

NetBackup™ for MongoDB Administrator's Guide

Release 11.2

NetBackup™ for MongoDB Administrator's Guide

Last updated: 2026-05-28

Legal Notice

Copyright © 2026 Cohesity, Inc. All rights reserved.

© 2026 Cohesity, Inc. All Rights Reserved. Cohesity, the Cohesity Logo and other Cohesity Marks are trademarks of Cohesity, Inc. in the US and/or internationally. The information supplied herein is the confidential and proprietary information of Cohesity and may only be used (a) by the intended recipients and (b) in conjunction with validly licensed Cohesity software and services. Find the terms of Cohesity licenses at www.cohesity.com/agreements.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. COHESITY SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT 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	Overview of protecting MongoDB using NetBackup	6
	About protecting a sharded, replica set, or standalone MongoDB cluster using NetBackup	6
	Protecting MongoDB data using NetBackup	10
	NetBackup for MongoDB terminologies	12
	Limitations	13
	Prerequisites and the best practices for protecting MongoDB	14
Chapter 2	Verify the pre-requisites for the MongoDB plug-in for NetBackup	18
	Operating system and platform compatibility	18
	Prerequisites for configuring the MongoDB plug-in	18
Chapter 3	Configuring NetBackup for MongoDB	20
	About the MongoDB configuration tool	20
	Prerequisites for manually creating the mongod.conf file	22
	Configuring backup options for MongoDB using the mongod.conf file	23
	Including the configuration file path in the allowed list on the NetBackup primary server	30
	Obtaining the RSA key of the MongoDB nodes	31
	Adding MongoDB credentials in NetBackup	32
	About the credential configuration file	33
	How to add the MongoDB credentials in NetBackup	35
	About the MongoDB roles for protecting the data	36
	Host user requirements	36
	Managing backup hosts	37
	Including a NetBackup client on NetBackup primary server allowed list	37

Chapter 4	Backing up MongoDB using NetBackup	39
	About backing up MongoDB data	39
	Backing up a MongoDB cluster	41
	Prerequisites for backing up a MongoDB cluster	41
	Configuring NetBackup policies for MongoDB plug-in	43
	Creating a BigData backup policy for MongoDB clusters with web UI	43
Chapter 5	Restoring or recovering MongoDB data using NetBackup	45
	About restoring MongoDB data	45
	Prerequisites for MongoDB restore and recovery	46
	Restore the MongoDB data on the same cluster	48
	Restore the MongoDB data on an alternate cluster	49
	Restoring MongoDB data in a high availability setup to an alternate client	52
	Manual steps after the recovery process	53
Chapter 6	Troubleshooting	55
	About NetBackup for MongoDB debug logging	55
	Known limitations for MongoDB protection using NetBackup	56
Appendix A	Additional information	69
	Sample MongoDB configuration utility workflow to add and update MongoDB credentials	69

Overview of protecting MongoDB using NetBackup

This chapter includes the following topics:

- [About protecting a sharded, replica set, or standalone MongoDB cluster using NetBackup](#)
- [Protecting MongoDB data using NetBackup](#)
- [NetBackup for MongoDB terminologies](#)
- [Limitations](#)
- [Prerequisites and the best practices for protecting MongoDB](#)

About protecting a sharded, replica set, or standalone MongoDB cluster using NetBackup

NetBackup supports the protection of the following MongoDB configurations:

- Sharded MongoDB cluster
- Replica set MongoDB cluster
- Standalone MongoDB cluster without replica sets

Protecting a sharded, replica set, or standalone MongoDB cluster using NetBackup

Use the NetBackup for MongoDB plug-in to protect your sharded (MongoDB cluster with configuration server and shards), replica set, or standalone MongoDB cluster using the following high-level steps:




Table 1-1 Protecting a sharded, replica set, or standalone MongoDB cluster using NetBackup

Step overview	Details
Use NetBackup to protect MongoDB.	<p>On a very high level, to protect MongoDB you need:</p> <ul style="list-style-type: none"> ■ NetBackup primary server ■ NetBackup media server ■ A backup host (NetBackup media server or a NetBackup client) <p>Refer to the NetBackup compatibility lists for the supported primary and media server configurations. The backup host (NetBackup media server or a NetBackup client) is supported only on an RHEL or a SUSE host.</p> <p>NetBackup appliance including Flex appliance is also supported as a NetBackup primary, a media server, or as a client that can act as a backup host.</p> <p>Refer to the following topics to get a protection overview and the best practices:</p> <ul style="list-style-type: none"> ■ See “Protecting MongoDB data using NetBackup” on page 10. ■ See “NetBackup for MongoDB terminologies” on page 12. ■ See “Prerequisites and the best practices for protecting MongoDB” on page 14.
Verify the pre-requisites for the MongoDB plug-in.	<p>Refer to the following topics before you use the plug-in:</p> <ul style="list-style-type: none"> ■ See “Operating system and platform compatibility” on page 18. ■ See “Prerequisites for configuring the MongoDB plug-in” on page 18.

Table 1-1 Protecting a sharded, replica set, or standalone MongoDB cluster using NetBackup (*continued*)

Step overview	Details
Run the MongoDB configuration tool.	<p>Run MongoDB configuration tool to generate the following files automatically:</p> <ul style="list-style-type: none"> ■ The file for the MongoDB cluster topology credentials. ■ The MongoDB configuration file that configures the global NetBackup parameters for the MongoDB cluster. <p>You can access the MongoDB configuration tool using the <code>tpconfig</code> command line on the NetBackup primary server. The path to access the <code>tpconfig</code> command is <code>/usr/openv/volmgr/bin/</code> for UNIX and <code><install_path>\Volmgr\bin\</code> for Windows.</p> <p>For more information, See “About the MongoDB configuration tool” on page 20.</p>

Table 1-1 Protecting a sharded, replica set, or standalone MongoDB cluster using NetBackup (*continued*)

Step overview	Details
<p>Configure the MongoDB plug-in and the communication between NetBackup and MongoDB.</p> <p>Note: If you use the MongoDB configuration tool, some of these configuration steps are not required.</p>	<p>Create a <code>mongodb.conf</code> file to configure backup options in NetBackup:</p> <ul style="list-style-type: none"> ■ See “Prerequisites for manually creating the mongodb.conf file” on page 22. ■ See “Configuring backup options for MongoDB using the mongodb.conf file” on page 23. ■ See “Including the configuration file path in the allowed list on the NetBackup primary server” on page 30. <p>Note: If you use the MongoDB configuration tool, these steps are not required.</p> <p>Get the RSA key of the MongoDB node for adding MongoDB credentials to NetBackup:</p> <ul style="list-style-type: none"> ■ See “Obtaining the RSA key of the MongoDB nodes” on page 31. <p>Note: If you use the MongoDB configuration tool, these steps are not required.</p> <p>Add the MongoDB credentials to NetBackup to facilitate communication:</p> <ul style="list-style-type: none"> ■ See “Adding MongoDB credentials in NetBackup” on page 32. ■ See “About the credential configuration file” on page 33. ■ See “How to add the MongoDB credentials in NetBackup” on page 35. <p>Note: If you use the MongoDB configuration tool, these steps are not required.</p> <p>Give the appropriate permissions to a NetBackup user in MongoDB:</p> <ul style="list-style-type: none"> ■ See “About the MongoDB roles for protecting the data” on page 36. <p>To use a non-root user or a user without root permissions as a host user:</p> <ul style="list-style-type: none"> ■ See “Host user requirements” on page 36. <p>Identify and configure a backup host.</p> <ul style="list-style-type: none"> ■ See “Managing backup hosts” on page 37. ■ To use NetBackup client as a backup host, include the NetBackup client on the primary server on the allowed list. See “Including a NetBackup client on NetBackup primary server allowed list” on page 37.

Table 1-1 Protecting a sharded, replica set, or standalone MongoDB cluster using NetBackup (*continued*)

Step overview	Details
Back up your MongoDB database with NetBackup.	Overview of the backup process: <ul style="list-style-type: none"> ■ See “About backing up MongoDB data” on page 39. ■ See “Backing up a MongoDB cluster” on page 41. Prerequisites or the best practices for backing up MongoDB databases: <ul style="list-style-type: none"> ■ See “Prerequisites for backing up a MongoDB cluster” on page 41. ■ See “Creating a BigData backup policy for MongoDB clusters with web UI” on page 43.
Restore and recover the MongoDB database.	Overview of the restore and the recovery process: <ul style="list-style-type: none"> ■ See “About restoring MongoDB data” on page 45. Prerequisites or the best practices for backing up MongoDB databases: <ul style="list-style-type: none"> ■ See “Prerequisites for MongoDB restore and recovery” on page 46. ■ Restore and recovery using the web UI: <ul style="list-style-type: none"> ■ See “Restore the MongoDB data on the same cluster” on page 48. ■ See “Restore the MongoDB data on an alternate cluster” on page 49.

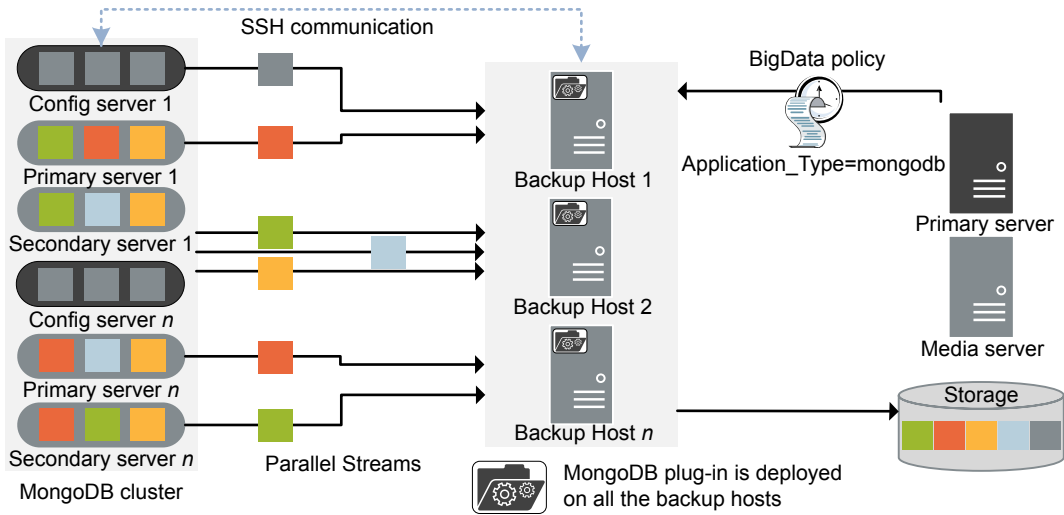
Protecting MongoDB data using NetBackup

Using the NetBackup Parallel Streaming Framework (PSF), MongoDB data can now be protected using NetBackup.

The following diagram provides an overview of how MongoDB data is protected by NetBackup.

Also, review the definitions of terminologies. See [“NetBackup for MongoDB terminologies”](#) on page 12.

Figure 1-1 Architectural overview



As illustrated in the diagram:

- The data is backed up in parallel streams wherein the data nodes stream data blocks simultaneously to multiple backup hosts. The job processing is accelerated due to multiple backup hosts and parallel streams.
- The communication between the MongoDB cluster and the NetBackup is enabled using the NetBackup plug-in for MongoDB.
- For NetBackup communication, you need to configure a BigData policy and add the related backup hosts.
- You can configure a NetBackup media server, client, or primary server as a backup host. Also, depending on the number of replica sets or sharding, you can add or remove backup hosts. You can scale up your environment easily by adding more backup hosts.
- The communication between the configuration server, secondary nodes and the backup hosts happens over SSH.
- The NetBackup Parallel Streaming Framework enables a thin client-based, agentless backup wherein the backup and restore operations run on the backup hosts. The NetBackup thin client binary is automatically pushed to the MongoDB cluster nodes during the backup and recovery operations. This thin client is automatically removed after the backup and recovery operations complete. There is no agent management required on the cluster nodes. Also, NetBackup is not affected by the MongoDB cluster upgrades or maintenance.

