



Backup for MongoDB Documentation

2.3.1 (December 4, 2023)

Percona Technical Documentation Team

Percona LLC, © 2023

Table of contents

1. Percona Backup for MongoDB Documentation	5
1.1 What you can do	5
1.2 What's in it for you?	6
1.3 Go further with Percona Backup for MongoDB	6
1.4 Get expert help	6
2. Release notes	7
2.1 Percona Backup for MongoDB release notes	7
2.2 Percona Backup for MongoDB 2.3.1 (2023-12-04)	9
2.3 Percona Backup for MongoDB 2.3.0 (2023-09-25)	10
2.4 Percona Backup for MongoDB 2.2.1 (2023-07-27)	12
2.5 Percona Backup for MongoDB 2.2.0 (2023-07-13)	13
2.6 Percona Backup for MongoDB 2.1.0 (2023-04-18)	15
2.7 Percona Backup for MongoDB 2.0.5 (2023-03-23)	17
2.8 Percona Backup for MongoDB 2.0.4 (2023-02-21)	18
2.9 Percona Backup for MongoDB 2.0.3 (2023-01-11)	19
2.10 Percona Backup for MongoDB 2.0.2 (2022-10-27)	21
2.11 Percona Backup for MongoDB 2.0.1 (2022-10-12)	22
2.12 Percona Backup for MongoDB 2.0.0 (2022-09-21)	24
2.13 Percona Backup for MongoDB 1.8.1 (2022-07-12)	26
2.14 Percona Backup for MongoDB 1.8.0 (2022-06-09)	27
2.15 Percona Backup for MongoDB 1.7.0 (2022-04-18)	29
2.16 <i>Percona Backup for MongoDB 1.6.1 (2021-11-04)</i>	31
2.17 Percona Backup for MongoDB 1.6.0	33
2.18 Percona Backup for MongoDB 1.5.0	35
2.19 Percona Backup for MongoDB 1.4.1	36
2.20 Percona Backup for MongoDB 1.4.0	37
2.21 Percona Backup for MongoDB 1.3.4	38
2.22 <i>Percona Backup for MongoDB 1.3.3</i>	39
2.23 Percona Backup for MongoDB 1.3.2	40
2.24 Percona Backup for MongoDB 1.3.1	42
2.25 Percona Backup for MongoDB 1.3.0	43
2.26 <i>Percona Backup for MongoDB 1.2.1</i>	44
2.27 <i>Percona Backup for MongoDB 1.2.0</i>	45
2.28 Percona Backup for MongoDB 1.1.3	46
2.29 Percona Backup for MongoDB 1.1.1	47

2.30	Percona Backup for MongoDB 1.1.0	48
2.31	Percona Backup for MongoDB 1.0.0	50
2.32	Percona Backup for MongoDB 0.5.0	51
3.	Features	54
3.1	Backup and restore types	55
3.2	Logical backups and restores	57
3.3	Physical backups and restores	59
3.4	Selective backup and restore	62
3.5	Incremental physical backups	64
3.6	Point-in-time recovery	66
3.7	Snapshot-based physical backups	69
3.8	Comparison with MongoDB	73
4.	Discover PBM	75
4.1	How Percona Backup for MongoDB works	75
4.2	Architecture	77
4.3	Authentication	88
4.4	Supported MongoDB deployments	90
4.5	Supported MongoDB versions	91
5.	Get started	92
5.1	Install	92
5.2	Initial setup	105
5.3	Make a backup	113
5.4	Restore a backup	119
5.5	Make a point-in-time restore	125
6.	Manage backups	130
6.1	List backups	130
6.2	View detailed information about a backup	132
6.3	Schedule backups	135
6.4	Cancel a backup	137
6.5	Delete backups	138
6.6	View restore progress	144
6.7	Replay oplog from arbitrary start time	146
6.8	View backup logs	148
7.	Administer PBM	150
7.1	Percona Backup for MongoDB configuration via pipelines	150
7.2	Automate access to S3 buckets for Percona Backup for MongoDB	152
7.3	Percona Backup for MongoDB diagnostics tools	154
7.4	Percona Backup for MongoDB status	158

7.5 Upgrade Percona Backup for MongoDB	160
7.6 Uninstall Percona Backup for MongoDB	166
8. FAQ	167
8.1 What's the difference between PBM and mongodump?	167
8.2 Why does Percona Backup for MongoDB use UTC timezone instead of server local timezone?	167
8.3 Can I restore a single collection with Percona Backup for MongoDB?	167
8.4 Can I back up specific shards in a cluster?	167
8.5 Do I need to stop the balancer for PITR restore?	167
8.6 Get expert help	168
9. Reference	169
9.1 pbm commands	169
9.2 Configuration file options	185
9.3 Glossary	199
10. Contribution guidelines	203
11. Submit bug reports or feature requests	203
11.1 Submit a code patch	203
11.2 Contribute to documentation	203
11.3 Get expert help	204
12. Copyright and licensing information	205
12.1 Documentation licensing	205
12.2 Get expert help	205
13. Trademark Policy	206
13.1 Get expert help	207

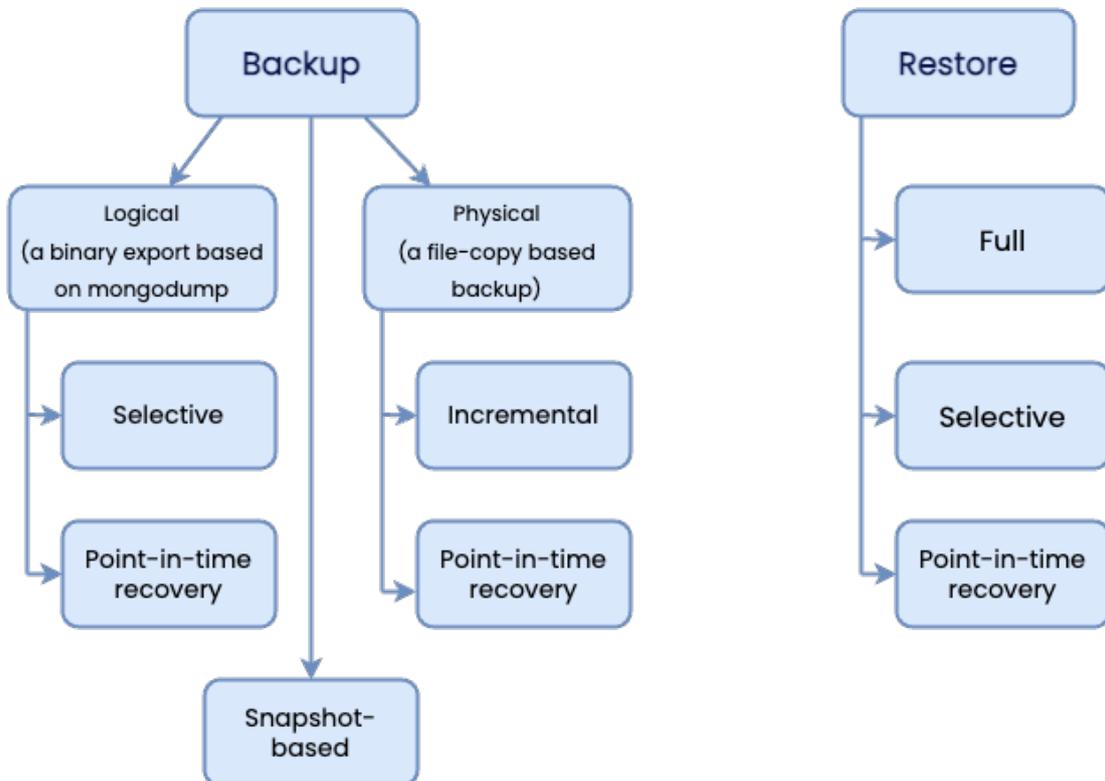
1. Percona Backup for MongoDB Documentation

Percona Backup for MongoDB (PBM) is an open source and distributed solution for consistent backups and restore of [MongoDB sharded clusters and replica sets](#). There is no notable performance nor operating degradation associated with PBM.

With Percona Backup for MongoDB, you can make backups on a running server and restore your database to a specific point in time using the command line. To do these tasks from a user interface, [use PBM with Percona Monitoring and Management](#).

This is the documentation for the latest release, **PBM 2.3.1** ([Release Notes](#)).

1.1 What you can do



- [Logical backups](#) to back up and / or migrate data to different platforms and database versions
- [Physical backups](#) to speed up performance for large (multi-terabyte) data sets
- [Selective backups](#) to work with the desired data set
- [Incremental physical backups](#) to ensure that critical data is regularly backed up and to save on costs for storage and transfer
- [Snapshot-based physical backups](#) using the technology of your choice.
- [Restore the full database or specific data set](#) from a backup
- [Restore the database to a specific point in time](#)
- [Replay oplog](#) on top of EBS-snapshots

[Explore features](#)

1.2 What's in it for you?

- Data consistency across clusters and replica sets
- A variety of [supported storage types](#) means no vendor lock-in
- Open source solution with [enterprise-grade features](#)

[Install and get started](#)

1.3 Go further with Percona Backup for MongoDB

- [Learn how Percona Backup for MongoDB works](#)
- [Manage PBM](#)
- [Contribute to PBM](#)

1.4 Get expert help

If you need assistance, visit the [community forum](#) for comprehensive and free database knowledge, or contact our [Percona Database Experts](#) for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: July 13, 2023

Created: August 18, 2022

2. Release notes

2.1 Percona Backup for MongoDB release notes

- [Percona Backup for MongoDB 2.3.1](#)
- [Percona Backup for MongoDB 2.3.0](#)
- [Percona Backup for MongoDB 2.2.1](#)
- [Percona Backup for MongoDB 2.2.0](#)
- [Percona Backup for MongoDB 2.1.0](#)
- [Percona Backup for MongoDB 2.0.5](#)
- [Percona Backup for MongoDB 2.0.4](#)
- [Percona Backup for MongoDB 2.0.3](#)
- [Percona Backup for MongoDB 2.0.2](#)
- [Percona Backup for MongoDB 2.0.1](#)
- [Percona Backup for MongoDB 2.0.0](#)
- [Percona Backup for MongoDB 1.8.1](#)
- [Percona Backup for MongoDB 1.8.0](#)
- [Percona Backup for MongoDB 1.7.0](#)
- [Percona Backup for MongoDB 1.6.1](#)
- [Percona Backup for MongoDB 1.6.0](#)
- [Percona Backup for MongoDB 1.5.0](#)
- [Percona Backup for MongoDB 1.4.1](#)
- [Percona Backup for MongoDB 1.4.0](#)
- [Percona Backup for MongoDB 1.3.4](#)
- [Percona Backup for MongoDB 1.3.3](#)
- [Percona Backup for MongoDB 1.3.2](#)
- [Percona Backup for MongoDB 1.3.1](#)
- [Percona Backup for MongoDB 1.3.0](#)
- [Percona Backup for MongoDB 1.2.1](#)
- [Percona Backup for MongoDB 1.2.0](#)
- [Percona Backup for MongoDB 1.1.3](#)
- [Percona Backup for MongoDB 1.1.1](#)
- [Percona Backup for MongoDB 1.1.0](#)
- [Percona Backup for MongoDB 1.0.0](#)
- [Percona Backup for MongoDB 0.5.0](#)

2.1.1 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: December 4, 2023

Created: August 18, 2022

2.2 Percona Backup for MongoDB 2.3.1 (2023-12-04)

Installation

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a non-sharded replica set), and for restoring those backups to a specific point in time.

2.2.1 Release Highlights

- Added support for Percona Server for MongoDB 7.0.
- The ability to define custom endpoints when using Microsoft Azure Blob Storage for backups
- Improved PBM Docker image to allow making physical backups with the shared `mongodb` data volume
- Updated Golang libraries that include fixes for the security vulnerability CVE-2023-39325.

2.2.2 Improvements

- [PBM-1195](#) - Allow custom endpointUrl for Azure storage
- [PBM-1205](#) - Change the user in PBM Docker image to `mongod`

2.2.3 Bugs Fixed

- [PBM-1193](#) - Fixed the `pbm-agent` failure when the connection to PSMDB is lost

2.2.4 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: December 4, 2023

Created: December 4, 2023

2.3 Percona Backup for MongoDB 2.3.0 (2023-09-25)

Installation

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a non-sharded replica set), and for restoring those backups to a specific point in time.

2.3.1 Release Highlights

- The support for MongoDB 4.2 is deprecated. Existing functionality in Percona Backup for MongoDB remains compatible with MongoDB 4.2 and Percona Server for MongoDB 4.2; however, further enhancements and bug fixes are no longer tested against this version.
- The ability to [view the backup contents](#) improves troubleshooting of backups in environments where databases are often created and / or dropped.
- The ability to [make physical backups in mixed deployments](#) with MongoDB Community and Percona Server for MongoDB (PSMDB) nodes streamlines the backup flow for organizations that are still evaluating or migrating their data sets against PSMDB.
- Improved validation of a base backup snapshot for point-in-time recovery aligns the recovery flow for logical and physical backups.

Supported platforms changes

- A Docker image for Percona Backup for MongoDB is now available for ARM64 architectures. This improves the user experience with PBM for developers with ARM-based MacBooks.
- Percona Backup for MongoDB packages are now available for Debian 12 (bookworm). Percona Server for MongoDB is adding support for Debian 12 with the 7.0.x version.

