■ `id`: is the cluster ID for which the preferred node is to be configured. ■ `details`: specify names along with their preferred nodes, with each value corresponding to a node name. For example: `"details": {"rsId1": ["prefnode_hostname1:port1"], "rsId2": ["prefnode_hostname2:port2"] }`

Once the preferred nodes of all replica sets/shard in the cluster are configured, the **Protection Ready** flag for the cluster will be displayed as **Yes**. Proceed to create the policies

Create a MongoDB Ops Manager Policy

Table 6-4 Create a MongoDB Ops Manager Policy

API	Important variables and options
`POST /netbackup/config/policies`	■ `policyType`: BigData. ■ `clients`: Clients is a required parameter to have a hostname with a value corresponding to the cluster ID in MongoDB Ops Manager. ■ `backupSelections`: Should have the following: ■ `Application_Type`: mongodbopsmanager ■ Cluster Information: Specify in the format: `group/<group id>/clusters/<cluster id>` ensuring that the cluster ID and group ID match those in Ops Manager. ■ `Backup_Host`: Represents the hostname of the preferred nodes in the cluster, with an entry for each node. For example: `"backupSelections": { "selections": ["Application_Type= mongodbopsmanager", "/group/<groupId>/clusters/<clusterId>", "Backup_Host=preferredNodeHostName" ] }`

Other processes, like scheduling within the policy or initiating the policy backup, remain unchanged.

MongoDB Ops Manager Recovery

Table 6-5 MongoDB Ops Manager Recovery

API	Important variables and options
`POST /netbackup/recovery/workloads/ mongodb-ops-manager/scenarios/cluster /recover`	■ `sourceClusterAssetId`: NetBackup asset id of the cluster whose backup is to be recovered. ■ `targetClusterAssetId`: is NetBackup asset id of the cluster where the recovery will be performed. ■ `recoverPointInTime`: this parameter is not mandatory, however you can specify as `null` when recovering databases to the current time. ■ `maxStreamsPerNode`: is Maximum number of backup streams per node. **Note:** If the value for `recoverPointInTime` parameter is not provided, the recovery of the databases will happen to current time. `recoverPointInTime` requires TLOG backups for the specified time range. If TLOG backups are not available, point-in-time recovery is not supported and the request may fail or be limited to snapshot recovery points.