For more information:

- See [“About backing up MongoDB data”](#) on page 39.
- See [“About restoring MongoDB data”](#) on page 45.
- See [“Limitations”](#) on page 13.
- For information about the NetBackup Parallel Streaming Framework (PSF) refer to the *NetBackup Administrator's Guide, Volume I*.

NetBackup for MongoDB terminologies

The following table defines the terms you come across when using NetBackup for protecting MongoDB cluster.

Table 1-2 NetBackup terminologies

Terminology	Definition
Compound job	<p>A backup job for MongoDB data is a compound job.</p> <ul style="list-style-type: none"> ■ The backup job runs a discovery job for getting information of the data to be backed up. ■ Child jobs are created for each backup host that performs the actual data transfer. ■ After the backup is complete, the job cleans up the snapshots on the backup nodes, removes the thin client and is marked complete.
Discovery job	<p>When a backup job is ran, first a discovery job is created. The discovery job communicates with the config server and gathers information of the shards that need to be backed up and the associated nodes.</p> <p>At the end of the discovery, the job populates a workload discovery file that NetBackup then uses to distribute the workload amongst the backup hosts.</p>
Child job	<p>For backup, a separate child job is created for each backup host to transfer data to the storage media. A child job can transfer data blocks from multiple secondary nodes.</p>
Workload discovery file	<p>During discovery, when the backup host communicates with the config server, a workload discovery file is created. The file contains information about the data files to be backed up and the associated data nodes.</p>
Workload distribution file	<p>After the discovery is complete, NetBackup creates a workload distribution file for each backup host. These files contain information of the data that is backed up by the respective backup host.</p>
Parallel streams	<p>The NetBackup parallel streaming framework allows data blocks from multiple secondary nodes to be backed up using multiple backup hosts simultaneously.</p>

Table 1-2 NetBackup terminologies (*continued*)

Terminology	Definition
Backup host	<p>The backup host acts as a proxy client. All the backup and the restore operations are ran through the backup host.</p> <p>You can configure media servers, clients, or a primary server as a backup host.</p> <p>The backup host is also used as destination client during restores.</p>
BigData policy	<p>The BigData policy is introduced to:</p> <ul style="list-style-type: none"> ■ Specify the application type. ■ Allow backing up distributed multi-node environments. ■ Associate backup hosts. ■ Perform workload distribution.
Application server	<ul style="list-style-type: none"> ■ Sharded MongoDB cluster: Application server is the MongoDB primary config server. ■ Replica set MongoDB cluster: Application server is the primary node of MongoDB. ■ Standalone cluster: Application server is the standalone node.
Primary config server	<p>In a high-availability scenario, the primary config server is the MongoDB instance running in a primary role on a config server replica set. The primary config server must have at least one associated <code>mongos</code> service running on the same host.</p>
Failover config server	<p>In a high-availability scenario, the config server other than the primary config server that is specified as <code>alternate_config_server</code> in the <code>mongodb.conf</code> file is referred as the failover config server.</p>

Limitations

Consider the following limitations before you deploy the MongoDB plug-in:

- For highly available MongoDB cluster, if fail-over happens during a backup, then the job fails.
- IP address is not supported for the application server and backup host field. You must enter the FQDN, host name, or the short name of the application server or the backup host.
- Encrypted MongoDB environments are not supported.
- English-only MongoDB environments are supported.

- Extended Access Control Lists (ACL) are not recovered after a recovery operation.
- Recovery is not supported for a shrunk MongoDB cluster.
- For standalone MongoDB nodes without replica sets, incremental backups are not supported.
- Protection of MongoDB environments that are deployed or managed using the MongoDB Ops Manager is not supported.
- If you change the Feature Compatibility Version between the full and differential incremental backups, the backups fail.
- The backup of a sharded MongoDB environment can be taken only as a sharded backup configuration and not a replica set or any other backup configuration.

Prerequisites and the best practices for protecting MongoDB

Prerequisites

- For sharded MongoDB clusters, `mongos` and `mongod` processes must be running on the application server that is specified as the client in the backup policy.
- Only RHEL and SUSE platforms are supported for backup hosts.
- The NetBackup for MongoDB plug-in requires that the NetBackup primary server, media server, and the backup host are on NetBackup version 8.2 or later.
- Verify that NetBackup supports the MongoDB version that you have. For more information, refer to the Software Compatibility List.
- NetBackup supports the MongoDB clusters that are configured or installed on RHEL, SUSE, and Linux-s390x, IBMzSeriesRedHat operating systems.
- NetBackup supports the following MongoDB configurations:
 - Sharded MongoDB cluster (MongoDB cluster with configuration server and shards)
 - Replica set MongoDB cluster
 - Standalone MongoDB without replica sets
- NetBackup supports the following authentication types for MongoDB:
 - No authentication
 - Simple authentication

- Certificate-based authentication
- NetBackup supports the following file systems for backup and restore:
 - XFS
 - ext4
- Install OpenSSH packages on all the MongoDB nodes. Enable SSH on all the MongoDB nodes.
- NetBackup supports the MongoDB clusters that are configured with the **WiredTiger** storage engine.
- NetBackup protects the MongoDB clusters that are configured or installed locally using the `.tar` files or installed using the MongoDB official repositories.
- NetBackup supports the Differential Incremental backup for MongoDB along with a Full Backup. Cumulative Incremental backups are not supported currently.
- NetBackup recommends that you have at least three configuration servers in your sharded MongoDB environment to support high availability of the backups.
- Do not install the MongoDB plug-in on a server that also has the MongoDB application. A server that has the MongoDB application cannot be used as a backup host.
- Ensure that the local time on the MongoDB server and the backup host are synchronized with the NTP server.
- For sharded MongoDB clusters, the query router role must be present on the config server.
- For the MongoDB cluster that has a SUSE operating system, on all the MongoDB nodes set the `PasswordAuthentication` field to **Yes** in the `/etc/ssh/sshd_config` file.
 After you update the file, restart `sshd`.
 Ensure that all the clusters support the same hash key algorithm (RSA).
- If the MongoDB cluster is running under the `mongod` account, a non-root sudoer account is required, and the same needs to be configured as a host user in `tpconfig` file.
- If the MongoDB cluster is running under the non-root or root account, ensure that the host user credentials that are configured using the `tpconfig` command are of the host user account that is used to configure the MongoDB cluster (MongoDB daemon's host user account that is either root or non-root).
 For more details, See "[Host user requirements](#)" on page 36.

Best practices for communication between MongoDB and NetBackup

- If you use a NetBackup client as a backup host, ensure to add the following value in the `bp.conf` file of the NetBackup primary server:


```
APP_PROXY_SERVER=NBU_CLIENT_FQDN
```
- If the MongoDB host user does not have root permissions, ensure that the user has access to all the temporary paths to copy the thin client (`mdbserver`), logs, snapshots, etc. Add the non-root user to the `sudoers` file in the operating system.
- If you install the MongoDB using the `.tar` file or to a non-default location, add the path of the MongoDB bin folder in the `bashrc` file of the operating system to ensure that you can run the MongoDB commands from the CLI.
- If your MongoDB server uses the SUSE 12.3 operating system, ensure that you can connect to `mongod` and `mongos` process with the `--host <FQDN>` option. For more information, refer to the MongoDB Administrator's Guide.
- When you use the `-host_password` option with the `tpconfig` command and `mongodb.conf HostPassword`, ensure that the password:
 - Does not exceed 63 characters
 - Contains one or more alphanumeric characters: a-z, A-Z, 0-9
 - Contains one or more of the following characters: - (hyphen), _ (underscore), ,(comma), . (period), ? (question mark)
- When you define the paths for logs, thin clients (`mdbserver`), snapshots, or anything else in the `mongodb.conf` file, ensure that the host user in the credentials file has valid permissions to access these paths.
- To enable SSH, add the following entry in the `sudoers` file:


```
Default <host_user> !requiretty
```

Best practices for protecting MongoDB using NetBackup

- Ensure that the MongoDB limits and thresholds are as per the official MongoDB guidelines.
- Ensure that the host name is consistently used in the `tpconfig` command, during the policy configuration, and in the `mongodb.conf` file. For example, if you use the FQDN, use it for all host name instances instead of short names.
- Ensure that `application_server` matches with the host name that is used in the MongoDB environment and verified using the `db.hostInfo()` command. For example, the host name value that displayed by `db.hostInfo()`:


```
"hostname" : "<hostname_value>:<port>"
```

- Ensure that there are no JSON format errors or typos in the `mongodb.conf` file before you run a backup or restore job.
- Ensure that the path of the security certificates that are added in the `mongod.conf` file and used with the `tpconfig` command are the same for all the MongoDB nodes.
- For simple authentication, configure the same user who is part of the root group from the admin database for every MongoDB node.
- If you use the `mongod.conf` or the `mongos.conf` file to start the MongoDB processes, run the `mongod` file using the absolute system path on the MongoDB cluster. For example, use the following command:

```
mongod --config /home/user1/mongod.conf
```
- NetBackup recommends that you run a full backup after making any configuration changes in the MongoDB instance. If an incremental backup is scheduled to run after you make the configuration changes, then run a full backup manually before the incremental backup.
For examples, when you modify the MongoDB Feature Compatibility Version (FCV), MongoDB version, authentication type, topology (addition of new shards or removal of existing shards), storage parameters, etc. then run a full backup.

Verify the pre-requisites for the MongoDB plug-in for NetBackup

This chapter includes the following topics:

- [Operating system and platform compatibility](#)
- [Prerequisites for configuring the MongoDB plug-in](#)

Operating system and platform compatibility

With this release along with RHEL and SUSE, a new platform Linux-s390x, IBMzSeriesRedHat is also supported for MongoDB clusters.

Note: For backup hosts only RHEL and SUSE platforms are supported.

Also, with this release NetBackup also supports backup and restore of MongoDB 5.0

For more information, see the:

- [NetBackup Database and Application Agent Compatibility List](#)
[NetBackup Master Compatibility List](#).

Prerequisites for configuring the MongoDB plug-in

Consider the following when you configure NetBackup for MongoDB:

Prerequisites:

- Add the MongoDB thin client package that is part of `vxupdate_nb_version_SJA` to the package repository on the NetBackup primary server.

Note: The required package should correspond to the NetBackup version of the backup host and the operating system of the MongoDB host.

To add the package, run the `nbrepo` command on the NetBackup primary server:

```
./nbrepo -add vxupdate_nb_version_suse_x64.sja
./nbrepo -add vxupdate_nb_version_redhat_x64.sja
./nbrepo -add vxupdate_nbclient_version_redhat_zseries.sja
```

For a MongoDB host with CentOS operating system, add the Linux RHEL VxUpdate package of the NetBackup version of the backup host in the package repository on the NetBackup primary server.

Note: If the package is not added, the MongoDB backups can fail with error - 6729: "Unable to download the thin client from the package repository."

- Use consistent conventions for host names of backup hosts, media servers, and primary server. For example, if you are using the host name as **MongoDB.veritas.com** (FQDN format) use the same everywhere, specially while running the `tpconfig` command.
- Ensure that the backup host can communicate with all the MongoDB nodes.
- Ensure that the `bindIp` setting in the configuration file of `mongod` instance on the MongoDB hosts has value `0.0.0.0`.

Best practices:

- Add the entries of all the nodes of the MongoDB cluster to the `/etc/hosts` file on all the backup hosts. You must add the host name in FQDN format.
Or
Add the appropriate DNS entries in the `/etc/resolv.conf` file.

Configuring NetBackup for MongoDB

This chapter includes the following topics:

- [About the MongoDB configuration tool](#)
- [Prerequisites for manually creating the mongod.conf file](#)
- [Configuring backup options for MongoDB using the mongod.conf file](#)
- [Obtaining the RSA key of the MongoDB nodes](#)
- [Adding MongoDB credentials in NetBackup](#)
- [Host user requirements](#)
- [Managing backup hosts](#)

About the MongoDB configuration tool

NetBackup provides a command-line configuration tool that enables you to accurately capture and update the information that is required to protect the MongoDB.