2.3.2 New Features

- [PBM-848](#) - List databases and collections from backup

2.3.3 Improvements

- [PBM-1147](#), [PBM-885](#) - Support symlinks as the `path` for filesystem storage type and properly retrieve data during resync
- [PBM-1162](#), [PBM-1166](#) - Improve validation of physical and incremental backups as the base for point-in-time recovery
- [PBM-1171](#) - Add ability to make physical backups in mixed deployments

2.3.4 Bugs Fixed

- [PBM-1161](#) - Rebuild indexes only for the selected namespaces after selective restore
- [PBM-1173](#), [PBM-1180](#) - Fixed the issue with the primary node crash after the point-in-time recovery from physical backup
- [PBM-1176](#) - Fixed the decoding of the physical restore metadata to ensure correct output of the `pbm describe-restore` command
- [PBM-1181](#) - Fixed the issue with logical restore failure by excluding the `config.changelog` collection from the restore

2.3.5 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: September 25, 2023

Created: September 25, 2023

2.4 Percona Backup for MongoDB 2.2.1 (2023-07-27)

Release date	July 27, 2023
Installation	Installing Percona Backup for MongoDB

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a non-sharded replica set), and for restoring those backups to a specific point in time.

2.4.1 Release Highlights

With this release you can increase the wait time for backup to start, thus eliminating the PBM failure when creating `$backupCursor` takes longer than usual.

2.4.2 Improvements

- [PBM-1154](#) - Add config option to change wait timeout for backup start status

2.4.3 Bugs Fixed

- [PBM-1151](#) - Fixed the issue with PBM crashing when Percona Server for MongoDB fails to open `$backupCursor`
- [PBM-1155](#) - Retry to open `$backupCursor` during incremental backups in case of failure

2.4.4 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: July 27, 2023

Created: July 27, 2023

2.5 Percona Backup for MongoDB 2.2.0 (2023-07-13)

Release date	July 13, 2023
Installation	Installing Percona Backup for MongoDB

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a non-sharded replica set), and for restoring those backups to a specific point in time.

2.5.1 Release Highlights

- [Point-in-time recovery from physical backups](#) is now automated similar to point-in-time recovery from logical ones. This offloads your DBAs on performing manual oplog replay on top of physical restore, ensures data consistency and unifies the user experience with PBM.
- Owners of large data sets can now use PBM to [create external physical backups](#) as EBS snapshots or via a technology of their choice and restore from those backups with the data consistency guaranteed by PBM. Thereby they benefit from increased performance and reduced downtime, and are sure that their data remains consistent. This is the technical preview feature.
- The ability to [restore from physical and incremental backups to a new environment](#) with different replica set names extends the set of compatible environments for physical restore.

2.5.2 New Features

- [PBM-952](#) - Add ability to make external backups
- [PBM-815](#) - Automated point-in-time recovery for physical backups

2.5.3 Improvements

- [PBM-1055](#) - Add ability to configure replica set remapping for physical and incremental restores
- [PBM-1118](#) - Support custom shard names for restores with replica set remapping
- [PBM-1132](#) - Disallow using the `--ns` flag for physical and incremental backups and restores
- [PBM-1133](#) - Add support for virtual-hosted-style URL for s3 storage buckets
- [PBM-1062](#) - Add a supported PSMDB version check per backup type

2.5.4 Bugs Fixed

- [PBM-1090](#) - Fixed the issue with the `E11000 duplicate key` error by applying the indexes building after the oplog replay operation
- [PBM-1095](#) - Fixed the authentication error during selective backup by changing the `authSource` to `admin` on config server replica set
- [PBM-1103](#) - Changed the warning message for incompatible backups in the `pbm status` output.
- [PBM-1123](#) - Fixed the issue with the restore by skipping the `config.image_collection` collection
- [PBM-1135](#) - Fixed the issue with physical backup failure on a sharded cluster

2.5.5 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: July 13, 2023

Created: July 13, 2023

2.6 Percona Backup for MongoDB 2.1.0 (2023-04-18)

Release date	April 18, 2023
Installation	Installing Percona Backup for MongoDB

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a non-sharded replica set), and for restoring those backups to a specific point in time.

2.6.1 Release Highlights

- [Incremental physical backups](#) are now generally available enabling you to use them in production environments. Note that due to the changes in metadata files required for the restore, backups made with previous PBM versions are incompatible for the restore with PBM 2.1.0.
- You can now [selectively back up and restore sharded collections](#). This improves the management of a desired subset of data in sharded clusters and saves you extra costs on data storage and transfer. This is the tech preview feature due to some [known limitations](#).
- Benefit from up to 19 times faster physical restore performance with the support of [parallel download of data chunks from the S3 storage](#).
- Improved deletion of old backups and point-in-time recovery oplog chunks simplifies the automation of backup storage cleanup.
- The improved handling of master keys for data at rest encryption in Percona Server for MongoDB and the retrieval of the key ID/secret path by PBM from a backup simplifies the environment preparation for the physical restore and improves the restore flow.
- The support of AWS tokens for the access to the S3 storage improves the security of your infrastructure and the integration with applications that interact with AWS resources via tokens.

2.6.2 New Features

- [PBM-1007](#) - Selective backup/restore for sharded collections

2.6.3 Improvements

- [PBM-850](#) - Add ability to delete older backups
- [PBM-875](#) - Concurrent download from S3 compatible storages
- [PBM-778](#) - Save physical restore logs on storage
- [PBM-951](#) - Get rid of redundant `pbm*.old` collections to improve resync speed
- [PBM-955](#) - Improve physical restore of data encrypted at rest
- [PBM-1033](#) - Flush previous incremental backup history after a new base backup is made
- [PBM-1034](#) - Restore only files listed in a target backup and remove unneeded files for incremental backups
- [PBM-1035](#) - Indicate a base incremental backup in `pbm status` output
- [PBM-1046](#) - Handle changes with an offset beyond the current file size for incremental backups
- [PBM-1053](#) - Stop mongod on shards before config servers during physical restore
- [PBM-1060](#) - Add support of AWS session tokens to access S3 storage
- [PBM-1066](#) - Restore from logical backup made on previous major PSMDB version

- [PBM-1070](#) - Improve handling of failed physical/incremental restores when running `pbm restore` with the `--wait` option
- [PBM-1078](#) - Preserve vault options during intermediate restarts
- [PBM-1085](#) - Add the `--yes` flag for commands where `--force` is used

2.6.4 Bugs Fixed

- [PBM-979](#) - Add support for delayed/arbiter nodes on physical restore
- [PBM-1012](#) - Drop a collection during selective oplog replay on the create collection stage
- [PBM-1030](#) - Fixed the restore from incremental backup for database with data-at-rest encryption if the data were added/removed during backup
- [PBM-1052](#) - Fixed PBM misbehavior during physical restore on a cluster that has more shards than in a backup
- [PBM-1061](#) - Fixed the restore with the replica set name remapping on PSMDB 4.2
- [PBM-1063](#) - Fixed the incremental restore failure if the `directoryPerDb` is configured in PSMDB
- [PBM-1069](#) - Check for incompatible incremental backups and notify accordingly in `pbm status` output
- [PBM-1073](#) - Improve performance of `pbm status` command on environments with large datasets and many backups
- [PBM-1077](#) - Save the `storage.bson` to ensure incremental backups and restores for database with data-at-rest encryption

2.6.5 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: July 13, 2023

Created: April 18, 2023

2.7 Percona Backup for MongoDB 2.0.5 (2023-03-23)

Release date	March 23, 2023
Installation	Installing Percona Backup for MongoDB

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a non-sharded replica set), and for restoring those backups to a specific point in time.

2.7.1 Release Highlights

This release introduces the following fixes:

- for the physical restore process for deployments where the `mongod --dbpath` option has a forward slash (`/`) as the last char
- security vulnerability [CVE-2022-41723](#) by updating the Golang library dependencies to the latest versions

2.7.2 Bugs Fixed

- [PBM-1058](#) - Fixed the physical restore issue if the `--dbpath` option has a trailing slash.
- [PBM-1076](#) - Update `golang.org/x/net` module due to [CVE-2022-41723](#)

2.7.3 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: March 29, 2023

Created: March 23, 2023

2.8 Percona Backup for MongoDB 2.0.4 (2023-02-21)

Release date	February 21, 2023
Installation	Installing Percona Backup for MongoDB

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a non-sharded replica set), and for restoring those backups to a specific point in time.

2.8.1 Release Highlights

The ability to [specify the custom path to mongod binaries](#) simplifies the physical restore process.

2.8.2 Bugs Fixed

- [PBM-1042](#) - Fixed the physical restore flow with the ability to set the custom path to `mongod` binaries
- [PBM-1047](#) - Fixed the issue with physical restore failure in a new environment.

2.8.3 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: September 25, 2023

Created: February 21, 2023

2.9 Percona Backup for MongoDB 2.0.3 (2023-01-11)

Release date	January 11, 2023
Installation	Installing Percona Backup for MongoDB

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a non-sharded replica set), and for restoring those backups to a specific point in time.

2.9.1 Release Highlights

- [Incremental physical backups](#) enable you to reduce storage costs and facilitate data safety for business crucial data. By saving only the differences results in faster completion time and makes IB much smaller in size compared to full backups. As such, you save on storage space and data transfer in case of cloud setups. This is the [technical preview feature](#), yet we welcome your feedback to improve the functionality.
- Now you can selectively back up and restore unsharded collections in sharded clusters. This extends the data set to work with.
- Added support of AWS IRSA (Identity Roles for Service Accounts) credentials allows Percona Backup for MongoDB running in a pod to access the AWS storage using the IAM roles. This increases the security of your cloud infrastructure and enables you to control the access to it from a single place.
- Percona Backup for MongoDB is now available on Red Hat Enterprise Linux 9 and compatible derivatives

2.9.2 New Features

- [PBM-999](#) - Incremental physical backups support
- [PBM-962](#) - Add support for selective backup / restore of unsharded collections in a sharded cluster

2.9.3 Improvements

- [PBM-998](#), [PBM-1002](#) - Ensure that PBM waits for the `mongod` to stop / start completely during physical restore
- [PBM-1009](#) - Add AWS IRSA credentials support (Thanks to the user `pastequo` for contributing to this feature)
- [PBM-1018](#) - Include the `pbm-agent-entrypoint` binary into the `percona-backup-mongodb` package and tarball

2.9.4 Bugs Fixed

- [PBM-932](#) - Fix the data access from `mongos` after selective restore on a new cluster
- [PBM-981](#) - Prevent showing the ERROR status while backup is running
- [PBM-1003](#) - Fix the ability to restore with remapping replica set names (Thanks to user `Piotr Krogulec` for reporting this issue and to user `ziollek` for contributing to it)
- [PBM-1006](#) - Fix the `pbm-agent` crash during physical restore by initializing the logger for storage
- [PBM-1008](#) - Fix the selective PITR restore for a collection created after the full backup
- [PBM-1016](#) - Fix the error handling for `pbm describe-backup` command

2.9.5 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: March 29, 2023

Created: January 11, 2023

2.10 Percona Backup for MongoDB 2.0.2 (2022-10-27)

Release date	October 27, 2022
Installation	Installing Percona Backup for MongoDB

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a non-sharded replica set), and for restoring those backups to a specific point in time.

This release fixes the usability issues for applications operating with Percona Backup for MongoDB by providing the error messages for the status output in the JSON format ([PBM-988](#)). It also includes the fix for the security vulnerability issue [CVE-2022-32149](#).

2.10.1 Bugs Fixed

[PBM-988](#): Show error message for JSON output of the `pbm status` command

2.10.2 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: October 27, 2022

Created: October 27, 2022

2.11 Percona Backup for MongoDB 2.0.1 (2022-10-12)

Release date	October 12, 2022
Installation	Installing Percona Backup for MongoDB

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a non-sharded replica set), and for restoring those backups to a specific point in time.

2.11.1 Release Highlights

- The support of [server-side encryption with customer-provided keys managed on the customer side \(SSE-C\)](#) enables you to use the S3-compatible storage of your choice thus preventing the vendor lock-in and saving your costs on AWS KMS (Key Management Service).
- The ability to [configure Percona Backup for MongoDB remotely](#) simplifies its management when PBM is deployed in Docker, Kubernetes or other cloud services.
- The ability to configure the sidecar mode for Percona Backup for MongoDB improves its operation as part of [Kubernetes Operator for MongoDB](#).
- Troubleshooting enhancements:
 - The ability to [define a timezone for logs and to follow the logs dynamically](#).
 - Indication of arbiter nodes as non-supported ones in `pbm status` output

2.11.2 New Features

- [PBM-733](#): Allow changing the timezone of logs in PBM
- [PBM-877](#): Enable PBM to read the config file from stdin

2.11.3 Improvements

- [PBM-949](#): Improve handling of arbiter nodes
- [PBM-948](#): Align `pbm describe-backup` and `pbm describe-restore` outputs
- [PBM-941](#): Add a sidecar mode for PBM to prevent the database crash when a `pbm-agent` is misconfigured in a container
- [PBM-922](#): Add support for SSE-C for S3 providers (Thanks to Richard Bateman for reporting this issue and contributing to it)
- [PBM-897](#): Add the ability to dynamically follow PBM logs

2.11.4 Bugs Fixed

- [PBM-978](#): Fixed the physical restore on replica sets having hidden nodes
- [PBM-975](#): Fixed the issue with the display of point-in-time recovery chunks without base snapshot after storage resync
- [PBM-966](#): Fixed the error handling of point-in-time recovery to non-existent timestamp by `pbm-agents`
- [PBM-858](#): Improve reporting of the error when an oplog chunk can't be created due to insufficient range

2.11.5 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: October 12, 2022

Created: October 12, 2022

2.12 Percona Backup for MongoDB 2.0.0 (2022-09-21)

Release date	September 21, 2022
Installation	Installing Percona Backup for MongoDB

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a non-sharded replica set), and for restoring those backups to a specific point in time.