You can use the MongoDB configuration tool to generate the following files automatically:

- The credentials file that configures the MongoDB cluster topology and credentials for NetBackup.
For more information about the credential configuration file and the manual method to create it, refer to the following topic:
See [“Adding MongoDB credentials in NetBackup”](#) on page 32.
- The MongoDB configuration file that configures the global NetBackup parameters for the MongoDB cluster.

For more information about the MongoDB configuration file and the manual method to create it, refer to the following topic:

See [“Configuring backup options for MongoDB using the mongod.conf file”](#) on page 23.

Note: You can create the two files manually, but you must ensure that the formatting and the parameters are correct.

You can access the MongoDB configuration tool using the `tpconfig` command line on the NetBackup primary server. The path to access the `tpconfig` command is `/usr/opensv/volmgr/bin/`.

- On a Windows primary server runs the `tpconfig -mongo_configuration` command to activate Mongo configuration interface.
- On a Linux and Solaris primary server, run `./tpconfig` and `tpconfig` and select the fourth option for MongoDB configuration.

See [“Sample MongoDB configuration utility workflow to add and update MongoDB credentials”](#) on page 69.

For more information about the `tpconfig` command, refer to the *NetBackup Commands Reference Guide*.

Allowed list the `mongod.conf` file path in `bp.conf` using the `bpcd_allowed_path` option. For more information, See [“Including the configuration file path in the allowed list on the NetBackup primary server”](#) on page 30.

Adding MongoDB credentials during recovery to an alternate MongoDB cluster

To recover to an alternate MongoDB cluster, use the configuration tool to add credentials of the alternate cluster in the existing cluster credentials.

Sharded MongoDB cluster

1. Use the configuration tool to update credentials of the existing cluster.
2. Add a new configuration server using the **Add new secondary config server option** and save.
3. Add shards of the new cluster using the **Add new shard host server option** and save.
4. Initiate the alternate recovery job.

Replica Set MongoDB

1. Use the configuration tool to update credentials of the existing replica set.

2. Add a new primary server using the **Add secondary server** option and save.
3. Add all of the secondary servers using the **Add secondary server** option and save.
4. Initiate the alternate recovery job.

Note: If you use the credentials file, you can manually update the file and upload the file using the `tpconfig` command.

Prerequisites for manually creating the `mongodb.conf` file

Note: If you use the MongoDB configuration tool, these manual steps are not required.

- If you do not specify any values in the `mongodb.conf` file for MongoDB cluster ports and paths to deploy the thin clients, create snapshots, or logs, the default values are considered.
- The minimum value of the `max_streams` field is 32. If `max_streams` is not defined, the default value is 32 parallel data streams per backup host.
For the `max_streams` field in the `mongodb.conf` file, the value of the backup host takes priority over the `global_default` value.
For example, the value 32 takes priority over the value 34 and the job runs 32 streams during a backup in this scenario:

```
"max_streams": {  
    "global_default": 34,  
    "Backup_Host": 32}
```

The backup streams are distributed across the backup hosts as defined in the backup policy. The streams are not distributed according to the backup hosts that are defined in the `max_streams` option in the `mongodb.conf` file.

- For a sharded MongoDB environment, ensure that the `mongodb.conf` file has the latest primary config server and secondary config server.
- Ensure that the folders or directories that are mentioned in the `mongodb.conf` file are available on the MongoDB cluster. For example, the folders or directories for `snapshot_mount_path`, `oplog_location`, `logdir`, etc.

- To the allowed list add the `mongodb.conf` file path. In the `bp.conf` file use the `bpcd_allowed_path` option.
See [“Including the configuration file path in the allowed list on the NetBackup primary server”](#) on page 30.
- Give the host user access to the port and the port range that is specified in the `mongodb.conf` file.

Configuring backup options for MongoDB using the `mongodb.conf` file

Note: If you use the MongoDB configuration tool, these manual steps are not required.

NetBackup uses the default options to back up MongoDB data. To specify custom options to use during a backup operation, you must create a `mongodb.conf` file in the `/usr/opensv/var/global/` directory on a UNIX and `<Install_Dir>\NetBackup\var\global\` on a Windows primary server.

Caution: The file name `mongodb.conf` is case-sensitive.

You do not need to specify all the options in the `mongodb.conf` file. NetBackup uses the default values for the options that do not have custom values.

Ensure that the `mongodb.conf` file uses JSON format and add the file path to the allowed list. (Use the `bpcd_allowed_path` option in the `bp.conf`.)

See [“Including the configuration file path in the allowed list on the NetBackup primary server”](#) on page 30.

Backup options in the `mongodb.conf` file

You can specify the following backup options and their values in the `mongodb.conf` file:

Caution: The options in the file are case-sensitive.

Options

`application_servers`

Details

Fully Qualified Domain Name (FQDN), or host name, or short name and the port number of the primary configuration server and `mongod` and `mongos` port in the following format:

```
clientFQDN_OR_hostname_OR_hostname:portnumber
```

Ensure that `application_server` matches with the host name value that is used in the MongoDB environment and verified using the `db.hostInfo()` command.

For example, the host name value that is displayed by `db.hostInfo()`:

```
"hostname" : "<hostname_value>:<port>"
```

Warning: Do not enter the node that acts as an Arbiter node for MongoDB.

`alternate_config_server`

Fully Qualified Domain Name (FQDN), or host name, or short name and the port number of the secondary or the alternate configuration server. You can add only one alternate configuration server for one cluster.

Use the following format for the value:

```
clientFQDN_OR_hostname_OR_hostname:portnumber
```

Ensure that `alternate_config_server` matches with the host name value that is used in the MongoDB environment and verified using the `db.hostInfo()` command.

For example, the host name value that is displayed by `db.hostInfo()`:

```
"hostname" : "<hostname_value>:<port>"
```

If a connection to the primary configuration server fails, the first, active alternate configuration server is used.

For sharded MongoDB clusters, both the `mongod` and `mongos` processes must be running on the alternate config server.

You must enter the value of the `alternate_config_server` separately for every `application_servers` entry.

`cleanup_time_in_min`

Specify the time in minutes to clean up the stale snapshots or `oplogstore` that are created because of canceled jobs.

The value must be an integer.

Options

Details

`free_space_percentage_snapshot`

Specifies the percentage of the free space on a volume group that can be used to create a snapshot. This option is used only in case of full backups.

The default value (if not specified) is 20%. The value must be between 0 and 100. Do not use the percentage symbol (%).

For example, run the `vgdisplay` command to verify the value for the "Free PE / Size" field. The `free_space_percentage_snapshot` value is the percentage of Free PE / Size of the volume group where the data path resides.

Adjust the `free_space_percentage_snapshot` value based on the data change rate of the MongoDB instance during the backup operation and the free space that is available on the volume group.

For example, when:

- Data change rate is 250 MB
- Volume group has 1 GB Free PE / Size
- The data change rate is 25% of Free PE / Size

Then specify the minimum value of `free_space_percentage_snapshot` as 25%.

Keeping the free space snapshot percentage value too low can result in snapshot (and subsequent backup) failure.

Keeping the free space snapshot percentage value too high can reduce the amount of available space on the volume group.

For more information and standard practices, refer to the Linux man page for the `lvcreate` command.

`data_channel_tls`

Use this parameter to disable or enable the data channel encryption between the MongoDB cluster and the backup host.

For example, use `"data_channel_tls": false` to disable the data channel encryption.

By default, all the traffic between the NetBackup backup host and the thin client (`mdbserver`) is over a TLS channel. You can disable this TLS for data movement from thin client (`mdbserver`) to the backup host to improve the performance.

Note: Control data and sensitive data such as credentials are still transferred over the TLS channel when this option is disabled.

Options

`logdir`

Details

Location where the thin client (`mdbserver`) logs are generated on the MongoDB nodes.

Default location is `/tmp`. If the directory path is mentioned, but the directory does not exist on the server, NetBackup creates a directory.

`loglevel`

Specify the logging level.

Default value is 3.

Refer to the following options for logging level values:

- `ESERROR = 1`
- `ESWARN = 2`
- `ESINFO = 3`
- `ESDEBUG = 4`
- `ESTRACE = 5`
- `ESCRITICAL = 6`

`max_log_mbsize`

Specify the maximum file size in MB for the NetBackup thin-client log file.

The default size is 10 MB. A new log file is created every day or if the existing log file exceeds the maximum size. The log file creation does not affect an ongoing job and the log roll-over happens during the next job that is performed by `mdbserver`.

The logs are cleaned up after 30 days.

Options

`max_streams`

Details

Note: This parameter is applicable only for sharded MongoDB clusters.

Defines the number of parallel data streams per backup host. The minimum value is 32.

If `max_streams` is not defined, the default value is 32 parallel data streams per backup host.

Add the following entry to the `mongodb.conf` file:

```
max_streams:
  {
    "global_default":<set_value>,
    "<backup_host>":<set_value>
  }
```

Where:

- `global_default`
Default upper limit of the parallel data streams for all the backup hosts.
- `backup_host`
Set the upper limit of the parallel data streams for a specified backup host.
The `backup_host` must be the same that is specified in the backup policy. If you have multiple backup hosts, the entry can be repeated for all backup hosts. If you do not specify a backup host, the `global_default` value is used.

Note: This option sets the upper limit on the number of parallel data streams per backup host. The backup or the recovery job might not use all the available streams.

`mdb_progress_loglevel`

Lets you print the progress logging information about the files that are restored to the activity monitor.

Default value is 0 (off).

To enable set "`mdb_progress_loglevel`": 1.

Note: Enabling this option can increase the recovery time.

Options

`mdbserver_location`

Details

Specify a location to copy the thin client (`mdbserver`) binaries on the MongoDB nodes that are required for the MongoDB backup and restore operation.

The files are copied to the servers that contain the data that needs to be protected and are removed once the backup operation completes.

Default location to copy the files is `/tmp`.

Note: Do not specify the mount path or the high-level Linux directories because that can cause conflicts in directory permissions. For example, avoid specifying the path as `/root`, `/etc`, `/usr`, `/bin`, `/home`, etc.

`mdbserver_port`

The port that the backup host uses to connect with the NetBackup thin client (`mdbserver`) that is running on the MongoDB node.

Default value is "11000".

This value is a string.

`mdbserver_port_range`

Use this parameter when multiple `mongod` instances are running on a single MongoDB node.

This option lets you use the next available port within the range for the backup and restore operation if the existing port is in use.

This option lets you run multiple backup jobs simultaneously on different ports by deploying multiple NetBackup thin clients (`mdbserver`).

Enter the value as `"mdbserver_port_range": range_value`, where the `range_value` is an integer to define the range of port numbers that can be used. For example, if you add the `range_value` as 10 and the `mdbserver_port` is defined 12000, the port range is used from 12000 to 12009.

The default value is 10.

Change this value based on the number of `mongod` instances on a MongoDB host that are backed up simultaneously.

`mdbserver_timeout_min`

Defines the time in minutes to wait before a NetBackup thin client (`mdbserver`) process is killed.

The default value is 300 (minutes).

Set the value higher than 300 minutes if your backup window requires more time.

Ideally, `mdbserver` is killed after the plug-in terminates or the backup is complete.

Options

`mongos_port`

Details

Port that the `mongos` process uses for communication.

This parameter is a mandatory parameter for sharded MongoDB clusters.

You must specify this value for each of the `application_servers` or `alternate_config_server` entry.

This value is a string.

`oplog_location`

For differential incremental backups, specify a custom directory to store the MongoDB `oplog` file.

The location is stored in the backup image.

Default location is `/tmp/oplogstore`.

Ensure that there is enough free space at this location for the `oplog` data of the incremental backups.

`snapshot_mount_path`

Specify the path on the MongoDB nodes for mounting LVM snapshots during full backups.

Default path is `/tmp`.

Note: Ensure that the `HostUser` that you have configured in the MongoDB credential file has the read and write permissions to all the paths mentioned in the `mongodb.conf` file.

If you do not add all the options, an entry is added in the logs about the missing options. The default values are used for the options that are not mentioned and the backup operation proceeds.

Sample of the `mongodb.conf` file contents

```
{
  "application_servers": {
    "FQDN_primary_configuration_server_1:port": {
      "alternate_config_server": [
        {
          "hostname:port": "FQDN_alternate_configuration_server_1:26051",
          "mongos_port": "26051"
        }
      ],
      "mongos_port": "26052"
    },
    "FQDN_primary_configuration_server_2:port": {
      "alternate_config_server": [
```

```

        {
            "hostname:port": "FQDN_alterate_configuration_server_2:26053",
            "mongos_port": "26053"
        }
    ],
    "mongos_port": "26054"
}
},
"mdbserver_location": "/path/to/store/mdbserver/",
"logdir": "/path/to/store/logdir/",
"mdbserver_port": "21020",
"loglevel": 5,
"max_log_mbsize": 4,
"oplog_location": "/path/to/store/oplog/",
"free_space_percentage_snapshot": "25",
"mdb_progress_loglevel": 1,
"snapshot_mount_path": "/path/to/mount/snapshot/",
"max_streams":
{
    "global_default":2,
    "FQDN_backup_host_1":1
}
}

```

Including the configuration file path in the allowed list on the NetBackup primary server

After you create the configuration file, you must include the path of the configuration file on the allowed list. This action allows the backup operation to run successfully. Perform the allowed list procedure on a NetBackup primary server.

Allowlisting is a security practice used for restricting systems from running software or applications unless these have been approved for safe execution.

To place the configuration file path on the allowed list:

- 1 Run the following command on the NetBackup primary server:

For UNIX:

```
bpsetconfig -h primaryserver_name
bpsetconfig bpcd_allowed_path = /usr/opensv/var/global/
```

Exit the command line.

- 2 For Windows:

```
bpsetconfig -h primaryserver_name
bpsetconfig bpcd_allowed_path = <install_dir>\NetBackup\var\global\
```

Exit the command line.

For more information about the `bpsetconfig` command, refer to the *NetBackup Commands Reference Guide*.

For more information about `bpcd_allowed_path`, see the *Configuration options for NetBackup servers* section in the *NetBackup Administrator's Guide, Volume I*.

Obtaining the RSA key of the MongoDB nodes

Use the following command on the MongoDB hosts to get the SHA256 based RSA of every MongoDB node from the MongoDB cluster:

```
cat /etc/ssh/ssh_host_rsa_key.pub |awk '{print $2}' |base64 -d |sha256sum |awk '{print $1}'
```

The output of the commands is the RSA key.

For example:

```
cat /etc/ssh/ssh_host_rsa_key.pub |awk '{print $2}' |base64 -d |sha256sum |awk '{print $1}'
```

Command output:

```
b2352722053ac9f40bc1XXXXXXXXXXXXXXXXXXXXXXXXXXXX419fa241ba9431fd6b9
```

Copy the RSA fingerprint. You need to provide this fingerprint when you add the MongoDB credentials.

Adding MongoDB credentials in NetBackup

Note: If you use the MongoDB configuration tool, these manual steps are not required.

To establish a seamless communication between MongoDB clusters and NetBackup for successful backup and restore operations, you must add and update MongoDB credentials to the NetBackup primary server.

About the authentication types for MongoDB that NetBackup supports

NetBackup supports the following authentication types for protecting the MongoDB data:

- No authentication
- Simple - Salted Challenge Response Authentication Mechanism (SCRAM)
- Certificate-based - x.509

Different options are required for each of the authentication types when you add the credentials using the `tpconfig` command.

The following table describes the options that are required for each authentication type:

Table 3-1 Required options for authentication types

Options	Option description	No authentication	Simple authentication	Certificate-based authentication
AppUserId	Specifies the username that is required to log into the application server.	✗	✓	✗
AppUserPassword	Specifies the user password that is required to log into the application server.	✗	✓	✗
HostUser	Specify the host's user ID for SSH implementation. See "Host user requirements" on page 36.	✓	✓	✓
HostPassword	Specify the host's user password for SSH implementation.	✓	✓	✓

Table 3-1 Required options for authentication types (*continued*)

Options	Option description	No authentication	Simple authentication	Certificate-based authentication
HostRSAKey	RSA key is required to perform password-less remote operations.	✓	✓	✓
ServerPemPath	Path to the PEM certificate file on the MongoDB node.	✗	✗	✓
CAPemPath	Path to the CA PEM certificate file on the MongoDB node.	✗	✗	✓
Passkey	Password of CA certificate.	✗	✗	✓
CADir	Path to the CA certificate.	✗	✗	✓
CARole	User role that is defined in the CA.	✗	✗	✓
CertificateUser	Specifies the details for the certificate user.	✗	✗	✓
application_server_conf	Specifies a path to the credential configuration file that contains the authentication type, user details, and the directory paths for the CA security certificates. See “About the credential configuration file” on page 33.	✓	✓	✓

About the credential configuration file

The credential configuration file is required for authentication. You can create this file at any location and use the path when you add the MongoDB node credentials.

Add the credentials of the MongoDB node using the `tpconfig` command for the application server that is specified in the **Clients** tab in the backup policy.

Multiple MongoDB nodes must be separated using a comma.

- For sharded MongoDB cluster, add all the mongod and mongos ports in the following format:

```
mongod_hostname:mongod_port
mongos_hostname:mongos_port
```

- For replica set MongoDB cluster and standalone MongoDB cluster, add the mongod port in the following format:

```
mongod_hostname:mongod_port
```

Refer to the following sample credential file that contains all the authentication types:

- No Authentication

```
{
"app.server.com:26050" : {
  "AuthType": "NOAUTH",
  "HostUser": "root",
  "HostPassword": "password",
  "HostRsaKey": "b2352722053ac9f40bc1XXXXXXXXXXXXXXXXXXXXXXXXX419fa241ba9431fd6b9"
}
}
```

- Simple Authentication

```
{
"app3.server.com:26051" : {
  "AuthType": "SIMPLEAUTH",
  "HostUser": "root",
  "HostPassword": "password",
  "HostRsaKey": "b2352722053ac9f40bc1XXXXXXXXXXXXXXXXXXXXXXXXX419fa241ba9431fd6b9",
  "AppUserId": "admin",
  "AppUserPassword": "password"
}
}
```

- Certificate-based authentication

```
{
"app4.server.com:26052" : {
  "AuthType": "CERTAUTH",
  "HostUser": "root",
  "HostPassword": "password",
  "HostRsaKey": "b2352722053ac9f40bc1XXXXXXXXXXXXXXXXXXXXXXXXX419fa241ba9431fd6b9",
  "ServerPemPath": "/root/certs/user1.pem",
  "CAPemPath": "/root/certs/mongo-CA-cert.crt",
  "Passkey": "password",
  "CA Dir": "/root/",
  "CARole": "root",
  "CertificateUser": "CN=user1,OU=nbu,O=vtas,L=pune,ST=mh,C=in"
}
}
```

How to add the MongoDB credentials in NetBackup

Use the `tpconfig` command to add credentials in NetBackup primary server.

For more information about the `tpconfig` command, see the *NetBackup Commands Reference Guide*.

Before you run the `tpconfig` command, ensure to remove all the earlier entries for the MongoDB nodes.

To run the `tpconfig` command:

- 1 Run `tpconfig` command from the following directory paths:

On UNIX systems, `<install_directory>/volmgr/bin/`

On Windows systems, `<install_path>\volmgr\bin\`

- 2 Run the `tpconfig --help` command. A list of options which are required to add, update, and delete MongoDB credentials is displayed.

To add credentials for all of the authentication types:

- 1 Run the following command by providing appropriate values for each options to add MongoDB credentials.

```
tpconfig -add -application_server  
app_server_name-mongod_port_number -application_type mongodb  
-requiredport mongod_port_number -application_server_conf  
/<install_directory>/var/global/mongodb_cred_file.conf
```

Where:

`application_server` is `mongodb_hostname-mongodport`

`application_server_conf` is a credential file to add support for single or multiple `mongod` per MongoDB cluster

You can use the `-update` or `-delete` options of the `tpconfig` command to update or delete the MongoDB credentials.

For more information, See [“About the credential configuration file”](#) on page 33.

- 2 Run the `tpconfig -dappservers` command to verify if the NetBackup primary server has the MongoDB credentials added.

Note: The encrypted credential configuration file (name:`appservername-portnumber.conf`) is created on the NetBackup primary server at the following location: `/usr/opensv/var/global/`.

About the MongoDB roles for protecting the data

For completing NetBackup operations, the application user that you add in NetBackup must have the required roles. The required role must have permissions to access, query, back up, and restore all the databases and manage the cluster. It is recommended that you assign the user a root role.

Note: Ensure that the same user is added for all the MongoDB nodes.

For more information about roles, refer to the MongoDB Manual.

Host user requirements

Host User requirements for clusters running under mongod account:

- Don't specify mongod service account as the host user in `tpconfig`.
- On the server where MongoDB is installed, add the host user in the `sudoers` file.
- Ensure to add a sudoer non-root account as the host user in `tpconfig`.
- Ensure that the host user has a home directory.
- Give that host user the ownership rights to all the directories that you have mentioned in the `mongodb.conf` file.
- This step ensures that the backup host can access the directories to copy the required files for backup operations. For example, the user needs rights to `mdbserver_location`, `logdir`, `oplog_location`.

If the host user that you want to use is a non-root user or does not have root permissions for the MongoDB server, then you must complete the following steps:

- Ensure that the host user credentials that are configured using the `tpconfig` command are of the host user account that is used to configure the MongoDB cluster (MongoDB daemon's host user account).
- Give that host user the ownership rights to all the directories that you have mentioned in the `mongodb.conf` file. This step ensures that the backup host can access the directories to copy the required files for backup operations. For example, the user needs rights to `mdbserver_location`, `logdir`, `oplog_location`.
- Ensure that the host user has ownership of the MongoDB database path and its contents. This ownership is required for the backup as well as restore operations.

- On the server where MongoDB is installed, add the host user in the `sudoers` file.

Managing backup hosts

A backup host acts as a proxy client which hosts all the backup and restore operations for MongoDB clusters. In case of MongoDB plug-in for NetBackup, backup host performs all the backup and restore operations without any separate agent installed on the MongoDB cluster.

The backup host must have a Linux operating system. NetBackup supports only RHEL and SUSE platforms as a backup host.

The backup host can be a NetBackup client or a media server or a primary server. NetBackup recommends that you have media server as a backup host.

Consider the following before adding a backup host:

- For backup and restore operations, you can add one or more backup hosts.
- A primary, media, or client can perform the role of a backup host.
- MongoDB plug-in for NetBackup is installed on all the backup hosts.
- When using multiple backup host, make sure that all backup hosts are communicating with the media server.

You can add a backup host while configuring **BigData** policy using the NetBackup web UI.

For more information on how to create a policy, see See [“Creating a BigData backup policy for MongoDB clusters with web UI”](#) on page 43.

Including a NetBackup client on NetBackup primary server allowed list

To use the NetBackup client as a backup host, you must include it in the allowed list. Perform the allow list procedure on the NetBackup primary server .

Allowlisting is a security practice used for restricting systems from running software or applications unless these have been approved for safe execution.

To place a NetBackup client on NetBackup primary server on the allowed list

- ◆ Run the following command on the NetBackup primary server:

- For UNIX

```
bpsetconfig -h primaryserver
```

```
bpsetconfig> APP_PROXY_SERVER = clientname1.domain.org
```

```
bpsetconfig> APP_PROXY_SERVER = clientname2.domain.org
bpsetconfig>
UNIX systems: <ctl+D>
```

- For Windows

```
bpsetconfig -h primaryserver
bpsetconfig> APP_PROXY_SERVER = clientname1.domain.org
bpsetconfig> APP_PROXY_SERVER = clientname2.domain.org
bpsetconfig>
Windows systems: <ctl+Z>
```

For more information about the `bpsetconfig` command, refer to the *NetBackup Commands Reference Guide*.

This command sets the `APP_PROXY_SERVER = clientname` entry in the backup configuration (`bp.conf`) file.