2.12.1 Release Highlights

- Physical backups and restores are now generally available. This enables you to use them in production environments.
- [Data-at-rest encryption](#) is supported for physical backups and restores. This enables you to comply to data security regulations and save time on operating with large data sets.
- By [tracking physical restore progress](#), you have a clear picture of your restore operations and can timely react to any changes or issues.
- [Logical backups and restores](#) can now be done selectively. This is a tech preview feature ¹ yet it enables you to work only with the desired subset of data and thereby save time on database maintenance and costs on storage.

2.12.2 New Features

- [PBM-790](#): Add the ability to make selective backups and restores. This is the tech preview feature ¹

2.12.3 Improvements

- [PBM-896](#): Enable physical restores for data encrypted at rest
- [PBM-898](#), [PBM-811](#), [PBM-812](#): Log temporary mongod process during physical restore
- [PBM-911](#): Rename backup's complete time in `pbm status` and `pbm list` outputs
- [PBM-813](#): Enable `pbm-agents` to send heartbeats via storage during physical restore

2.12.4 Bugs Fixed

- [PBM-909](#): Consider the storage settings during physical restore
- [PBM-943](#): Fix the display of `pbm status` for filesystem storage
- [PBM-929](#): Recreate UUID for timeseries collections during `oplog apply`

2.12.5 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

1. Tech Preview Features are not yet ready for enterprise use and are not included in support via SLA. They are included in this release so that users can provide feedback prior to the full release of the feature in a future GA release (or removal of the feature if it is deemed not useful). This functionality can change (APIs, CLIs, etc.) from tech preview to GA. ←←

Last update: March 29, 2023

Created: September 21, 2022

2.13 Percona Backup for MongoDB 1.8.1 (2022-07-12)

Release date	July 12, 2022
Installation	Installing Percona Backup for MongoDB

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a non-sharded replica set), and for restoring those backups to a specific point in time.

2.13.1 Release Highlights

- [PBM-871](#) - Fixed the restore failure on a different cluster. Now the UUID of users and system collections are not preserved when replaying the oplog.
- [PBM-881](#) - The point-in-time recovery chunks display is now consistent in both `pbm status` and `pbm list` outputs.

2.13.2 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: August 18, 2022

Created: August 18, 2022

2.14 Percona Backup for MongoDB 1.8.0 (2022-06-09)

Release date	June 9, 2022
Installation	Installing Percona Backup for MongoDB

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a single replica set), and for restoring those backups to a specific point in time.

2.14.1 Release Highlights

- Ability to restore data to a replica set with a different name and configuration. This extends the list of environments compatible for the restore.
- When you use EBS-snapshots or other tools for physical backups, you no longer have to create a mandatory base backup snapshot in [Percona Backup for MongoDB] as the starting point for Point-in-Time Recovery oplog slicing. This reduces time and effort on managing excessive backups and makes Point-in-Time Recovery from physical or storage-level backups more straightforward.
- The ability to wait for the backup operation to finish before doing further actions through the session lock. This simplifies the automation of operations with Percona Backup for MongoDB.
- Ability to define backup compression level and method in Percona Backup for MongoDB configuration.
- To simplify the Percona Backup for MongoDB configuration, the example configuration file is now included in the Percona Backup for MongoDB package.
- Ubuntu 22.04 (Jammy Jellyfish) is added to the list of [supported platforms](#)

2.14.2 New Features

- [PBM-776](#): Allow data restore into the replica set with a different name
- [PBM-866](#): Add the ability to wait for the backup operation to finish and print the result
- [PBM-782](#): Allow saving Point-in-Time Recovery oplog without base snapshot
- [PBM-838](#): Add the ability to configure default compression method and level for backups

2.14.3 Improvements

- [PBM-828](#): Add the full reference configuration file to packages
- [PBM-751](#): Format timestamps according to [RFC3339](#) (Thanks to Damiano Albani for reporting this issue)

2.14.4 Bugs Fixed

- [PBM-820](#): Fix a bug where PBM crashed if backup cancelled right after it started by cancelling the backup gracefully

2.14.5 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 **Community Forum**  **Get a Percona Expert**

Last update: August 18, 2022

Created: August 18, 2022

2.15 Percona Backup for MongoDB 1.7.0 (2022-04-18)

Release date	April 18, 2022
Installation	Installing Percona Backup for MongoDB

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a single replica set), and for restoring those backups to a specific point in time.

2.15.1 Release highlights

- Support for physical backups in Percona Server for MongoDB starting from versions 4.2.15-16 and 4.4.6-8 and higher. Physical backups drastically speed up backup and restore performance for large databases (several terabytes). This is a technical preview feature .
- Olog replay from the arbitrary start time. This reduces Recovery Point Objective (RPO) when database is recovered from physical or storage-level backups.
- Ability to configure compression method and level for Point-in-Time Recovery chunks and compression level for backups.
- Ability to configure the number of S3 multipart upload chunks to comply with various S3-compatible storage provider requirements.
- Ability to configure the number of upload retries. This facilitates data upload in case of unstable network connection.

2.15.2 New Features

- [PBM-734](#): Add the config option to set debug log levels for S3 requests
- [PBM-805](#): Implement physical backups to improve performance for large databases
- [PBM-742](#): Add the ability to replay olog from arbitrary start time. This reduces Recovery Point Objective (RPO) when database is recovered from physical backups.

2.15.3 Improvements

- [PBM-680](#): Skip TLS verification for object storage. This can be useful for private object storage with self-signed certificates.
- [PBM-770](#): Support configurable compression method / level for Point-in-Time Recovery chunks (Thanks to Damiano Albani for reporting this issue and contributing to it)
- [PBM-764](#): Support Zstandard compression format (Thanks to Damiano Albani for reporting this issue and contributing to it)
- [PBM-750](#): Make max number of S3 upload parts configurable (Thanks to Damiano Albani for reporting this issue and contributing to it)
- [PBM-777](#): Expand / fix the configuration API to support compression method for Point-in-Time Recovery chunks (Thanks to Damiano Albani for reporting and contributing to this issue)
- [PBM-756](#): Add the ability to configure logging levels for S3 requests to debug issues with object storage (Thanks to Damiano Albani for reporting this issue and contributing to it)
- [PBM-577](#): It is now possible to choose an S3 storage class for granular control over various S3 tiers (Thanks to Damiano Albani for the contribution)

2.15.4 Bugs Fixed

- **PBM-721:** Fixed a bug where an upload of the backup to S3-storage was failing due to unstable network connection. Percona Backup for MongoDB can now be configured to retry the upload with flexible timeouts.
- **PBM-773:** Check distributed transactions on all participating shards to avoid commit timestamp inconsistency upon restore

2.15.5 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: August 18, 2022

Created: August 18, 2022

2.16 Percona Backup for MongoDB 1.6.1 (2021-11-04)

Release date	November 4, 2021
Installation	Installing Percona Backup for MongoDB

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a single replica set), and for restoring those backups to a specific point in time.

2.16.1 Release Highlights

- Deprecated support for MongoDB 3.6. Percona Backup for MongoDB remains compatible with MongoDB 3.6 and Percona Server for MongoDB 3.6; however, further enhancements and bug fixes are no longer tested against this version.
- Improved backup and PITR routines alignment by using sequential `delete-pitr / install-backup` operations instead of in-memory backup intent. This fixes the inability of a backup to start.
- Added support for automated access to S3 buckets using an EC2 instance profile. When Percona Backup for MongoDB is deployed using an EC2 instance, EC2 environment variables and metadata are used for S3 authentication, saving you from explicitly specifying S3 credentials in the Percona Backup for MongoDB configuration file. This comes handy for architectures deployed using the services like Amazon EC2, kiam, kube2iam or irsa.
- Extended logging for `pbm-agents`. This improves user experience with Percona Backup for MongoDB.

2.16.2 Improvements

- [PBM-740](#): Use AWS EC2 instance profile to simplify access to S3 buckets for PBM

2.16.3 Bugs Fixed

- [PBM-714](#): Fix backup and point-in-time recovery routines alignment algorithm to avoid backup failure
- [PBM-722](#): Fix `pbm-agent`'s crash during the `delete-pitr` request execution if there is nothing to delete (Thanks to Daniel Oliver for reporting this issue)
- [PBM-735](#): Fix a possible failure of a PITR catchup process to copy backup slices
- [PBM-712](#): Fix an empty time value in JSON formatted log records by using Unix timestamps for time output

2.16.4 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: August 18, 2022

Created: August 18, 2022

2.17 Percona Backup for MongoDB 1.6.0

Release date	August 16, 2021
Installation	Installing Percona Backup for MongoDB

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a single replica set), and for restoring those backups to a specific point in time.

2.17.1 Release Highlights:

- Support for Percona Server for MongoDB and MongoDB Community 5.0
- Point-in-time recovery enhancements: ability to restore from any previous snapshot and configurable span of oplog events
- JSON output for PBM commands to simplify interfacing PBM with applications

2.17.2 New Features

- [PBM-617](#): Ability to restore from previous snapshots to point-in-time

2.17.3 Improvements

- [PBM-543](#): Configure the size of the span of oplog events for point-in-time recovery
- [PBM-403](#): Mask user credentials in `ps` output of `pbm-agent`
- [PBM-700](#): Improve backup/pitr tasks synchronization and align oplogs creation
- [PBM-697](#): Add support of MongoDB 5.0 TS collections
- [PBM-686](#): Do not show the starting second of a PITR range which cannot be used for PITR restore
- [PBM-652](#): Add a command to delete PITR chunks
- [PBM-632](#): Add JSON output for all commands

2.17.4 Bugs Fixed

- [PBM-694](#): Fix restoring from a backup when it contains VIEWS collection (Thanks to Danish Qamar for reporting this issue)
- [PBM-647](#): Reduce frequency of S3 header GET requests during agent health checks (Thanks to Ryan Gunner for reporting this issue)
- [PBM-708](#): Ignore `config.system.indexBuilds` collection
- [PBM-705](#): Avoid writing the "Read/Write on closed pipe" error in logs on expected connection closure
- [PBM-703](#): PITR restore fails due to error "Failed to apply operation due to missing collection `config.transactions`"
- [PBM-701](#): Prevent restore to time which is not covered by PITR chunks
- [PBM-683](#): Show PITR restore as failed if an error occurred during data retrieval from storage
- [PBM-640](#): Remove `cancelBackup` and fix `pitrestore` filters for `pbm logs` command
- [PBM-480](#): Make `path` attribute mandatory for backups on local storage

2.17.5 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: August 18, 2022

Created: August 18, 2022

2.18 Percona Backup for MongoDB 1.5.0

Release date	May 10, 2021
Installation	Installing Percona Backup for MongoDB

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a single replica set), and for restoring those backups to a specific point in time.

2.18.1 New Features

- [PBM-596](#): Azure Blob Storage support
- [PBM-488](#): Create weight or tag method to influence with `pbm-agent` node will do backups

2.18.2 Improvements

- [PBM-662](#): Show PITR Status based on `admin.pbmLock` instead of config settings
- [PBM-494](#): Prefer a (healthy) hidden secondary to any other node in automatic selection

2.18.3 Bugs Fixed

- [PBM-642](#): Display `priority=0` members on agent list in `pbm status` output
- [PBM-636](#): Different collection UUID after restore (Thanks to Nikolay for reporting this issue and Dmitry Kuzmin for contributing)
- [PBM-646](#): Stop the balancer during backup to make sure it doesn't start running during restore
- [PBM-635](#): Wait for the leader's metadata before starting backups
- [PBM-490](#): Use cluster time for the snapshot start time

2.18.4 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: August 18, 2022

Created: August 18, 2022

2.19 Percona Backup for MongoDB 1.4.1

Release date	January 28, 2021
Installation	Installing Percona Backup for MongoDB

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a single replica set), and for restoring those backups to a specific point in time.

2.19.1 Improvements

- [PBM-621](#): Show incomplete backups in `pbm status` output
- [PBM-619](#): Optimize response time from storage for `pbm status`
- [PBM-615](#): Check backup validity for current cluster
- [PBM-608](#): Enable Kerberos authentication for PBM by adding support for GSSAPI
- [PBM-478](#): Prevent restore from incomplete backup
- [PBM-610](#): Fix response time from GCS for `pbm status` command

2.19.2 Bugs Fixed

- [PBM-618](#): Check for the complete file set in backup snapshot before processing it

2.19.3 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: August 18, 2022

Created: August 18, 2022

2.20 Percona Backup for MongoDB 1.4.0

Release date	December 24, 2020
Installation	Installing Percona Backup for MongoDB

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a single replica set), and for restoring those backups to a specific point in time.