For more information about the `APP_PROXY_SERVER = clientname`, refer to the *Configuration options for NetBackup clients* section in *NetBackup Administrator's Guide, Volume I*

Backing up MongoDB using NetBackup

This chapter includes the following topics:

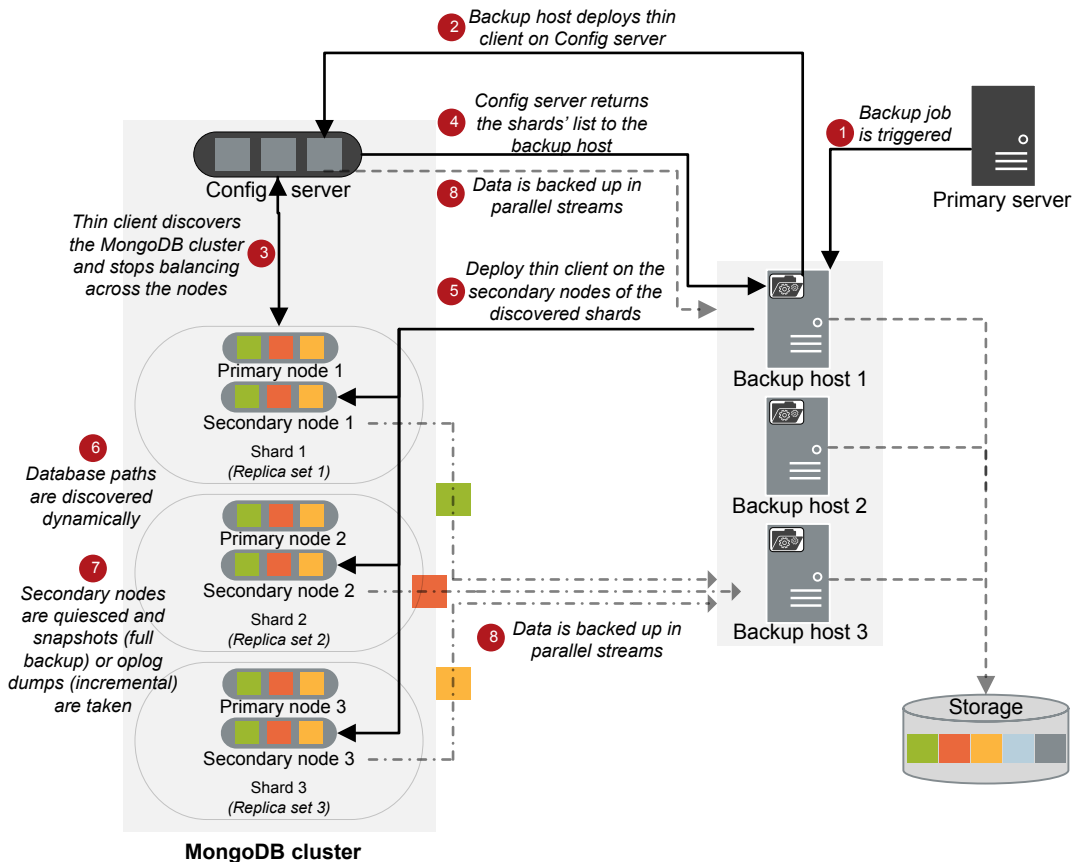
- [About backing up MongoDB data](#)
- [Prerequisites for backing up a MongoDB cluster](#)
- [Configuring NetBackup policies for MongoDB plug-in](#)

About backing up MongoDB data

MongoDB data is backed up in parallel streams wherein MongoDB data nodes stream data blocks simultaneously to multiple backup hosts.

The following diagram provides an overview of the backup flow:

Figure 4-1 Backup flow



As illustrated in the above diagram:

1. A scheduled backup job is triggered from the primary server.
2. Backup job for MongoDB data is a compound job. When the backup job is triggered, first a discovery job runs.
3. During discovery, the backup host deploys a transient thin client (`mdbserver`) on the configuration server and obtains the details of the shards in the MongoDB cluster. The thin client also stops the balancing across the nodes in a replica set.
4. After receiving the information about the cluster, the backup host deploys a thin client on the secondary node of a replica set in the MongoDB cluster.

5. The thin client discovers the database paths dynamically, quiesces the secondary nodes, and takes snapshots for full backups and captures `oplog` for incremental backups.
6. Individual child jobs run for each backup stream and data is backed up.
7. Data blocks are streamed simultaneously from different secondary nodes to multiple backup hosts.

Once the backup operation is completed, the thin client is removed from the servers.

The compound backup job is not completed until all the child jobs are completed. After the child jobs are completed, NetBackup cleans all the snapshots from the secondary nodes. Only after the cleanup activity is completed, the compound backup job is completed.

See [“Backing up a MongoDB cluster”](#) on page 41.

Backing up a MongoDB cluster

You can either schedule a backup job or run a backup job manually. See the [NetBackup Administrator's Guide, Volume I](#).

For overview of the backup process, See [“About backing up MongoDB data”](#) on page 39.

Prerequisites for backing up a MongoDB cluster

- NetBackup chooses the node in a MongoDB cluster to take a backup in the following order:
 - Active hidden node
 - Active secondary node
 - Active primary nodeIf you want NetBackup to select a particular backup node in the MongoDB cluster, set it as a hidden node.
- Before you run a backup job, ensure that you get a successful ping response from all the MongoDB nodes on the backup host. Check and update the firewall settings so that the backup hosts can communicate with the MongoDB cluster.
- Ensure that the MongoDB cluster that you want to protect lets you take LVM snapshots.
- Logical volume requirement for snapshot:
 - Ensure that the MongoDB database directory is mounted on a logical volume to complete the snapshot operations.

- Use the `vgdisplay` command to ensure that the free physical extent size is sufficient in the logical volume group to complete the snapshot operations.
- Renaming the volume group or the physical and logical volumes of the LVM for MongoDB database paths causes the backup to fail. If you rename the volume group or the physical and logical volumes of the LVM, ensure that the MongoDB database is mounted on the new path before taking a backup.
- The backup shuts down the balancer on the `mongos` process and blocks all the other operations. Hence, during a backup process, ensure that you do not run any other operation that uses the `mongos` process. For example, import the database.
- Always run a full backup when you make change to the database path, or modify the configuration file for the `mongod` or the `mongos` processes, or change the MongoDB topology.
- If there are multiple MongoDB clients in a single NetBackup backup policy increase the **Client read timeout** parameter for primary server, media server, and the client to ensure that the all the backups are successful.
For more information, refer to the NetBackup™ Administrator's Guide, Volume I and the [Timeouts properties section](#).
- Incremental backup jobs use consistent backup images as a reference for determining the incremental changes. If the previous backup has failed, or was partially successful (failed for one of the nodes), it is skipped completely and a previous backup image is considered. In such cases, the backup operation can take longer and the image that is created might be of a larger size.
- The `oplog` file has a capped or rolling cache and you can configure the size of the file. NetBackup uses `oplog` to capture incremental data. `oplog` roll-over can cause the incremental backups to fail. To prevent this, make sure that `oplog` file size is sufficient to hold the incremental data that is generated between the incremental backups.
- If the MongoDB cluster is running under the `mongod` account, we need a non-root sudoer account, and the same needs to be configured as a host user in `tpconfig`.
- If the MongoDB cluster is running under the non-root or root account, then ensure that the host user credentials that are configured using the `tpconfig` command are of the host user account that is used to configure the MongoDB cluster (MongoDB daemon's host user account that is either root or non-root).
- If you are using *oplog retention feature* on replica set, ensure that the scheduled time between incremental backups is less than the minimum oplog retention period. This ensures capturing the correct incremental backup.

- NetBackup supports full backup only for MongoDB 4.4 and later supported versions on sharded cluster.

Configuring NetBackup policies for MongoDB plug-in

Backup policies provide the instructions that NetBackup follows to back up clients. For configuring backup policies for MongoDB plug-in, use the **BigData** policy as the **Policy Type**.

You can create **BigData** policy using the **NetBackup web UI**.

Creating a BigData backup policy for MongoDB clusters with web UI

- 1 On the left, select **Protection > Policies**.
- 2 Click **Add**.
- 3 On the **Attributes** tab, select the following:
 - **Policy type:** BigData
- 4 On the **Schedules** tab, configure all the necessary schedules. For example, Full and incremental schedules.

Name	Type	Frequency	Retention
Full-backup	Full backup	1 week	2 weeks
Incremental-backup	Differential incremental backup	1 day	2 weeks

- 5 On the **Clients** tab, based on your MongoDB setup enter the following value:

Sharded MongoDB cluster:

 - The client name as seen in the MongoDB shell and the `mongod` port number of the primary configuration server in the following format:
`MongoDBNode-portnumber`

Replica set MongoDB cluster:

 - The client name as seen in the MongoDB shell and the `mongod` port number of the primary node for replica set in the following format:
`MongoDBNode-portnumber`

Standalone MongoDB setup:

- The client name as seen in the MongoDB shell and the `mongod` port number of the standalone node in the following format:

```
MongoDBNode-portnumber
```

Warning: Warning: Do not enter the node that acts an Arbiter node for MongoDB.

- 6 On the **Backup selections** tab, add the application type, the backup hosts, and manually add the ALL_DATABASES directive.

Backup selection list	Notes
<code>Application_Type=mongodb</code>	The parameter values are case-sensitive.
<code>mongodbhost=mongodbhost.domain.com</code>	Use the format <code>Backup_Host=<FQDN_or_hostname></code> . The backup host can be a NetBackup client or a media server.
<code>ALL_DATABASES</code>	

- 7 Click **Create**.

Restoring or recovering MongoDB data using NetBackup


This chapter includes the following topics:

- [About restoring MongoDB data](#)
- [Prerequisites for MongoDB restore and recovery](#)
- [Restore the MongoDB data on the same cluster](#)
- [Restore the MongoDB data on an alternate cluster](#)
- [Restoring MongoDB data in a high availability setup to an alternate client](#)
- [Manual steps after the recovery process](#)

About restoring MongoDB data

For restore you require a backup host. Backup hosts can be a primary server, a media server, or a NetBackup client.

The following diagram provides an overview of the restore flow.

Figure 5-1 Restore flow

The diagram describes the following process:

1. The restore job is triggered from the primary server.
2. The backup host connects with the `config` server. Backup host is also the destination client.
3. The actual data restore from the storage media starts.
4. The data blocks are restored on the MongoDB cluster.

Prerequisites for MongoDB restore and recovery

Review the following prerequisites and limitations before you begin the restore or the recovery process:

- Ensure that your source MongoDB cluster and the target MongoDB cluster have the same:
 - MongoDB version
 - Authentication type
- If the target MongoDB cluster runs under the `mongod` account, you need a non-root `sudoer` account. Also configure it as a host user in `tpconfig` before you start the restore. NetBackup recovers and runs the instance under the target cluster's user.
- (Restoring MongoDB data to an alternate cluster) Ensure that you add the alternate MongoDB cluster credentials in the source cluster credentials file. See [“Adding MongoDB credentials in NetBackup”](#) on page 32.

- Ensure that the PEM files or security certificates are available on the destination cluster before you start the recovery operation.
- The authentication type of the target cluster during the recovery process must be the same as the authentication type during the backup.
- Ensure that the target MongoDB cluster has sufficient free storage space for restoring the data.
- Select only one full backup image group and its appropriate incremental images. If you select more than one full backup image group, the recovery may fail as the restored data can get corrupted.
- NetBackup for MongoDB plug-in does not support cross-platform file system restore. For example, XFS to ext4 or vice versa is not supported.
- Ensure that the `HostUser` value that is defined in the `tpconfig` command is the same as the host user account that is used to configure the MongoDB cluster. (MongoDB daemon's host user account.)
- For the destination client, select a backup host before you submit the restore job.
- Point-in-time recovery is valid only for recovery from incremental backups.
- Canceling a parent job in a compound restore job does not cancel the child restore jobs. You must manually cancel the child restore jobs as well.
- After you start a restore or a recovery job in the web UI, look for the job record and status in the **Task progress** section in the **Recover** node. The job might take some time to appear in the list and the compound job can take time to start the parent pre-recovery check. Click **Refresh** to refresh the task list.
- In a **Restore Only** operation for a sharded cluster, follow the standard **Restore Only** steps:
 - Shut down all the MongoDB processes (`mongos` or `mongod`) before starting the restore.
 - The `tpconfig` host user must have permission to the target folder in the original restore and alternate restore.
- The MongoDB log file paths remain the same from the original configuration. If you do an alternate restore:
 - Make sure that the same path is available during restore.
 - Change the log file path in the configuration file for `mongod` or `mongos` process after a successful recovery.
- The path to the `.pid` files must be available on the destination MongoDB cluster to ensure that the recovery operation is successful.

- Special consideration is needed when you back up multiple MongoDB clusters that run on the same server. (Backups that use the same or different backup policies.) Ensure that you select the correct application server that you want to restore.

For example, if there are multiple clusters with the following configuration:

```
Replica1
Primary:   host1:26050
Secondary: host1:26060
```

```
Replica2
Primary:   host1:26055
Secondary: host1:26066
```

And you want to recover `Replica1`, ensure that you specify the correct application server and its port (`host1-26050`) as the **Source client**.

- Before you start a recovery operation, ensure that stale `mdbserver` (thin client) processes are not running on your MongoDB instance at the following path:

```
/<mdbserver_location>/<Host>-<MongodPort>-<mdbserver_port in range>/mdbserver
```

If any of these stale processes are running, the recovery operation is unable to shut down the MongoDB instance which causes the recovery job to stop responding.

- Restore and recovery of the MongoDB cluster requires the same security mode that is used at the time of the backup. Ensure that the security mode is the same for the original cluster and the target cluster.

For example, if SSL is used during the backup, then recovery is done using SSL and the target configuration is changed to SSL. Similarly, if TLS is used during the backup, then recovery is done using TLS and the target configuration is changed to TLS.

- Restore and recovery of the MongoDB cluster requires the same value of the Feature Compatibility Version (FCV) that is used at the time of the backup. Ensure that FCV is the same for the original cluster and the target cluster.

For example, if FCV is 4.2 during the backup, then restore uses FCV 4.2 and the target cluster has FCV 4.2 after the recovery process completes. Similarly, if FCV is 4.0 during the backup, then restore uses FCV 4.0 and the target cluster has FCV 4.0 after the recovery process completes.

Restore the MongoDB data on the same cluster

This topic describes how to restore MongoDB data on the same cluster.

To restore the MongoDB data on the same cluster

- 1 Open the NetBackup web UI.
- 2 On the left, click **Recovery**.
- 3 On the **Regular recovery** card, click **Start recovery**.
- 4 In **Basic properties**, enter the following:
 - Select the **Policy type** as **BigData > MongoDB**
 - From the **Source client** list, select the required Application server.
 - From the **Destination client** list, select the required backup host. The restore is faster if the backup host is the media server that backed up the node.
 - Click **Next**.
- 5 In **Recovery details**, do the following:
 - If necessary, click **Edit** to select the appropriate date range to restore the complete data set. Or select **Use backup history** and select the backup images that you want to restore.
 - Granular restore and recovery are not supported for MongoDB. Therefore, make sure to select all files and folders for the restore process.
 - Click **Next**.
- 6 In **Recovery options**, do the following:
 - Select **Restore everything to original location** to restore your files to the same location where you performed your backup.
 - In **MongoDB options**, select **Restore and recover** and recover databases for the current time or select the specific schedule.
 - Click **Next**.
- 7 Review the recovery details and click **Start recovery**.

Restore the MongoDB data on an alternate cluster

NetBackup supports the following alternate recovery scenarios for MongoDB:

- Redirected restore and recovery to an alternate cluster
- Redirected restore and recovery to an alternate node or port or database path in an existing cluster

To restore the MongoDB data on an alternate cluster

- 1 Run the `tpconfig` command to update the original cluster's credentials with the alternate application server's credentials.

For example, to recover source client `Host1-26050` to an alternate application server `Host2` that is running on port `28001`:

- Add the credentials of `Host2:28001` and its related nodes in the original cluster's credential configuration file.
See [“About the credential configuration file”](#) on page 33.
- Run the update `tpconfig` command for `application_server` that you want to recover (`Host1-26050`)

Here is a sample command:

```
/usr/opensv/volmgr/bin/tpconfig -update -application_server  
Host1-26050 -application_type mongodb -requiredport 26050  
-application_server_conf /usr/opensv/var/global/credential.conf
```

- 2 Open the NetBackup web UI.
- 3 On the left, click **Recovery**.
- 4 On the **Regular recovery** card, click **Start recovery**.
- 5 In **Basic properties**, enter the following:
 - Select the **Policy type** as **BigData > MongoDB**
 - From the **Source client** list, select the required Application server.
 - From the **Destination client** list, select the required backup host. Restore is faster if the backup host is the media server that had backed up the node.
 - Click **Next**.
- 6 In **Recovery details**, do the following:
 - If necessary, click **Edit** to select the appropriate date range to restore the complete data set. Or select **Use backup history** and select the backup images that you want to restore.
 - Granular restore and recovery are not supported for MongoDB. Therefore, make sure to select all files and folders for the restore process.
 - Click **Next**.
- 7 In **Recovery options**, do the following:
 - Rename the application server and its nodes and set the value for the alternate application server.

Select **Restore individual directories and files to different locations**.

To change the folder paths, click **Edit file paths**.

See [the section called “Alternate restore from a nested database path”](#) on page 51.

- In the **MongoDB options**, select **Restore and recover** and recover databases for current time or select the specific schedule.
- Click **Next**.

8 Review and click **Start recovery**.

You can check the status in the **Activity monitor**.

Restoring the MongoDB oplog file to an alternate temporary location

You can restore the MongoDB `oplog` files from an incremental backup to an alternate path. The files and their path are seen in the web UI.

You must specify the paths during the alternate restore using the **Restore individual directories and files to different locations** option.

If you want to retain the original MongoDB path but change the `oplog` file path, click **Edit file paths** and specify the source and alternate paths.

For example, Source `/host:port/tmp` and Destination `/host:port/alternate_tmp`.

Alternate restore from a nested database path

For an alternate restore from a nested database path, use the Add Destination dialog box and for every subfolder, add an appropriate target alternate path.

For example, to change the path from `/host:port/usr/mongodb/db1` to `/host:port/alt-dir/dbpath/mydb`:

- Specify the source and the destination path:
Source `/host:port/usr/mongodb/db1` and Destination `/host:port/alt-dir/dbpath/mydb`
- Specify the source and the destination path for the parent folder:
Source `/host:port/usr/mongodb` and Destination `/host:port/alt-dir/dbpath`
- Specify the source and the destination path for the base parent folder:
Source `/host:port/usr` and Destination `/host:port/alt-dir`

Note: When you do an alternate restore to a non-root path, the restore is partially successful if the database path contains multiple subfolders.

In such a scenario, when you do an alternate restore to a different location, you must add an entry for each directory level.

For example:

Source: /hostname1:port1/Config_Data

Destination: /hostname2:port3/mongo_inst2

Source: /hostname1:port1/Config_Data/data

Destination: /hostname2:port3/mongo_inst2/data

Source: /hostname2:port2/Shard1_Primary

Destination: /hostname2:port3/mongo_inst2

Source: /hostname2:port2/Shard1_Primary/data

Destination: /hostname2:port3/mongo_inst2/data

Restoring a MongoDB cluster where the backups are taken from different MongoDB nodes in the same replica set

You can restore a MongoDB cluster (sharded or replica set) that was backed up from different nodes. This capability exists because of the role switch (between primary and secondary nodes) within a shard or a replica set. In this scenario, the full backup can be taken from one host and the incremental backup is taken from another host in the same shard or replica set.

During restore, you must redirect the restore of these backup images to the same MongoDB host.

For example, to restore backups from /host1:port1/dbpath and /host2:port1/tmp, specify the following:

Source /host1:port1/dbpath and Destination /alhost:port1/dbpath

Source /host2:port1/tmp and Destination /alhost:port1/tmp

Restoring MongoDB data in a high availability setup to an alternate client

If a client (MongoDBnode-port) that is defined in the backup policy is not available use the following procedure to restore to a different client in the same MongoDB cluster (MongoDBnode-port).

In a high availability setup, you can restore the MongoDB data as follows:

- Sharded MongoDB cluster
Restore to an alternate config server in the same MongoDB cluster.
- Replica set MongoDB cluster
Restore to an alternate node of the replica set that is in the same MongoDB cluster.

To restore MongoDB data in a high availability setup to an alternate client

- 1 Open the web UI.
- 2 On the left, click **Recovery**. On the **Regular recovery** card, click **Start recovery**.
- 3 Rename the application server and its nodes and set the value for the alternate application server.
 - Select **Restore individual directories and files to different locations** and click **Edit file paths** to add the alternate application servers.
 - If the `application_server (Host1-port1)` is different from the target `application_server (Host2:Port2)`, the rename entry must contain `ALT_APPLICATION_SERVER=Host2:Port2`.
In the following example, the **Source client** that is defined in the backup policy is `endu79-26050`, and the backup was performed by the MongoDB node `endu79-26055`. In this scenario, as part of restore and recovery, append `endu79:26055` as follows:`ALT_APPLICATION_SERVER=endu79:26055`.

Manual steps after the recovery process

- After you recover the backup images that were taken from the hidden MongoDB nodes, the hidden nodes become primary nodes. Update all such primary nodes in the shard list and restart the `mongos` process using the following command:

```
db.getSiblingDB('config').shards.updateOne({ "_id" : "shard1" }, {
  $set : { "host" :
    "ShardName/rep11.example.net:27018,rep12.example.net:27018,rep13.example.net:27018"
  } })
```
- After the recovery process is complete, manually add the secondary nodes to the cluster. Before adding the nodes, ensure that ownership and permissions on the MongoDB data path on secondary nodes are set properly. For more information, refer to the following article:
[add-members-to-the-replica-set](#)

- After the recovery operation, the `mongod` or `mongos` process is started using the configuration files from the `/tmp` location. Ensure that you move the configuration files to a selected location and restart the services from that location. Remove the configuration files from the `/tmp` location so that the restore or recovery operations can restore files at the `/tmp` location with the same name for different users. If you do not remove the files, the subsequent recovery operation using a different user fails with error 2850 because the configuration files cannot be restored at the `/tmp` location. You can add more MongoDB configuration parameters if there are any changes from the backup data that is restored.
- Before recovery, if you start the MongoDB services with `systemctl` command, then `systemctl status mongod` command may show the `mongod` status as dead after recovery. This happens because after recovery, the `mongod` services are brought up with the `config` file under `/tmp` location. In such scenarios, bring up the service again with a use of `systemctl start mongod` commands.

Troubleshooting

This chapter includes the following topics:

- [About NetBackup for MongoDB debug logging](#)
- [Known limitations for MongoDB protection using NetBackup](#)

About NetBackup for MongoDB debug logging

NetBackup maintains process-specific logs for the various processes that are involved in the backup and restore operations. Examining these logs can help you to find the root cause of an issue.

These log folders must already exist in order for logging to occur. If these folders do not exist, you must create them.

Additionally, after every job (backup/restore), `mdbserver` log(s) that are created on the MongoDB nodes are copied to the respective backup hosts from where it was processed. These logs are kept in the `nbaapireq_handler` folder so that they can be collected by `nbsu` or `nbcplogs` utilities. To retain uniqueness of log file names collected from different hosts to a single folder, each log file name is prefixed with the hostname. For example, if the log files generated by `mdbserver` on hosts `MDBSERVER1` and `MDBSERVER2` for a backup job are "root.mdbserver.121219_00001.log", they are copied back to the backup host as `MDBSERVER1-root.mdbserver.121219_00001.log` and `MDBSERVER2-root.mdbserver.121219_00001.log`.

The log folders reside on the following directories

- On Windows: `install_path\NetBackup\logs`
- On UNIX or Linux: `/usr/opensv/netbackup/logs`

Table 6-1 NetBackup logs related to MongoDB

Log Folder	Messages related to	Logs reside on
<code>install_path/NetBackup/logs/nbaapidiscv</code>	BigData framework, discovery, and MongoDB configuration file logs	Backup host
<code>install_path/NetBackup/logs/bpbrm</code>	Policy validation, backup, and restore operations	Media server
<code>install_path/NetBackup/logs/bpbkar</code>	Backup	Backup host
<code>install_path/NetBackup/logs/tar</code>	Restore and MongoDB configuration file	Backup host
<code>install_path/NetBackup/logs/nbaapireq_handler</code>	<code>nbaapireq_handler</code> and <code>mdbserver</code>	Backup host

For more details, refer to the [NetBackup Logging Reference Guide](#).