2.20.1 New Features

- [PBM-345](#): Centralize logs
- [PBM-435](#): `pbm status` command

2.20.2 Improvements

- [PBM-572](#): Change backup 'name' in 'pbm list' etc to be consistent time (~= end time) rather than start time
- [PBM-556](#): Introduce operation ID

2.20.3 Bugs Fixed

- [PBM-595](#): Shard backup with different rset name
- [PBM-604](#): Compression flag for 'pbm list' command doesn't change the output
- [PBM-602](#): Empty PITR files are created on storage if PBM fails to upload oplog chunk due to insufficient range
- [PBM-597](#): Properly handle mongo fail while PITR slicing is enabled

2.20.4 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: May 9, 2023
Created: August 18, 2022

2.21 Percona Backup for MongoDB 1.3.4

Release date	November 19, 2020
Installation	Installing Percona Backup for MongoDB

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a single replica set), and for restoring those backups to a specific point in time.

2.21.1 Improvements

- [PBM-586](#): Add a request timeout to the S3 downloader during the restore
- [PBM-584](#): Ignore shard configuration during the restore

2.21.2 Bugs Fixed

- [PBM-555](#): Fix the " error demultiplexing archive" error during restore by downloading backup from s3 storage in chunks
- [PBM-460](#): Restore fails with conflicting namespace destinations (Thanks to user pedroalb for reporting this issue)

2.21.3 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: August 18, 2022

Created: August 18, 2022

2.22 Percona Backup for MongoDB 1.3.3

Release date	November 4, 2020
Installation	Installing Percona Backup for MongoDB

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a single replica set), and for restoring those backups to a specific point in time.

2.22.1 Bugs Fixed

- [PBM-575](#): `mongodump` connects to the primary node

2.22.2 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: August 18, 2022

Created: August 18, 2022

2.23 Percona Backup for MongoDB 1.3.2

Release date	October 14, 2020
Installation	Installing Percona Backup for MongoDB

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a single replica set), and for restoring those backups to a specific point in time.

2.23.1 New Features

- [PBM-426](#): Add AWS KMS key encryption/decryption for S3 buckets

Config format

```
storage:
  s3:
    serverSideEncryption:
      sseAlgorithm: "aws:kms"
      kmsKeyID: "....."
```

(Thanks to user pedroalb for reporting this issue)

2.23.2 Improvements

- [PBM-568](#): Print uploadPartSize value to log during backup
- [PBM-560](#): Use s2 compression as default for `pbm-speed-test` instead of gzip

2.23.3 Bugs Fixed

- [PBM-485](#): Fix backups to S3 failing with `MaxUploadParts` limit by auto-adjusting `uploadPartSize` value (Thanks to user pedroalb for reporting this issue)
- [PBM-559](#): `pbm-agent` runs out of memory while doing restore of large backup (Thanks to user Simon Bernier St-Pierre for reporting this issue)
- [PBM-562](#): Correct calculation of available PITR time ranges by `pbm list`
- [PBM-561](#): Fix setting of numeric options in config
- [PBM-547](#): Allow deleting backups from local filesystem by moving delete operations to `pbm-agents`

2.23.4 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: August 18, 2022

Created: August 18, 2022

2.24 Percona Backup for MongoDB 1.3.1

Release date	September 3, 2020
Installation	Installing Percona Backup for MongoDB

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a single replica set), and for restoring those backups to a specific point in time.

2.24.1 Bugs Fixed

- [PBM-542](#): Fix backup folder permissions on filesystem storage for Point-in-Time recovery

2.24.2 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: August 18, 2022

Created: August 18, 2022

2.25 Percona Backup for MongoDB 1.3.0

Release date	August 26, 2020
Installation	Installing Percona Backup for MongoDB

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a single replica set), and for restoring those backups to a specific point in time.

2.25.1 New Features

- [PBM-455](#): Add oplog archiver thread for PITR
- [PBM-491](#): Modify “pbm restore” to accept arbitrary point in time when PITR oplog archives available

2.25.2 Improvements

- [PBM-526](#): Add pbm version information to the backup metadata

2.25.3 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: August 18, 2022

Created: August 18, 2022

2.26 Percona Backup for MongoDB 1.2.1

Release date	July 27, 2020
Installation	Installing Percona Backup for MongoDB

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a single replica set).

2.26.1 Bugs Fixed

- [PBM-509](#): Include “pbm-speed-test” binary for debian packages

2.26.2 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: August 18, 2022

Created: August 18, 2022

2.27 Percona Backup for MongoDB 1.2.0

Release date	May 13, 2020
Installation	Installing Percona Backup for MongoDB

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a single replica set).

2.27.1 New Features

- [PBM-348](#): Add ability to delete old backups
- [PBM-447](#): `pbm-speed-test`: Add a tool to field-test compression and upload speeds

2.27.2 Improvements

- [PBM-431](#): Raise dump output speed through compression tuning, parallelization
- [PBM-461](#): `s2` is set as the default compression mechanism
- [PBM-429](#): Periodic backup progress messages added to `pbm-agent` logs
- [PBM-140](#): Added ability to cancel a backup

2.27.3 Bugs Fixed

- [PBM-451](#): Resync didn't work if storage type was set to filesystem

2.27.4 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: August 18, 2022

Created: August 18, 2022

2.28 Percona Backup for MongoDB 1.1.3

Release date	April 14, 2020
Installation	Installing Percona Backup for MongoDB

2.28.1 Improvements

- [PBM-424](#): Remove the `--mongodb-uri` arg from `pbm-agent.service` unit file
- [PBM-419](#): Resolve restore-blocking issues related to `admin.system.version`
- [PBM-417](#): Improve `pbm` control collection etc. metadata for restores

2.28.2 Bugs Fixed

- [PBM-425](#): `pbm-agent` could fail when restoring
- [PBM-430](#): S3 store resync didn't work if the store had a prefix
- [PBM-438](#): `pbm list --size=5` worked in reverse

2.28.3 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: August 18, 2022

Created: August 18, 2022

2.29 Percona Backup for MongoDB 1.1.1

Release date	January 31, 2020
Installation	Installing Percona Backup for MongoDB

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a single replica set), and for restoring those backups to a specific point in time. The project was inspired by (and intends to replace) the Percona-Lab/mongodb_consistent_backup tool.

Percona Backup for MongoDB supports [Percona Server for MongoDB](#) or [MongoDB Community Server](#) version 3.6 or higher with [MongoDB replication](#) enabled. Binaries for the supported platforms as well as the tarball with source code are available from the [Percona Backup for MongoDB download page](#). For more information about Percona Backup for MongoDB and the installation steps, see the documentation.

2.29.1 Bugs Fixed

- [PBM-407](#): Very large collections experienced timeout due to full-collection scan for a preliminary count
- [PBM-414](#): The upload on Google cloud storage was broken with "InvalidArgument: Invalid argument. status code: 400"
- [PBM-409](#): Restore failed with "incompatible auth version with target server"

2.29.2 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: August 18, 2022

Created: August 18, 2022

2.30 Percona Backup for MongoDB 1.1.0

Percona is happy to announce the release of Percona Backup for MongoDB 1.1.0 on January 16, 2020.

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a single replica set), and for restoring those backups to a specific point in time. The project was inspired by (and intends to replace) the Percona-Lab/mongodb_consistent_backup tool.

Percona Backup for MongoDB supports [Percona Server for MongoDB](#) or [MongoDB Community Server](#) version 3.6 or higher with [MongoDB replication](#) enabled. Binaries for the supported platforms as well as the tarball with source code are available from the [Percona Backup for MongoDB download page](#). For more information about Percona Backup for MongoDB and the installation steps, see the documentation.

Percona Backup for MongoDB 1.1.0 introduces the new `pbm config` command to enable configuring the store from the command line in addition to the configuration file. This command effectively replaces `pbm store` which was only able to read store configuration from the configuration file.

```
$ pbm config --set storage.s3.bucket="operator-testing"
```

2.30.1 New Features

- [PBM-344](#): New `pbm config` command to support configuring the store from the command line.

2.30.2 Improvements

- [PBM-361](#): Improved the processing of timestamps when using `oplog`.

2.30.3 Bugs Fixed

- [PBM-214](#): `pbm-agent` could crash with `restore` command running forever, if the primary node became unavailable during the restore operation.
- [PBM-279](#): `pbm-agent` could be started with an invalid config file.
- [PBM-338](#): Backups that failed could appear in the output of the `pbm list` command.
- [PBM-362](#): The `pbm backup` could fail when called from the primary node if there were no healthy secondaries.
- [PBM-369](#): ReplicaSets could not establish connections when TLS was used in the cluster.

2.30.4 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: August 18, 2022

Created: August 18, 2022

2.31 Percona Backup for MongoDB 1.0.0

Percona is happy to announce the GA release of our latest software product Percona Backup for MongoDB 1.8 on September 19, 2019.

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a single replica set), and for restoring those backups to a specific point in time. The project was inspired by (and intends to replace) the Percona-Lab/mongodb_consistent_backup tool.

Percona Backup for MongoDB supports [Percona Server for MongoDB](#) or [MongoDB Community Server](#) version 3.6 or higher with [MongoDB replication](#) enabled. Binaries for the supported platforms as well as the tarball with source code are available from the [Percona Backup for MongoDB download page](#). For more information about Percona Backup for MongoDB and the installation steps, see the documentation.

Percona Backup for MongoDB 1.0.0 features the following:

- The architecture and the authentication of Percona Backup for MongoDB have been simplified compared to the previous release.
- Stores backup data on [Amazon Simple Storage Service](#) or compatible storages, such as [MinIO](#).
- The output of **pbm list** shows all backups created from the connected MongoDB sharded cluster or replica set.

2.31.1 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: August 18, 2022

Created: August 18, 2022

2.32 Percona Backup for MongoDB 0.5.0

Percona is pleased to announce the early release of Percona Backup for MongoDB 0.5.0 of our latest software product on June 17, 2019. The GA version of Percona Backup for MongoDB is scheduled to be released later in 2019.

Percona Backup for MongoDB is a distributed, low-impact solution for consistent backups of MongoDB sharded clusters and replica sets. This is a tool for creating consistent backups across a MongoDB sharded cluster (or a single replica set), and for restoring those backups to a specific point in time. Percona Backup for MongoDB uses a distributed client/server architecture to perform backup/restore actions.

The project was inspired by (and intends to replace) the [Percona-Lab/mongodb_consistent_backup tool](#).

Percona Backup for MongoDB supports Percona Server for MongoDB or MongoDB Community Server version 3.6 or higher with MongoDB replication enabled. Binaries for the supported platforms as well as the tarball with source code are available from the [Percona Backup for MongoDB download page](#). For more information about Percona Backup for MongoDB and the installation steps, see the [documentation](#).

Percona Backup for MongoDB 0.5.0 features the following:

- Enables storing backup metadata on Amazon Simple Storage Service storages.
- The API of Percona Backup for MongoDB introduces HTTP basic authentication to prevent an unauthorized user from running backups or restoring data if they manage to access the API port.
- To optimize the usage of network resources, the `pbm-agent` on `mongos` is not needed any more and backup-coordinator automatically establishes connection to the appropriate `mongos` instance.
- The output of `pmbctl list nodes` now includes the replica set name and informs the backup status of the node.

Percona doesn't recommend this release for production as its API and configuration fields are still likely to change. It only features a basic API level security. Please report any bugs you encounter in [our bug tracking system](#).

2.32.1 New Features and Improvements

- **93:** Support storage of backup metadata on AWS S3.
- **99:** **pbm-agent** is deprecated on `mongos`.
- **105:** Log a warning if a Primary node-type is used for a backup
- **122:** Include the replica set name to the output of `pmbctl list nodes`
- **130:** Add HTTP Basic Authentication to gRPC servers (API and RPC)
- **139:** Support listing backup status in the output of `pmbctl list nodes`
- **170:** Enable setting the 'stopOnError' attribute in `mongorestore` to ensure consistency of the data being restored.