Known limitations for MongoDB protection using NetBackup

The following table lists the known limitations for MongoDB protection using NetBackup:

Table 6-2 Known limitations

Limitation	Workaround
<p>Consider a configuration with a sharded MongoDB cluster with high availability that contains multiple <code>mongos</code> processes. Before you start a restore and recover operation, only the <code>mongos</code> process on the restore destination for the Config Server Replica Set (CSRS) image should be running.</p> <p>Manually stop any other <code>mongos</code> processes in the cluster before you start a restore and recover operation.</p> <p>After recovery reconfigure the <code>mongos</code> services to point to the recovered cluster.</p> <p>If the <code>mongos</code> process is not shut down on all nodes except one, the additional <code>mongos</code> processes might conflict with the restore operation. This situation causes the data that is restored to be inaccessible with a connection to <code>mongos</code>.</p>	<p>If you do not shut down the <code>mongos</code> processes before you start the restore and recovery, then after recovery you must manually shut down the stale <code>mongos</code> processes. Then restart all the recovered <code>mongod</code> and <code>mongos</code> processes under the cluster.</p>

Table 6-2 Known limitations (*continued*)

Limitation	Workaround
<p>You must start the MongoDB processes with an absolute path to the configuration files. You must use the absolute paths for the certificate files and the CA file as well. You must specify the absolute paths for the CA file, PEM file, and key files as well.</p>	<p>N/A</p>
<p>If the authentication type that was present during backup changes and you run a recovery job that requires a different authentication, the recovery process might fail.</p>	<p>Ensure that the authentication type during recovery remains the same as the type that was used during the backup.</p>
<p>If after you run a backup you then rename the volume group or the logical volume, the subsequent backup may fail.</p>	<p>N/A</p>
<p>During recovery, ensure that you select only one full backup image and the subsequent incremental images that are relevant. If you select more than one image, the recovery may fail as the restored data could be corrupted.</p>	<p>N/A</p>
<p>After your recover the MongoDB cluster, the cluster information for only the restored node is available.</p>	<p>After the recovery process is complete, manually add the secondary nodes to the cluster.</p> <p>For more information, refer to the following article: add-members-to-the-replica-set</p>
<p>During the backup process, if the MongoDB import operation is running, it can become unresponsive. Avoid the MongoDB import operation during the backup or restore process.</p>	<p>N/A</p>
<p>During the restore process, the message The restore was successfully initiated is displayed, but the restore job does not start.</p>	<p>This issue occurs when you enter the Application server for both the Source client and the Destination client in the web UI.</p> <p>Ensure that Source client and Destination client are entered correctly. The Source client must be the Application server and the Destination client must be the backup host.</p>
<p>If your environment has DNAT, ensure that the backup host or the restore host and all the MongoDB nodes are in the same private network.</p>	<p>N/A</p>
<p>The NetBackup for MongoDB plug-in does not support the command line <code>bprestore options -w</code> and <code>-print_jobid</code>.</p>	<p>N/A</p>

Table 6-2 Known limitations (*continued*)

Limitation	Workaround
<p>MongoDB restores are not supported from the backup hosts. All the restore operations for MongoDB must be initiated from the NetBackup primary.</p>	<p>N/A</p>
<p>If your restore job submission does not display the restore job, verify that your destination node has a MongoDB plug-in that is installed on it.</p>	<p>N/A</p>
<p>If you restore the MongoDB database to a non-LVM location and then try to take a backup from this non-LVM location, the backup fails.</p>	<p>Restore the data to an LVM location and then try to take a backup of the restored data.</p>
<p>The NetBackup for MongoDB plug-in does not support hard or soft links in the data path folders. Do not add any hard or any soft links that point to locations in a different logical volume or a non-logical volume.</p> <p>NetBackup cannot ensure that the data is consistent at the time of backups if you have hard or soft links in the data path folder. During the restore process, the hard or the soft links are created as folders and not links.</p>	<p>N/A</p>
<p>When you cancel a child restore job during the MongoDB restore and recovery process, the thin client (<code>mdbserver</code>) is not removed immediately. The thin client is removed after the next restore operation.</p>	<p>N/A</p>
<p>MongoDB restore fails and displays error 2850.</p>	<p>Consider the following solutions:</p> <ul style="list-style-type: none"> ■ Ensure that the destination host and port are valid and that the credentials were configured with the <code>tpconfig</code> command and the credentials file. For more information, refer to the <code>tar</code> logs. ■ The target database path does not exist and there are insufficient permissions for the non-root user. Workaround: Ensure that the target database path exists and there are sufficient permissions for the non-root user. ■ Ensure that there are no special characters in the <code>rename</code> and <code>filelist</code> file. Also, if the primary server is a Windows computer then make sure that the EOL conversion of the file is <code>Unix Style (LF)</code>.

Table 6-2 Known limitations (*continued*)

Limitation	Workaround
<p>After recovery, the MongoDB shard node fails to restart manually and the following error is seen in the MongoDB logs:</p> <pre>NoSuchKey: Missing expected field "configsvrConnectionString"</pre>	<p>On the MongoDB shard where the problem occurs, start MongoDB in the maintenance mode and run the following method on the <code>system.version</code> collection in the admin database:</p> <pre>use admin db.system.version.deleteOne ({ _id: "minOpTimeRecovery" })</pre>
<p>In a restore operation that contains one or more replica sets, replica set members are restored to the replica set that uses the default <code>"cfg.members[#].host"</code> value that <code>rs.config()</code> provides.</p> <p>If this value was previously updated from the default value after the restore and recover completes, this value may need to be updated to match the original configuration. (For example, from shortname to FQDN.)</p>	<p>Workaround:</p> <ol style="list-style-type: none"> 1 Log on to the replica set MongoDB cluster 2 Use the following command to verify the configuration: <pre>rs.config()</pre> 3 Use the following command to update the configuration for the replica set: <pre>Update configuration for replica set member 0: cfg = rs.config(); cfg.members[0].host = '<hostname.domain.com>: <port-number>'; rs.reconfig(cfg)</pre> 4 Verify the changes using the following command: <pre>rs.config()</pre> 5 Repeat the steps for the other replica sets and the members, or only the replica set members.

Table 6-2 Known limitations (continued)

Limitation	Workaround
<p>Backup jobs fail and the following error codes are displayed:</p> <ul style="list-style-type: none"> ■ (50) client process aborted ■ (1) The requested operation was partially successful. ■ (112) no files specified in the file list 	<p>Ensure that the backup windows for incremental backups are different for the same MongoDB cluster. The backup windows must not overlap each other for incremental backups for the same MongoDB cluster.</p> <p>Ensure that permissions are in place for the <code>mdbserver</code> location, <code>oplog</code> location, and snapshot mount location. For more information, See “Host user requirements” on page 36.</p> <p>In a sharded MongoDB cluster environment, a 112 error can indicate that the <code>mongos</code> process is not running on the client that is defined in the backup policy.</p> <p>An error 112 can also indicate that same hosts names for multiple backup hosts are added to the BigData policy. Use unique host names for multiple backup hosts that are running the backup operations.</p>
<p>After a restore operation, if you try to stop and restart the <code>mongod</code> or <code>mongos</code> services (service <code>mongod stop</code> or <code>service mongod restart</code>), the commands fail.</p> <p>This error occurs when the <code>mongod</code> or <code>mongos</code> processes are launched as service using the <code>service</code> or <code>systemctl</code> commands and not using a direct command.</p>	<p>Workaround:</p> <p>Stop the <code>mongod</code> or <code>mongos</code> services using alternative methods. For example, <code>mongod -f /etc/mongod.conf --shutdown</code> or <code>kill <PID></code>. After stopping the services, you can use the <code>service</code> or <code>systemctl</code> commands again.</p> <p>Note: When you stop the services after restore and recovery, the <code>.pid</code> or <code>.sock</code> files remain when you shut down the <code>mongod</code> or <code>mongos</code> processes. You must delete the files if the <code>mongod</code> or <code>mongos</code> services do not start after shutting them down.</p> <p>The default location of the <code>.sock</code> files is <code>/tmp</code></p> <p>The default location of the <code>.pid</code> files is <code>/var/run/mongodb/</code></p>
<p>Backup operation fails if a command that generates output in <code>.bashrc</code> is added.</p> <p>Backup fails with error 6646 and displays the following error:</p> <p>Error: Unable to communicate with the server.</p>	<p>Ensure that no output is generated by <code>.bashrc</code> (<code>echo</code> or any other output generating command). The output should not return <code>STDERR</code> or <code>STDOUT</code> when the shell is non-interactive.</p>

Table 6-2 Known limitations (*continued*)

Limitation	Workaround
<p>When you select two full backup images and try to restore to a point-in-time image that is between the two full backup images, the latest full backup image is restored.</p>	<p>Workaround:</p> <p>During the restore operation, do not select more than one full backup image.</p> <p>For an effective point-in-time recovery, ensure that you run differential incremental backups of shorter duration.</p>
<p>Unable to see the restore job progress in the Recover node.</p>	<p>Workaround:</p> <p>For a compound restore job that uses a non-primary server as the restore host, look for the job record and status in the Task progress section in the Recover node. Click Refresh to refresh the task list.</p>
<p>Backup fails with the following error:</p> <p>(6625) The backup host is either unauthorized to complete the operation or it is unable to establish a connection with the application server.</p>	<p>Workaround:</p> <p>On the server where MongoDB is installed, ensure that <code>PasswordAuthentication</code> is not disabled in <code>/etc/ssh/sshd_config</code> file.</p> <p>Run the <code>sudo service sshd restart</code> command.</p>
<p>Backup fails with the following error:</p> <p>(6646) Unable to communicate with the server.</p>	<p>Workaround:</p> <p>Ensure that the backup host can access the defined port in <code>mongodb.conf</code> file or the default <code>mongodbserver_port</code> (11000).</p> <p>There can be an error when you copy the thin client files on the MongoDB server because of the following issues:</p> <ul style="list-style-type: none"> ■ Connectivity issues with the MongoDB server ■ User does not have permissions to the location for copying the thin-client files.
<p>The following error is displayed in the <code>mongodbserver</code> logs:</p> <pre>error-sudo: sorry, you must have a tty to run sudo</pre>	<p>Workaround:</p> <ul style="list-style-type: none"> ■ To disable the <code>requiretty</code> option globally in the <code>sudoers</code> file, replace <code>Defaults requiretty</code> with <code>Defaults !requiretty</code>. This action changes the global <code>sudo</code> configuration. ■ You can change the <code>sudo</code> configuration for the user, group, or command. On the server where MongoDB is installed, add the host user, or group, or command in the <code>sudoers</code> file. <pre>Add Defaults /path/to/my/bin !requiretty Add Default <host_user> !requiretty</pre>

Table 6-2 Known limitations (*continued*)

Limitation	Workaround
<p>The <code>nbaapireq_handler</code> log folder is not created on a Flex Container, even you run the <code>mkdir</code> command.</p>	<p>Workaround:</p> <p>When a Flex Appliance is upgraded from version 8.1.2 to 8.2 and the Flex media server is used as backup host, the MongoDB plug-in creates the following log directory:</p> <pre>/usr/openv/netbackup/logs/nbaapireq_handler</pre>
<p>The snapshot size as described by the <code>free_space_percentage_snapshot</code> parameter must be set according to the MongoDB cluster size and must be large enough. If these criteria are not met, the backup fails and displays the following error:</p> <pre>invalid command parameter (20)</pre>	<p>Validate the <code>free_space_percentage_snapshot</code> value with the MongoDB cluster.</p>
<p>Backup fails with the following error:</p> <pre>(13) file read failed for Media</pre>	<p>Ensure that the:</p> <ul style="list-style-type: none"> ■ NetBackup version on the primary server is the latest. ■ NetBackup version on the media server is the same as the primary server but newer than the NetBackup client version on the backup host. ■ NetBackup client version on the backup host is the same as or older than the media server.
<p>The <code>mdb_progress_loglevel</code> parameter is missing from the MongoDB configuration tool.</p>	<p>To modify the <code>mdb_progress_loglevel</code> parameter, update the <code>mongodb.conf</code> file after the MongoDB configuration tool creates it.</p> <p>For more information, refer to the <i>MongoDB Administrator's Guide</i>.</p>
<p>Snapshots are not deleted and stale <code>mdbserver</code> instances are seen. This scenario might cause <code>Cannot lstat</code> errors during backup and partially successful backups.</p>	<p>Change the configuration settings for the following parameters in the <code>mongodb.conf</code> file:</p> <ul style="list-style-type: none"> ■ <code>cleanup_time_in_min</code> ■ <code>mdbserver_timeout_min</code> <p>Set the values such that the stale snapshots and stale instances of <code>mdbserver</code> are cleared before the next full or incremental backup schedule.</p>

Table 6-2 Known limitations (*continued*)

Limitation	Workaround
<p>If the backup host has NetBackup version earlier than 8.3 and primary and media server have the latest version of NetBackup, the following invalid error codes can be seen for various scenarios:</p> <p>13302, 13303, 13304, 13305, 13306, 13307, 13308, 13309, 13310, 13311, 13312, 13313, 13314, 13315</p>	

Table 6-2 Known limitations (*continued*)

Limitation	Workaround
	<p>Workaround:</p> <p>Refer to the following list of corresponding actual error codes if you see the invalid error codes for the actual scenarios and recommended actions:</p> <ul style="list-style-type: none"> ■ Invalid error code: 13302 Actual error: 6724 Message: Restore node count is invalid. ■ Invalid error code: 13303 Actual error: 6725 Message: Unable to find information about the MongoDB replica set. ■ Invalid error code: 13304 Actual error: 6704 Message: Restoring multiple MongoDB nodes on one replica set is invalid. ■ Invalid error code: 13305 Actual error: 6705 Message: Restoring MongoDB data on an arbiter node is invalid. ■ Invalid error code: 13306 Actual error: 6706 Message: A discovered shard was found in a drain state, cannot proceed with backup. ■ Invalid error code: 13307 Actual error: 6707 Message: An unsupported MongoDB storage engine is detected. ■ Invalid error code: 13308 Actual error: 6708 Message: Unable to parse command output ■ Invalid error code: 13309 Actual error: 6709 Message: Unable to run the command. ■ Invalid error code: 13310 Actual error: 6710 Message: Pre-check for recovery has failed as WiredTiger log files are present at

Table 6-2 Known limitations (*continued*)

Limitation	Workaround
	<p>the database path.</p> <ul style="list-style-type: none"> ■ Invalid error code: 13311 Actual error: 6711 Message: Unable to backup MongoDB configuration file. ■ Invalid error code: 13312 Actual error: 6712 Message: Unable to find operation log for previous backup. ■ Invalid error code: 13313 Actual error: 6713 Message: Operations log roll-over detected. ■ Invalid error code: 13314 Actual error: 6714 Message: Error while collection was iterated. ■ Invalid error code: 13315 Actual error: 6715 Message: Operation log verification error. <p>For detailed information and recommended actions, refer to the <i>NetBackup Status Codes Reference Guide</i>.</p>
The Restore button in the NetBackup web UI is disabled for the imported MongoDB backup images.	<p>Workaround:</p> <p>If you import the images to the same NetBackup primary server that was originally used to back them up, use either of the following methods:</p> <ul style="list-style-type: none"> ■ Perform the restore operation using the <code>bprestore</code> command. ■ Restore the catalog backup that enables the Restore button in the web UI and then restore the images. <p>If you import the images to a different NetBackup primary server than the one that was originally used to back them up, use the <code>bprestore</code> command to run the restore operation.</p>

Table 6-2 Known limitations (*continued*)

Limitation	Workaround
Recovery operation fails on an alternate, sharded MongoDB cluster. The following error is displayed: Unable to find the configuration parameter. (6661)	

Table 6-2 Known limitations (*continued*)

Limitation	Workaround
	<p>This issue occurs during an alternate cluster recovery because the pre-recovery check is unable to find the <code>mongos</code> port for the alternate cluster in the <code>mongodb.conf</code> file. This issue occurs because of the way the MongoDB configuration tool creates the <code>mongodb.conf</code> file when you add the alternate MongoDB cluster details using the Update option from the tool.</p> <p>Workaround:</p> <p>Before you start the recovery process, update the <code>mongodb.conf</code> file to separate the alternate cluster from the original cluster.</p> <p>For example:</p> <p>Existing <code>mongodb.conf</code> file:</p> <pre> "application_servers": { "original.mongodb.cluster.com:26050": { "alternate_config_server": [{ "hostname:port": "alt.mongodb.cluster.com:26000", "mongos_port": "26001" }], "mongos_port": "26051" } } </pre> <p>Suggested update to the <code>mongodb.conf</code> file:</p> <pre> "application_servers": { "original.mongodb.cluster.com:26050": { "mongos_port": "26051" }, "alt.mongodb.cluster.com:26000": { "mongos_port": "26001" } } </pre>

Table 6-2 Known limitations (*continued*)

Limitation	Workaround
	<pre> } } </pre>
<p>The MUI tool displays the following error: Unable to delete configuration.</p>	<p>Recommended action:</p> <ul style="list-style-type: none"> ■ Verify that the <code><hostname-port>.conf</code> file still exists in the <code>/usr/opensv/var/global</code> directory. ■ Refer to the <code>tpconfig</code> logs and check for error: Translate <code>EMM_ERROR_MachineNotExist(2000000)</code> to 88 in the Device Config context. <p>Work Around:</p> <p>Delete the <code><hostname-port>.conf</code> file manually from <code>/usr/opensv/var/global</code>.</p>
<p>In case of certificate-based authentication enabled on MongoDB, a differential incremental backup fails with error 6709: Unable to run the command.</p>	<p>Workaround:</p> <p>Refer to the <code>mdbserver</code> logs to find the error code and command details. Then perform one of the following actions:</p> <ul style="list-style-type: none"> ■ If <code>mdbserver</code> logs indicate <code>mongodump</code> command failure, try running <code>mongodump</code> command manually on the MongoDB host and check the error. ■ If <code>mongodump</code> command fails with X509 certificate-related connection errors, make sure to fix these errors by updating MongoDB server certificates with <code>subjectAltName</code> property as per MongoDB documentation. Then re-run the differential incremental backup.

Additional information

This appendix includes the following topics:

- [Sample MongoDB configuration utility workflow to add and update MongoDB credentials](#)

Sample MongoDB configuration utility workflow to add and update MongoDB credentials

Adding MongoDB credentials

```
Device Management Configuration Utility
```

- 1) Drive Configuration
- 2) Robot Configuration
- 3) Credentials Configuration
- 4) MongoDB Configuration
- 5) Print Configuration
- 6) Help
- 7) Quit

```
Enter option :4
```

```
MongoDB Application Configuration
```

- 1) Configure MongoDB Application Topology & Credentials
- 2) Configure NetBackup Global Parameters for MongoDB Application
- 3) Quit

```
Enter option :1
```

```
Configure the MongoDB cluster credentials
```

- 1) ADD Credentials

Sample MongoDB configuration utility workflow to add and update MongoDB credentials

- 2) UPDATE Credentials
- 3) DELETE Credentials
- 4) Return to previous menu

Select the operation :1

Please select your MongoDB cluster type.

- 1) Standalone node
- 2) Sharded Cluster
- 3) Replica set
- 4) Return to main menu

Select the type of your MongoDB cluster :3

Select MongoDB host credentials type

- 1) No Auth
- 2) Simple Auth
- 3) Certificate based
- 4) Return to main menu

Select the authentication type used in the MongoDB cluster :2

Configure Replica Set MongoDB Cluster

Enter the hostname of primary server : host1.fqdn.com

Enter the mongod port of primary server [On the MongoDB Shell, run the command "rs.status()" for replica set and "sh.status()" for sharded environment] : 28000

Enter the name of MongoDB host user : root

Enter the password of MongoDB host user :

Enter the RSA key of the MongoDB host [On the MongoDB host, run the command "cat /etc/ssh/ssh_host_rsa_key.pub | awk '{print \$2}' | base64 -d | openssl dgst -sha256 | awk '{print \$2}'] : RSA-KEY-OF-THE-HOST

Enter MongoDB database user : mongodb-shell-login-user

Enter MongoDB database user password :

Does this primary server has replicas?(y/n) :y

Enter the hostname of secondary server : host2.fqdn.com

Enter the mongod port of secondary server [On the MongoDB Shell, run the command "rs.status()" for replica set and "sh.status()" for sharded environment] : 28001

Enter the name of MongoDB host user : root

Enter the password of MongoDB host user :

Sample MongoDB configuration utility workflow to add and update MongoDB credentials

```

Enter the RSA key of the MongoDB host [On the MongoDB host, run the
command "cat /etc/ssh/ssh_host_rsa_key.pub | awk '{print $2}' | base64
-d| openssl dgst -sha256 | awk '{print $2}']: RSA-KEY-OF-THE-HOST
Enter MongoDB database user : mongodb-shell-login-user
Enter MongoDB database user password :

```

```

Do you have more secondary servers for this primary server? (y/n) :n

```

Summary is displayed after you add the credentials.

```

-----REPLICA SET MONGODB CONFIGURATION SUMMARY-----

```

```

Primary Server :
  Server Hostname           : host1.fqdn.com
  Server Mongod Port       : 28000
  No of Secondary Servers   : 1
  HostUser: root
  HostPassword: *****
  AppUserId: mongodb-shell-login-user
  AppUserPassword: *****
  HostRsaKey: RSA-KEY-OF-THE-HOST

```

```

Secondary Server number 1:
  Secondary Server Hostname       : host2.fqdn.com
  Secondary Server Mongod Port    : 28001
  HostUser: root
  HostPassword: *****
  AppUserId: mongodb-shell-login-user
  AppUserPassword: *****
  HostRsaKey: RSA-KEY-OF-THE-HOST

```

```

*****Please make sure to save this entered config and credentials. If you
don't save it now, you will have to enter it again.*****

```

```

Do you want to save this cluster configuration and credential info?(y/n) :

```

Please wait while we save the cluster configuration.

Successfully saved config and credentials for this cluster.

Please use Client name as "host1.fqdn.com-28000" under 'Clients' tab in

mongodb backup policy.

Press any key to to return to main menu...

Updating MongoDB credentials

Device Management Configuration Utility

- 1) Drive Configuration
- 2) Robot Configuration
- 3) Credentials Configuration
- 4) MongoDB Configuration
- 5) Print Configuration
- 6) Help
- 7) Quit

Enter option :4

MongoDB Application Configuration

- 1) Configure MongoDB Application Topology & Credentials
- 2) Configure NetBackup Global Parameters for MongoDB Application
- 3) Quit

Enter option :1

Configure the MongoDB cluster credentials

- 1) ADD Credentials
- 2) UPDATE Credentials
- 3) DELETE Credentials
- 4) Return to previous menu

Select the operation :2

Please select your MongoDB cluster type.

- 1) Standalone node
- 2) Sharded Cluster
- 3) Replica set
- 4) Return to main menu

Select the type of your MongoDB cluster :3

Update replica set MongoDB cluster configuration

Enter the hostname of primary server : host1.fqdn.com

Enter the mongod port of primary server [On the MongoDB Shell, run the command "rs.status()" for replica set and "sh.status()" for

Sample MongoDB configuration utility workflow to add and update MongoDB credentials

```
sharded environment] : 28000
```

```
[Note- similar steps can be followed for deleting creds for cluster]
```

```
--
```

```
Update host1.fqdn.com:28000 replica set MongoDB cluster configuration
```

- 1) Update primary server credentials
- 2) Add secondary server
- 3) Update secondary server config & credentials
- 4) Delete secondary Replica server
- 5) Return to previous menu

```
Enter option: option as applicable
```