2.32.2 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 **Community Forum**

 **Get a Percona Expert**

Last update: January 19, 2023

Created: August 18, 2022

3. Features

3.1 Backup and restore types

You can use several types of database backups and restores to develop your backup strategy. The following table provides an overview of each type

Type	Status	Version added	Supported deployments	Percona Server for MongoDB	MongoDB Community / Enterprise Edition
Logical	GA	1.0.0	Sharded clusters and non-sharded replica sets	Yes	Yes
Physical	GA	2.0.0	Sharded clusters and non-sharded replica sets	Yes (starting with 4.2.15-16, 4.4.6-8 and higher, 5.0.x, 6.0.x)	No
Selective	Tech preview	2.0.0	Sharded clusters (starting with version 2.0.3) and non-sharded replica sets. Sharded collections are not yet supported.	Yes	Yes
Incremental	GA	2.1.0 (Supported for PSMDB versions 4.2.24-24, 4.4.18-18, 5.0.2-1, 6.0.2-1 and higher)	Sharded clusters and non-sharded replica sets	Yes	No
Point-in-time recovery	GA	1.3.0 (logical), 2.0.0 (physical manually), 2.2.0 (physical automated)	Sharded clusters and non-sharded replica sets	Yes	Yes
Snapshot-based	Tech preview	2.2.0	Sharded clusters and non-sharded replica sets	Yes	No

3.1.1 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: July 13, 2023

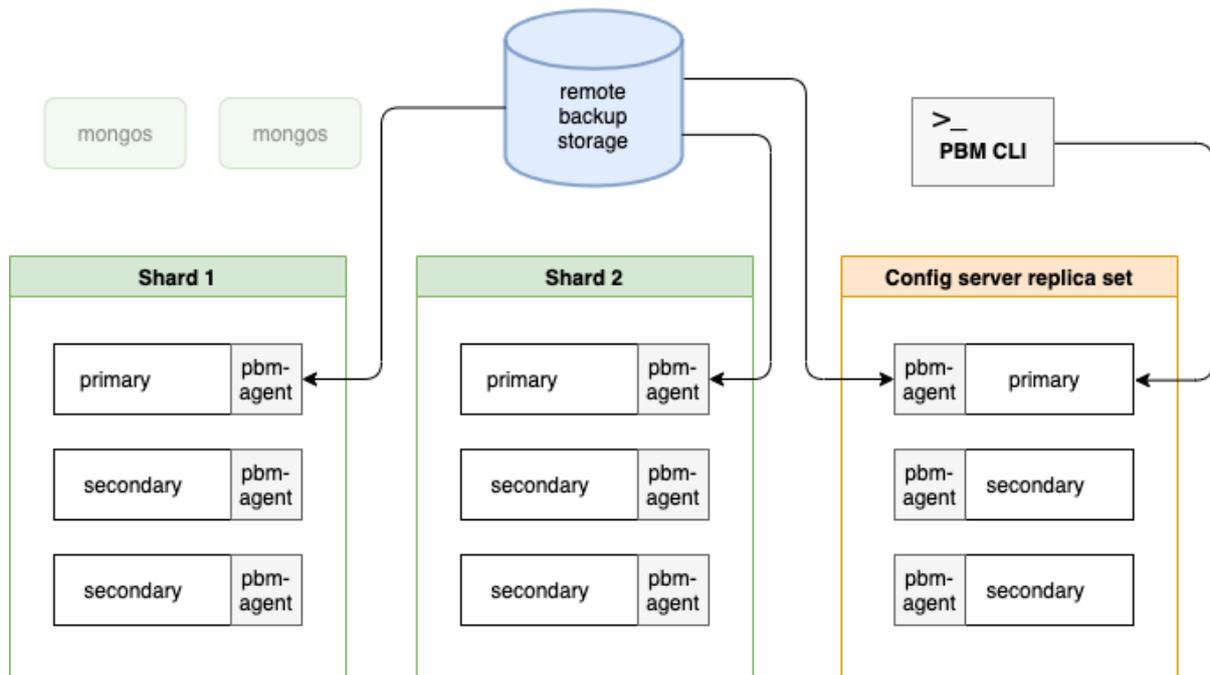
Created: August 18, 2022

3.2 Logical backups and restores

Logical backup is the copying of the actual database data. A `pbm-agent` connects to the database, retrieves the data, and writes it to the remote backup storage.

Logical restore is the reverse process: The `pbm-agent` retrieves the backup data from the storage and inserts it on every primary node in the cluster. The remaining nodes receive the data during the replication process.

The following diagram shows the restore flow.



Logical backups allow for point-in-time recovery.

Advantages

- Easy to operate with, using a single command
- Support for point-in-time recovery
- The backup size is smaller as it includes only the data

Disadvantages

- Much slower than physical backup / restore
- Adds database overhead on reading and inserting the data

[Make a backup](#)

[Restore a backup](#)

3.2.1 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)

 [Get a Percona Expert](#)

Last update: March 29, 2023

Created: January 19, 2023

3.3 Physical backups and restores

 **Version added: 1.7.0**

3.3.1 Availability and system requirements

- Percona Server for MongoDB starting from versions 4.2.15–16, 4.4.6–8, 5.0 and higher.
- WiredTiger is used as the storage engine in Percona Server for MongoDB, since physical backups heavily rely on the WiredTiger `$backupCursor` functionality.

 **See also**

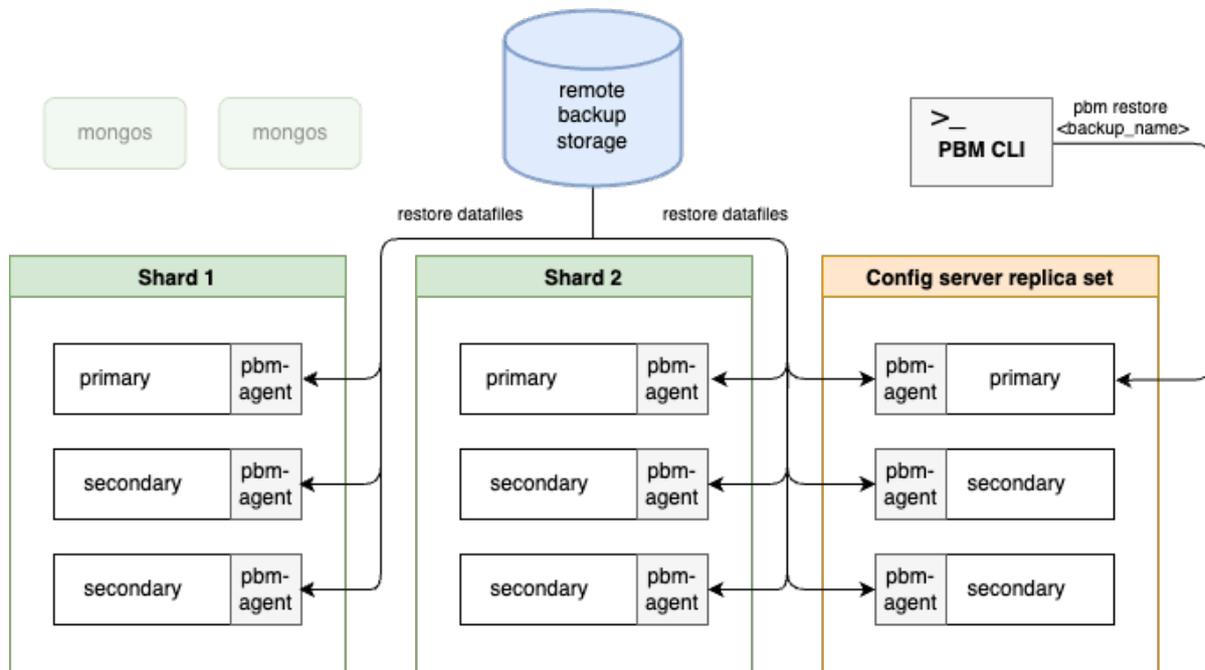
Percona Blog

- [Physical Backup Support in Percona Backup for MongoDB](#)
- [\\$backupCursorExtend in Percona Server for MongoDB](#)

Physical backup is copying of physical files from the Percona Server for MongoDB `dbPath` data directory to the remote backup storage. These files include data files, journal, index files, etc. Starting with version 2.0.0, Percona Backup for MongoDB also copies the WiredTiger storage options to the backup's metadata.

Physical restore is the reverse process: `pbm-agents` shut down the `mongod` nodes, clean up the `dbPath` data directory and copy the physical files from the storage to it.

The following diagram shows the physical restore flow:



During the restore, the `pbm-agents` temporarily start the `mongod` nodes using the the WiredTiger storage options retrieved from the backup's metadata. The logs for these starts are saved to the `pbm.restore.log` file inside the `dbPath`. Upon successful restore, this file is deleted. However, it remains for debugging if the restore were to fail.

During physical backups and restores, `pbm-agents` don't export / import data from / to the database. This significantly reduces the backup / restore time compared to logical ones and is the recommended backup method for big (multi-terabyte) databases.

Advantages

- Faster backup and restore speed
- Recommended for big, multi-terabyte datasets
- No database overhead

Disadvantages

- The backup size is bigger than for logical backups due to data fragmentation extra cost of keeping data and indexes in appropriate data structures
- Extra manual operations are required after the restore
- Point-in-time recovery requires manual operations

[Make a backup](#)

[Restore a backup](#)

3.3.2 Physical backups in mixed deployments

 **Version added: 2.3.0**

You may run both MongoDB Community / Enterprise Edition nodes and Percona Server for MongoDB (PSMDB) nodes in your environment, for example, when migrating to or evaluating PSMDB.

You can make a physical, incremental or a snapshot-based backup in such a mixed deployment using PBM. This saves you from having to reconfigure your deployment for a backup, and keeps both your migration and backup strategies intact.

Physical, incremental and snapshot-based backups are only possible from PSMDB nodes since their implementation is based on the `$backupCursorExtend` functionality. When it's time to make a backup, PBM searches the PSMDB node and makes a backup from it. The PSMDB node must not be an arbiter nor a delayed node.

If more than 2 nodes are suitable for a backup, PBM selects the one with a higher `priority`. Note that if you override a priority for at least one node, PBM assigns `priority 1.0` for the remaining nodes and uses the new priority list .

Consider the following flow for **incremental backups**: By default, PBM picks the node from where it made the incremental base backup when it makes subsequent backups. PBM assigns `priority 3.0` to this node ensuring that it is the first in the list. If you change the node priority, make a new incremental base backup to ensure data continuity.

The physical restore in mixed deployments has no restrictions except the versions in backup and in the source cluster must match.

3.3.3 Physical restores with data-at-rest encryption

 **Version added: 2.0.0**

You can back up and restore the data encrypted at rest. Thereby you ensure data safety and can also comply with security requirements such as GDPR, HIPAA, PCI DSS, or PHI.

This is how it works:

During a backup, Percona Backup for MongoDB stores the encryption settings in the backup metadata. This allows you to verify them using the `pbm describe-backup` command. Note that the encryption key is not stored nor shown.

Important

Make sure that you know what master encryption key was used and store it, as this key is required for the restore.

Starting with [Percona Server for MongoDB version 4.4.19-19, 5.0.15-13, 6.0.5-4](#) and higher, the master key rotation for data-at-rest encrypted with HashiCorp Vault has been improved to use the same secret key path on every server in your entire deployment. For the restore with earlier versions of Percona Server for MongoDB and PBM 2.0.5 and earlier, see the [Restore for Percona Server for MongoDB **before** 4.4.19-19, 5.0.15-13, 6.0.5-4 using HashiCorpVault](#) section.

To restore the encrypted data from the backup, configure data-at-rest encryption settings on all nodes of your destination cluster or replica set to match the settings of the target cluster where you made the backup

During the restore, Percona Backup for MongoDB restores the data all nodes using the same master key. To meet the security policy requirements in your organization, we recommend to rotate the master encryption keys afterwards.

To learn more about master key rotation, refer to the following documentation:

- [Master key rotation in HashiCorp Vault server](#)
- [KMIP master key rotation](#)

Restore for Percona Server for MongoDB before 4.4.19-19, 5.0.15-13, 6.0.5-4 using HashiCorpVault

In Percona Server for MongoDB version **before** 4.4.19-19, 5.0.15-13, 6.0.5-4 with the Vault server used for data-at-rest encryption, the master key rotation with the same key used for 2+ nodes is not supported. If you run these versions of Percona Server for MongoDB and PBM before 2.1.0, consider using the scenario where PBM restores the data on one node of every replica set. The remaining nodes receive the data during the initial sync.

Here's how it works:

Configure data-at-rest encryption on one node of every shard in your destination cluster or a replica set.

During the restore, Percona Backup for MongoDB restores the data on the node where the encryption key matches the one with which the backed up data was encrypted. The other nodes are not restored, so the restore has the "partially done" status. You can start this node and initiate the replica set. The remaining nodes receive the data as the result of the initial sync from the restored node.

3.3.4 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: September 25, 2023

Created: March 29, 2023

3.4 Selective backup and restore

 **Version added: 2.0.0**

 **Important**

Selective backup and restore is the [technical preview feature](#).

You can back up and restore certain namespaces - databases or collections. For example, if your “Payments” collection in the “Staff” database was corrupted, you can restore only this collection from your full backup up to a specific point in time. Or, if your “Invoices” database contains sensitive data and must be backed up frequently, you can configure the backup of only this database. This way you work only with the desired subset of data without disrupting the operations of your whole cluster.

You also drastically reduce time on backup / restore operations of the whole data set and save on storage consumption.

With the selective backup and restore functionality, you have the following options:

1. Backup a single database or a specific collection and restore all data from it.
2. Restore a specific collection from a single database backup
3. Restore certain databases and / or collections from a full backup
4. Make a point-in time recovery for the specified databases / collections.

3.4.1 Sharded collections

 **Version added: 2.1.0**

You can back up and restore sharded collections. During backup, `pbm-agents` on each shard save the documents for the specified databases/collections and the full oplog for the period of the backup process. A `pbm-agent` on the config server replica set saves router config documents from the `config` database required for restoring the selected namespaces.

During the restore, the reverse process occurs:

- A `pbm-agent` on each shard restores only the specified databases/collections and replays the oplog that relates only to the specified namespaces. The operations for other namespaces are ignored.
- On the config server replica set, the `pbm-agent` restores the router configuration only for the specified sharded collections. The router configuration for other databases, collections and chunks remains intact.

The restore for sharded time series collections is not supported.

Note that selective backups and restores operate only with data and router configuration. The cluster configuration and topology-related settings are ignored. Therefore, we recommended to restore the databases/collections on the same environment.

Implementation specifics

During the selective restore, the primary shard for a database is set to the state it had during the backup. For example, the primary shard for the database “Staff” during backup was A. After you restore the “Staff” database, the primary shard will be set to A even if you moved the primary from A to B before the restore. All

non-sharded collections will be restored on A; however, they will not be deleted from B. You must take needed actions (cleanup or move the primary back to B) to maintain them.

3.4.2 Known limitations of selective backups and restores

1. Only **logical** backups and restores are supported.
2. Selective backups and restores are supported in sharded clusters for non-sharded collections starting with version 2.0.3. Sharded collections are supported starting with version 2.1.0.
3. Sharded time series collections are not supported.
4. Multiple namespaces are not yet supported for selective backups. However, you can specify several namespaces for the restore (e.g., restore several collections of a database).
5. Multi-collection transactions are not yet supported for selective restore.
6. System collections in `admin`, `config`, and `local` databases cannot be backed up and restored selectively. You must make a full backup and restore to include them.
7. Selective point-in-time recovery is not supported for sharded clusters.

[Make a backup](#)[Restore a backup](#)

3.4.3 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: July 13, 2023

Created: March 29, 2023

3.5 Incremental physical backups

 **Version added: 2.0.3**

3.5.1 Considerations

-  Incremental backups made with Percona Backup for MongoDB prior to 2.1.0 are incompatible for restore with Percona Backup for MongoDB 2.1.0. This is because of the changed set of metadata files that are now stored in backups. These files are absent in backups made with previous PBM versions but are required for the restore with PBM 2.1.0.

We recommend to make a new incremental base backup and start the incremental backup chain from it after the upgrade to Percona Backup for MongoDB 2.1.0

- Incremental backup implementation is based on the `$backupCursor` aggregation stage that is available only in Percona Server for MongoDB. Therefore, you must be running Percona Server for MongoDB in your deployment to use incremental physical backups.
- Incremental backups are supported for Percona Server for MongoDB starting with the following versions: 4.2.24-24, 4.4.18, 5.0.14-12, 6.0.3-2 and higher.
- Due to [WiredTiger restrictions in Log-Structured Merge \(LSM\) trees](#) behavior when the `$backupCursor` is opened, incremental backups are not available if the LSM tree is configured in the database.

Owners of large datasets may need to back up data frequently. Making full physical backups every time is costly in terms of storage space. Incremental physical backups come in handy in this scenario, enabling you to optimize backup strategy and reduce storage costs.

During incremental backups, Percona Backup for MongoDB saves only the data that was changed after the previous backup was taken. This results in faster backup / restore performance. Since incremental backups are smaller in size compared to full backups, you also save on costs for their storage and transfer in case of cloud deployments.

```
graph LR
  A[Full physical ] --> B[Increment 1 ];
  B --> C[Increment 2 ];
  C --> |.....| D[Increment n ];
```

3.5.2 Implementation specifics

Percona Backup for MongoDB tracks the backup history only on the node where the base incremental backup was taken. This means that subsequent incremental backups must always be run on that very node. To make this happen, Percona Backup for MongoDB tries to schedule backups on that same node.

If the node with the base incremental backup is down or unavailable, you must start the incremental backup chain anew on another node.

[Make a backup](#)

[Restore a backup](#)

3.5.3 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 **Community Forum**  **Get a Percona Expert**

Last update: July 4, 2023

Created: March 29, 2023

3.6 Point-in-time recovery

 **Version added: 1.3.0**

Point-in-time recovery is restoring a database up to a specific timestamp. This includes restoring the data from a backup snapshot and replaying all events that occurred to this data up to a specified time from [oplog slices](#).

Advantages

Helps you prevent data loss during a disaster such as a crashed database, accidental data deletion or drop of tables, and unwanted update of multiple fields instead of a single one.

Disadvantages

Restore takes longer since it requires you to restore the backup and then replay oplog events on top of it.

3.6.1 Enable point-in-time recovery

Set the `pitr.enabled` configuration option to `true`.

Command line	Configuration file
<pre>\$ pbm config --set pitr.enabled=true</pre>	<pre>pitr: enabled: true</pre>

The `pbm-agent` starts [saving consecutive slices of the oplog](#) periodically. A method similar to the way replica set nodes elect a new primary is used to select the `pbm-agent` that saves the oplog slices. (Find more information in [pbm-agent](#).)

[Restore to a point-in-time](#)

3.6.2 Oplog slicing

To start saving [oplog slices](#), the following preconditions must be met:

Logical backups Physical backups

- A full logical backup snapshot is required. Make sure that a [backup exists](#). See the [Make a backup](#) guide to make a backup snapshot.
- Point-in-time recovery routine is [enabled](#).

Enable point-in-time recovery routine and configure it to save oplog slices without the base backup.

```
pitr:
  enabled: true
  oplogOnly: true
```

If you just enabled point-in-time recovery, it requires 10 minutes for the first chunk to appear in the `pbm list` output.

Important**For in MongoDB 5.0 and higher versions**

If you [reshard](#) a collection, make a fresh backup and re-enable point-in-time recovery oplog slicing to prevent data inconsistency and restore failure.

Olog duration**Version added: 1.6.0**

By default, a slice covers a 10-minute span of oplog events. It can be shorter if point-in-time recovery is disabled or interrupted by the start of a backup snapshot operation.

You can change the duration of an oplog span via the configuration file. Specify the new value (in minutes) for the `pittr.oplogSpanMin` option.

Command line	Configuration file
--------------	--------------------

```
$ pbm config --set pittr.oplogSpanMin=5
```

```
pittr:
  oplogSpanMin: 5
```

If you set the new duration when the `pbm-agent` is making an oplog slice, the slice's span is updated right away.

If the new duration is shorter, this triggers the `pbm-agent` to make a new slice with the updated span immediately. If the new duration is larger, the `pbm-agent` makes the next slice with the updated span in its scheduled time.

Compressed oplog slices**Version added: 1.7.0**

The oplog slices are saved with the `s2` compression method by default. You can specify a different compression method via the configuration file. Specify the new value for the `pittr.compression` option.

Command line	Configuration file
--------------	--------------------

```
$ pbm config --set pittr.compression=gzip
```

```
pittr:
  compression: gzip
```

Supported compression methods are: `gzip`, `snappy`, `lz4`, `s2`, `pgzip`, `zstd`.

Additionally, you can override the compression level used by the compression method by setting the `pittr.compressionLevel` option. The default values differ for each compression level.

Note that the higher value you specify, the more time and computing resources it will take to compress the data.

 **Note**

You can use different compression methods for backup snapshots and point-in-time recovery slices. However, backup snapshot-related oplog is compressed with the same compression method as the backup itself.

View oplog slices

The oplog slices are stored in the `pbmPitr` subdirectory in the [remote storage defined in the config](#). A slice name reflects the start and end time this slice covers.

The `pbm list` output includes the following information:

- Backup snapshots. As of version 1.4.0, it also shows the completion time (renamed to the `restore_to_time` in version 2.0.0)
- Valid time ranges for recovery
- Point-in-time recovery status

```
$ pbm list

2021-08-04T13:00:58Z [restore_to_time: 2021-08-04T13:01:23Z]
2021-08-05T13:00:47Z [restore_to_time: 2021-08-05T13:01:11Z]
2021-08-06T08:02:44Z [restore_to_time: 2021-08-06T08:03:09Z]
2021-08-06T08:03:43Z [restore_to_time: 2021-08-06T08:04:08Z]
2021-08-06T08:18:17Z [restore_to_time: 2021-08-06T08:18:41Z]

PITR <off>:
2021-08-04T13:01:24 - 2021-08-05T13:00:11
2021-08-06T08:03:10 - 2021-08-06T08:18:29
2021-08-06T08:18:42 - 2021-08-06T08:33:09
```

3.6.3 Get expert help

If you need assistance, visit the [community forum](#) for comprehensive and free database knowledge, or contact our [Percona Database Experts](#) for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: September 25, 2023

Created: March 29, 2023

3.7 Snapshot-based physical backups

 **Version added: 2.2.0**

3.7.1 Considerations

1. This is a [technical preview feature](#).
2. Supported only for full physical backups
3. Available only if you run Percona Server for MongoDB in your environment as PBM uses the [\\$backupCursor](#) and [\\$backupCursorExtended](#) aggregation stages.

While a physical backup is a physical copy of your data directory, a snapshot is a point in time copy of your disk or a volume where the data files are stored. Restoring from snapshots is much faster and allows almost immediate access to data, while the database is unavailable during physical restore. Snapshot-based backups are especially useful for owners of large data sets with terabytes of data. Yet the snapshots don't guarantee data consistency in sharded clusters.

This is where Percona Backup for MongoDB steps in. It provides the interface to make snapshot-based physical backups and restores and ensures data consistency. As a result, database owners benefit from increased performance and reduced downtime, and are sure that their data remains consistent.

The snapshot-based physical backup / restore flow consists of three distinct stages:

- Preparing the database — done by PBM
- Copying files — done by the user / client app
- Completing the backup / restore — done by PBM.

This is the first stage of the snapshot-based backups where you can make them manually. Automated snapshot-based backups are planned for the future.

3.7.2 Make a backup

1. Refer to the [Before you start](#) section and make sure that you have made all the preparation steps for the backup.
2. To make a snapshot-based backup, run the `pbm backup` command with the type `external`:

```
$ pbm backup -t external
```

When executing the command, PBM does the following:

- opens the `$backupCursor`
- prepares the database for file copy
- stores the backup metadata on the storage and adds it to the files to copy
- prints the prompt similar to the following:

```
Ready to copy data from:
<node-list>
```

You also see the backup name.

3. (Optional) You can check the backup progress with the `pbm describe-backup`. The command output provides the backup state and what nodes are running backup.

- At this stage, you can copy the `dataDir` contents to the storage / make a snapshot using the technology of your choice.
- After the file copy, run the following command to close the `$backupCursor` and complete the backup:

```
$ pbm backup-finish <backup_name>
```

3.7.3 Restore a backup

Before you start:

- Shut down all `mongos` nodes. If you have set up the automatic restart of the database, disable it.
- Stop the arbiter nodes manually since there's no `pbm-agent` on these nodes to do that automatically.

Restore from a backup made through PBM

The following procedure describes the restore from backups [made through PBM](#). See [Restore from a backup made outside PBM](#) for how to restore from a backup made outside of PBM.

- To make a restore, run the following command:

```
$ pbm restore --external
```

Percona Backup for MongoDB stops the database, cleans up data directories on all nodes, provides the restore name and prompts you to copy the data:

```
Starting restore <restore_name> from
'[external]'.....
to copy data to the nodes data directory.
  After the copy is done, run: pbm restore-finish <restore_name> -c </path/to/
pbm.conf.yaml>
  Check restore status with: pbm describe-restore <restore_name> -c </path/to/
pbm.conf.yaml>
  No other pbm command is available while the restore is running!
```

- Copy the data. While a backup is made from a single node of a replica set, for the restore you must **copy the data on every node of a corresponding replica set in a cluster**. For example, copy files from a backup for a replica set `rs1` to all nodes in `rs1` in the target cluster and so on.
- After you copied the files to the nodes, complete the restore with the following command:

```
$ pbm restore-finish <restore_name> -c </path/to/pbm-conf.yaml>
```

At this stage, Percona Backup for MongoDB reads the metadata from the backup, prepares the data for the cluster / replica set start and ensures its consistency. The database is restored to the timestamp specified in the `restore_to_time` of the metadata.

- Optional. You can track the restore progress by running the `pbm describe-restore` command.

POST-RESTORE STEPS

After the restore is complete, do the following:

1. Start all `mongod` nodes
2. Start all `pbm-agents`
3. Run the following command to resync the backup list with the storage:

```
$ pbm config --force-resync
```

4. Start the balancer and start `mongos` nodes.
5. Make a fresh backup to serve as the new base for future restores.

Restore form a backup made outside PBM **Important**

For external backups made through PBM, PBM performs compatibility checks for the backup and the target cluster. If you restore the backup made outside PBM, it cannot ensure that the backup was made properly and in a consistent manner. Therefore, the backup compatibility is your responsibility.

To restore an external backup made outside PBM, you need to specify the following for the `pbm restore` command:

- a path to the configuration file of the `mongod` node on the source cluster from where the backup was made. This is the configuration file that PBM uses during the restore. It should contain the [storage options](#) per replica set name, for example:

```
rs1:
  storage:
    directoryPerDB: true
rs2:
  storage:
    directoryPerDB: true
```

To restore the data encrypted at rest, make sure data-at-rest encryption settings on the source and target clusters are the same.

- a timestamp to restore to

To restore from a backup, do the following:

1. Start a restore

```
$ pbm restore --external -c </path/to/mongod.conf> --ts
```

If the path to the source cluster `mongod.conf` is undefined, PBM tries to retrieve the required configuration options from the `mongod.conf` of the target cluster.

If the timestamp to restore to is undefined, PBM looks into the actual data during the restore and defines the most recent common cluster time across all shards. PBM restores the database up to this time.

2. Next, copy the data files. Note that you must copy the data **on every data-bearing node of your cluster / replica set**.

3. Complete the restore by running:

```
$ pbm restore-finish <restore_name> -c </path/to/pbm.conf.yaml>
```

At this stage, Percona Backup for MongoDB prepares the data for the cluster / replica set start and ensures its consistency.

4. Don't forget to complete the [post-restore steps](#)

3.7.4 Get expert help

If you need assistance, visit the [community forum](#) for comprehensive and free database knowledge, or contact our [Percona Database Experts](#) for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: July 13, 2023

Created: July 13, 2023

3.8 Comparison with MongoDB

Percona Backup for MongoDB is a fully supported community backup solution that can perform cluster-wide consistent backups in MongoDB. The following table compares Percona Backup for MongoDB with the MongoDB backup solutions:

Feature name	Percona Backup for MongoDB	MongoDB Community <code>mongodump</code>	MongoDB Enterprise	MongoDB Atlas
Open source backup	Yes	No	No	No
Binary database export (logical backup)	Yes	Yes	Yes	Yes
Built-in point-in-time recovery support	Yes	No	Yes	Yes
Physical backup	Yes	No	Yes	Yes
Incremental backup (physical)	Yes	No	Yes	Yes
Backup management interfaces	Percona Backup for MongoDB (CLI) PMM <code>mongodump</code> / <code>mongorestore</code> (CLI)	- - <code>mongodump</code> / <code>mongorestore</code> (CLI)	Ops Manager Cloud Manager <code>mongodump</code> / <code>mongorestore</code> (CLI)	Atlas backups <code>mongodump</code> / <code>mongorestore</code> (CLI)
Sharded cluster restores supported	Yes	No	Yes	Yes

3.8.1 What you get with Percona Backup for MongoDB

- [Enterprise features without extra costs](#)
- [Works for both sharded clusters and non-sharded replica sets](#)
- [Simple command-line management utility](#). For backup management via a user interface, consider [using PBM through Percona Monitoring and Management](#)
- [Simple, integrated-with-MongoDB authentication](#)
- [Distributed transaction consistency with MongoDB 4.2+](#)
- [Compatibility with different storage types: S3-compatible storage, Microsoft Azure Blob storage, filesystem storage type for locally-mounted remote filesystem backup servers](#)

3.8.2 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: March 29, 2023

Created: March 29, 2023

4. Discover PBM

4.1 How Percona Backup for MongoDB works

Even in a highly-available architecture, such as with MongoDB replication, backups are still required even though losing one server is not fatal. Whether for a complete or partial data disaster, you can use PBM (Percona Backup for MongoDB) to go back in time to the best available backup snapshot.

Percona Backup for MongoDB is a command line interface. It provides [the set of commands](#) to manage backup and restore operations in your database.

4.1.1 Usage example

Let's have a look at how Percona Backup for MongoDB works.

With [Percona Backup for MongoDB up and running](#) in your environment, make a backup:

```
$ pbm backup
```

To save all events that occurred to the data between the backups, enable saving oplog slices:

```
$ pbm config --set pitr.enabled=true
```

Now, imagine that your web application's update was released on February 7 at 03:00 UTC. By 15:23 UTC, someone realizes that this update has a bug that is wiping the historical data of any user who logged in. To remediate this negative impact on data, it's time to roll back up to the time of the application's update - up to February 7, 03:00 UTC.

```
$ pbm list
```

The output lists the valid time ranges for the restore. The desired time (February 7, 03:00 UTC) falls within the `2021-02-03T08:08:36Z-2021-02-09T12:20:23Z` range, so let's restore the database up to that time.

Since the restore and saving oplog slices are exclusive operations and cannot run together, let's stop the oplog slicing first:

```
$ pbm config --set pitr.enabled=false
```

Now, let's restore the database:

```
$ pbm restore --time 2021-02-07T02:59:59
```

To be on the safe side, it is a good practice to make a fresh backup after the restore is complete.

```
$ pbm backup
```

This backup refreshes the timeline and serves as the base for saving oplog slices. To re-enable this process, run:

```
$ pbm config --set pitr.enabled=true
```

4.1.2 Next steps

[Install and get started with Percona Backup for MongoDB](#)

4.1.3 Useful links

- [PBM architecture](#)
- [Backup types](#)

4.1.4 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: March 29, 2023

Created: August 18, 2022

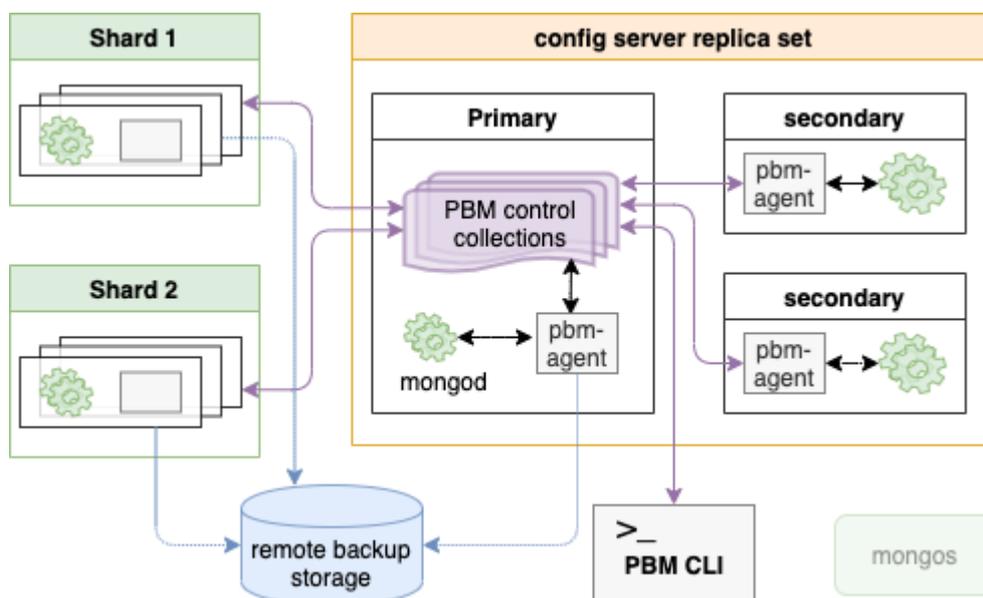
4.2 Architecture

4.2.1 Architecture

Percona Backup for MongoDB consists of the following components:

- `pbm-agent` is a process running on every `mongod` node within the cluster or within a replica set that performs backup and restore operations.
- `pbm CLI` is a command-line utility that instructs `pbm-agents` to perform an operation.
- **PBM Control collections** are special collections in MongoDB that store the configuration data and backup states. Both `pbm CLI` and `pbm-agent` use PBM Control collections to check backup status in MongoDB and communicate with each other.
- Remote backup storage is where Percona Backup for MongoDB saves backups. It can be either [an S3 compatible storage](#) or a filesystem-type storage.

The following diagram illustrates how Percona Backup for MongoDB components communicate with MongoDB.



Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

[Community Forum](#) [Get a Percona Expert](#)

Last update: March 29, 2023

Created: August 18, 2022

4.2.2 pbm-agent

A `pbm-agent` is a process that runs backup, restore, delete, and other operations available with Percona Backup for MongoDB.

A `pbm-agent` instance must run on each `mongod` instance. This includes replica set nodes that are currently secondaries and config server replica set nodes in a sharded cluster.

An operation is triggered when the `pbm` CLI makes an update to the `PBM Control` collection. All `pbm-agents` monitor changes to the PBM control collections, but only one `pbm-agent` in each replica set will be elected to execute an operation. The elections are done by a random choice among secondary nodes. If no secondary nodes respond, then the `pbm-agent` on the primary node is elected for an operation.

The elected `pbm-agent` acquires a lock for an operation. This prevents mutually exclusive operations like backup and restore to be executed simultaneously.

When the operation is complete, the `pbm-agent` releases the lock and updates the PBM control collections.

A single `pbm-agent` is involved with only one cluster (or non-sharded replica set). The `pbm` CLI utility can connect to any cluster to which it has network access, so it is possible for one user to list and launch backups or restores on many clusters.

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: March 29, 2023

Created: March 29, 2023

4.2.3 PBM Command Line Utility (pbm)

`pbm` CLI is the command line tool with which you operate Percona Backup for MongoDB. `pbm` provides the **pbm** command that you will use manually in the shell. It will also work as a command that can be executed in scripts (for example, by `crond`).

The set of `pbm` sub-commands enables you to manage backups in your MongoDB environment.

`pbm` uses [PBM Control collections](#) to communicate with `pbm-agent` processes. It starts and monitors backup or restore operations by updating and reading the corresponding PBM control collections for operations, log, etc. Likewise, it modifies the PBM config by saving it in the PBM Control collection for config values.

`pbm` CLI does not have its own config and/or cache files. Setting the `PBM_MONGODB_URI` environment variable in your shell is a configuration-like step that should be done for practical ease though. (Without `PBM_MONGODB_URI`, the `--mongodb-uri` command line argument will need to be specified each time.)

To learn how to set the `PBM_MONGODB_URI` environment variable, see [Set the MongoDB connection URI for pbm CLI](#). For more information about MongoDB URI connection strings, see [Authentication](#).

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: March 29, 2023

Created: March 29, 2023

4.2.4 PBM Control Collections

The config and state (current and historical) for backups is stored in collections in the MongoDB cluster or non-sharded replica set itself. These are put in the system `admin` db to keep them cleanly separated from user db namespaces.

In sharded clusters, this is the `admin` db of the config server replica set. In a non-sharded replica set, the PBM Control Collections are stored in `admin` db of the replica set itself.

- `admin.pbmBackups` - Log / status of each backup.
- `admin.pbmAgents` - Contains information about `pbm-agents` statuses and health.
- `admin.pbmConfig` - Contains configuration information for Percona Backup for MongoDB.
- `admin.pbmCmd` - Is used to define and trigger operations.
- `admin.pbmLock` - **pbm-agent** synchronization-lock structure.
- `admin.pbmLockOp` - Is used to coordinate operations that are not mutually exclusive such as make backup and delete backup.
- `admin.pbmLog` - Stores log information from all `pbm-agents` in the MongoDB environment. Available in Percona Backup for MongoDB as of version 1.4.0.
- `admin.pbmOpLog` - Stores [operation IDs](#).
- `admin.pbmPITRChunks` - Stores [Point-in-Time Recovery](#) oplog slices.
- `admin.pbmPITRState` - Contains current state of Point-in-Time Recovery incremental backups.
- `admin.pbmRestores` - Contains restore history and the restore state for all replica sets.
- `admin.pbmStatus` - Stores Percona Backup for MongoDB status records.

The `pbm` command line tool creates these collections as needed. You do not have to maintain these collections, but you should not drop them unnecessarily either. Dropping them during a backup will cause an abort of the backup.

Filling the config collection is a prerequisite to using Percona Backup for MongoDB for executing backups or restores. (See config page later.)

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: March 29, 2023

Created: March 29, 2023

4.2.5 Remote backup storage

Remote backup storage

Percona Backup for MongoDB saves your files to a directory. Using `pbm list`, a user can scan this directory to find existing backups even if they never used `pbm` on their computer before.

The files are prefixed with the (UTC) starting time of the backup. For each backup, there is one metadata file. For each replica set, a backup includes the following:

- A mongodump-format compressed archive that is the dump of collections
- A (compressed) BSON file dump of the oplog covering the time span of the backup

The end time of the oplog slice(s) is the data-consistent point in time of a backup snapshot.

SUPPORTED STORAGE TYPES

Percona Backup for MongoDB supports the following storage types:

- [S3-compatible storage](#)
- [Filesystem type storage](#)
- [Microsoft Azure Blob storage](#)

S3-compatible storage

Percona Backup for MongoDB should work with other S3-compatible storages, but was only tested with the following ones:

- [Amazon Simple Storage Service](#)
- [Google Cloud Storage](#)
- [MinIO](#)

Server-side encryption

 **Version added: 1.3.2**

Percona Backup for MongoDB supports [server-side encryption](#) for [S3 buckets](#) with customer-provided keys stored in AWS KMS (SSE-KMS).

To use the SSE-KMS encryption, specify the following parameters in the Percona Backup for MongoDB configuration file:

```
serverSideEncryption:
  kmsKeyID: <kms_key_ID>
  sseAlgorithm: aws:kms
```

 **Version added: 2.0.1**

Percona Backup for MongoDB also supports server-side encryption with customer-provided keys that are stored on the client side (SSE-C). Percona Backup for MongoDB provides the encryption keys as part of the requests to the S3 storage. The S3 storage uses them to encrypt/decrypt the data using the [AES-256](#) encryption algorithm. In such a way you save on subscribing to AWS KMS services and can use the server-side encryption with the S3-compatible storage of your choice.

SSE-C encryption should work with other S3-compatible storage types, but was only tested with the AWS and MinIO. Check the support of this functionality with your S3 storage provider.

Warning

1. Enable/disable the server-side encryption only for the empty bucket. Otherwise, Percona Backup for MongoDB fails to save/retrieve objects to/from the storage properly.
2. S3 storage doesn't manage or store the encryption key. It is your responsibility to track what key was used to encrypt what object in the bucket. If you lose the key, any request for an object without the encryption key fails and you lose the object.

To use the SSE-C encryption, specify the following parameters in the Percona Backup for MongoDB configuration file:

```
serverSideEncryption:
  sseCustomerAlgorithm: AES256
  sseCustomerKey: <your_encryption_key>
```

See also

AWS Documentation:

- [Protecting Data Using Server-Side Encryption with CMKs Stored in AWS Key Management Service \(SSE-KMS\)](#)
- [Protecting data using server-side encryption with customer-provided encryption keys \(SSE-C\)](#)

Debug logging

Version added: 1.7.0

You can enable debug logging for different types of S3 requests in Percona Backup for MongoDB. Percona Backup for MongoDB prints S3 log messages in the `pbm logs` output so that you can debug and diagnose S3 request issues or failures.

To enable S3 debug logging, set the `storage.s3.DebugLogLevel` option in Percona Backup for MongoDB configuration. The supported values are: `LogDebug`, `Signing`, `HTTPBody`, `RequestRetries`, `RequestErrors`, `EventStreamBody`.

Storage classes

Version added: 1.7.0

Percona Backup for MongoDB supports [Amazon S3 storage classes](#). Knowing your data access patterns, you can set the S3 storage class in Percona Backup for MongoDB configuration. When Percona Backup for MongoDB uploads data to S3, the data is distributed to the corresponding storage class. The support of S3 bucket storage types allows you to effectively manage S3 storage space and costs.

To set the storage class, specify the `storage.s3.storageClass` option in Percona Backup for MongoDB configuration file

```
storage:
  type: s3
```

```
s3:
  storageClass: INTELLIGENT_TIERING
```

When the option is undefined, the S3 Standard storage type is used.

Configure upload retries

 **Version added: 1.7.0**

You can set up the number of attempts for Percona Backup for MongoDB to upload data to S3 storage as well as the min and max time to wait for the next retry. Set the options `storage.s3.retryer.numMaxRetries`, `storage.s3.retryer.minRetryDelay` and `storage.s3.retryer.maxRetryDelay` in Percona Backup for MongoDB configuration.

```
retryer:
  numMaxRetries: 3
  minRetryDelay: 30
  maxRetryDelay: 5
```

This upload retry increases the chances of data upload completion in cases of unstable connection.

Data upload for storage with self-issued TLS certificates

 **Version added: 1.7.0**

Percona Backup for MongoDB supports data upload to S3-like storage that supports self-issued TLS certificates. To make this happen, disable the TLS verification of the S3 storage in Percona Backup for MongoDB configuration:

```
$ pbm config --set storage.s3.insecureSkipTLSVerify=True
```

Warning

Use this option with caution as it might leave a hole for man-in-the-middle attacks.

Remote filesystem server storage

This storage must be a remote file server mounted to a local directory. It is the responsibility of the server administrators to guarantee that the same remote directory is mounted at exactly the same local path on all servers in the MongoDB cluster or non-sharded replica set.

Warning

Percona Backup for MongoDB uses the directory as if it were any normal directory, and does not attempt to confirm it is mounted from a remote server.

If the path is accidentally a normal local directory, errors will eventually occur, most likely during a restore attempt. This will happen because **pbm-agent** processes of other nodes in the same replica set can't access backup archive files in a normal local directory on another server.

Local filesystem storage

This cannot be used except if you have a single-node replica set. (See the warning note above as to why). We recommend using any object store you might be already familiar with for testing. If you don't have an object store yet, we recommend using MinIO for testing as it has simple setup. If you plan to use a remote filesystem-type backup server, please see the [Remote Filesystem Server Storage](#) above.

Microsoft Azure Blob storage

 **Version added: 1.5.0**

You can use [Microsoft Azure Blob Storage](#) as the remote backup storage for Percona Backup for MongoDB.

This gives users a vendor choice. Companies with Microsoft-based infrastructure can set up Percona Backup for MongoDB with less administrative efforts.

Permissions setup

Regardless of the remote backup storage you use, grant the `List/Get/Put/Delete` permissions to this storage for the user identified by the access credentials.

The following example shows the permissions configuration to the `pbm-testing` bucket on the AWS S3 storage.

```
{
  "Version": "2021-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:ListBucket"
      ],
      "Resource": "arn:aws:s3:::pbm-testing"
    },
    {
      "Effect": "Allow",
      "Action": [
        "s3:PutObject",
        "s3:PutObjectAcl",
        "s3:GetObject",
        "s3:GetObjectAcl",
        "s3:DeleteObject"
      ],
      "Resource": "arn:aws:s3:::pbm-testing/*"
    }
  ]
}
```

Please refer to the documentation of your selected storage for the data access management.

 **See also**

- AWS documentation: [Controlling access to a bucket with user policies](#)
- Google Cloud Storage documentation: [Overview of access control](#)
- Microsoft Azure documentation: [Assign an Azure role for access to blob data](#)
- MinIO documentation: [Policy Management](#)

GET EXPERT HELP

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: October 6, 2023

Created: August 18, 2022

Remote backup storage configuration examples

Provide the remote backup storage configuration as a YAML config file. The following are the examples of config files for [supported remote storages](#). For how to insert the config file, see [Insert the whole Percona Backup for MongoDB config from a YAML file](#).

S3-COMPATIBLE REMOTE STORAGE

Amazon Simple Storage Service

```
storage:
  type: s3
  s3:
    region: us-west-2
    bucket: pbm-test-bucket
    prefix: data/pbm/backup
    credentials:
      access-key-id: <your-access-key-id-here>
      secret-access-key: <your-secret-key-here>
    serverSideEncryption:
      sseAlgorithm: aws:kms
      kmsKeyID: <your-kms-key-here>
```

GCS

```
storage:
  type: s3
  s3:
    region: us-east1
    bucket: pbm-testing
    prefix: pbm/test
    endpointUrl: https://storage.googleapis.com
    credentials:
      access-key-id: <your-access-key-id-here>
      secret-access-key: <your-secret-key-here>
```

MinIO

```
storage:
  type: s3
  s3:
    endpointUrl: "http://localhost:9000"
    region: my-region
    bucket: pbm-example
    prefix: data/pbm/test
    credentials:
      access-key-id: <your-access-key-id-here>
      secret-access-key: <your-secret-key-here>
```

REMOTE FILESYSTEM SERVER STORAGE

```
storage:
  type: filesystem
  filesystem:
    path: /data/local_backups
```

MICROSOFT AZURE BLOB STORAGE

```
storage:
  type: azure
```

```
azure:  
  account: <your-account>  
  container: <your-container>  
  prefix: pbm  
  credentials:  
    key: <your-access-key>
```

For the description of configuration options, see [Configuration file options](#).

GET EXPERT HELP

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: March 29, 2023

Created: August 18, 2022

4.3 Authentication

Percona Backup for MongoDB has no authentication and authorization subsystem of its own — it uses that of MongoDB. This means that `pbm` CLI and `pbm-agent` require only a valid MongoDB connection URI string for the `pbm` user.

For the S3-compatible remote storage authentication config, see [Percona Backup for MongoDB configuration in a cluster \(or non-sharded replica set\)](#).

4.3.1 MongoDB connection strings - A Reminder (or Primer)

Percona Backup for MongoDB uses [MongoDB Connection URI](#) strings to open MongoDB connections. Neither `pbm` CLI nor `pbm-agent` accept legacy-style command-line arguments for `--host`, `--port`, `--user`, `--password`, etc. as the `mongo` shell or `mongodump` command does.

The `pbm-agent` connection string The `pbm` CLI connection string

The `pbm-agent` processes should connect to their localhost `mongod` with a standalone type of connection.

```
pbm-agent --mongodb-uri "mongodb://pbmuser:secretpwd@localhost:27017/?authSource=admin"
```

Alternatively:

```
export PBM_MONGODB_URI="mongodb://pbmuser:secretpwd@localhost:27017/?authSource=admin"
pbm-agent
```

Replace the `pbmuser:secretpwd` with the credentials of [the user who owns the pbm process](#).

```
pbm list --mongodb-uri "mongodb://
pbmuser:secretpwd@mongocsvr1:27017,mongocsvr2:27017,mongocsvr3:27017/?
replicaSet=configrs&authSource=admin"
```

Alternatively:

```
export PBM_MONGODB_URI="mongodb://
pbmuser:secretpwd@mongocsvr1:27017,mongocsvr2:27017,mongocsvr3:27017/?
replicaSet=configrs&authSource=admin"
$ pbm list
```

Replace the `pbmuser:secretpwd` with the credentials of [the user who owns the pbm process](#)

The `pbm` CLI will ultimately connect to the replica set with PBM Control Collections.

- In a non-sharded replica set it is simply that replica set.
- In a cluster it is the config server replica set.

You do not necessarily have to provide that connection string. If you provide a connection to any live node (shard, configsvr, or non-sharded replica set member), `pbm` CLI will automatically determine the right hosts and establish a new connection to those instead.

The connection URI above is the format that MongoDB drivers have accepted universally since approximately the release time of MongoDB server v3.6. The `mongo` shell [has accepted it too since v4.0](#). Using a v4.0+ `mongo` shell is a recommended way to debug connection URI validity from the command line.

Since Percona Backup for MongoDB must authenticate in MongoDB, we recommend specifying the authentication database associated with the `pbm` user's credentials in the connection URI string using the `authSource` option.

The [MongoDB Connection URI](#) specification also allows specifying the authentication database via the `defaultauthdb` component. However, in this case, Percona Backup for MongoDB makes a backup of only this specified database.

If both `authSource` and `defaultauthdb` are unspecified, the authentication database defaults to the `admin` database.

The [MongoDB Connection URI](#) specification includes several non-default options you may need to use. For example, the TLS certificates/keys needed to connect to a cluster or non-sharded replica set with network encryption enabled are `tls=true` plus `tlsCAFile` and/or `tlsCertificateKeyFile` (see [tls options](#)).

4.3.2 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: March 29, 2023

Created: August 18, 2022

4.4 Supported MongoDB deployments

Percona Backup for MongoDB works with sharded clusters and replica sets. It doesn't work on standalone MongoDB instances. This is because Percona Backup for MongoDB requires [oplog](#) to guarantee backup consistency. Oplog is available on nodes with replication enabled.

For testing purposes, you can deploy Percona Backup for MongoDB on a single-node replica set. To convert a standalone server into a replica set, specify the `replication.replSetName` option in the configuration file and then initiate the replica set.

 **Version added: 2.1.0**

Physical restores are available for deployments with arbiter nodes. However, since these nodes don't have the `pbm-agent` installed, you must manually stop them before the restore.

 **See also**

MongoDB Documentation: [Convert a Standalone to a Replica Set](#)

4.4.1 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: April 18, 2023

Created: March 29, 2023

4.5 Supported MongoDB versions

Percona Backup for MongoDB is compatible with the following MongoDB versions:

PBM version	MongoDB Community / Enterprise	Percona Server for MongoDB
2.3.0	Logical backups - version 4.4 and higher with MongoDB Replication enabled	- Logical backups - version 4.4 and higher - Physical backups - version 4.4.6-8, 5.0 and higher with MongoDB Replication enabled and WiredTiger configured as the storage engine.
1.7.0	Logical backups - version 4.2 and higher with MongoDB Replication enabled	- Logical backups - version 4.2 and higher - Physical backups (tech preview) - version 4.2.15-16, 4.4.6-8, 5.0 and higher with MongoDB Replication enabled and WiredTiger configured as the storage engine.
1.6.1	Logical backups - version 3.6 and higher with MongoDB Replication enabled	Logical backups - version 3.6 and higher

4.5.1 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: September 25, 2023

Created: August 18, 2022

5. Get started

5.1 Install

5.1.1 Install Percona Backup for MongoDB

Supported platforms

Percona Backup for MongoDB is available for the most 64-bit Linux distributions. Find the list of supported platforms on the [Percona Software and Platform Lifecycle](#) page.

What nodes to install on

PBM-AGENT

Install `pbm-agent` on all servers that have `mongod` nodes in the MongoDB cluster (or non-sharded replica set). You don't need to start it on the `arbiter` node, since it doesn't have the data set.

PBM CLI

You can install `pbm CLI` on any or all servers or desktop computers you wish to use it from. Those computers must not be network-blocked from accessing the MongoDB cluster.

Installation tutorials

Install Percona Backup for MongoDB in one of the following ways:

- [from Percona repositories](#) using the package manager of your operating system. This is the recommended way
- [build from source code](#) if you want full control over the installation
- [download tarballs from Percona website](#)
- [run Percona Backup for MongoDB in a Docker container](#)
- [deploy Percona Backup for MongoDB on Kubernetes](#)

After the installation completes, you have the following tools:

Tool	Purpose
<code>pbm</code>	Command-line interface for controlling the backup system
<code>pbm-agent</code>	An agent for running backup/restore actions on a database host
<code>pbm-speed-test</code>	An interface for field-testing compression and backup upload speed
<code>pbm-agent-entrypoint</code>	An entry point application that allows starting <code>pbm-agent</code> and also restarts it in case of any faults

Next steps

[Install from Percona repositories](#)

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: March 29, 2023

Created: August 18, 2022

5.1.2 Install from Percona repositories

Use the package manager of your operating system to install Percona Backup for MongoDB:

- `apt` – for Debian and Ubuntu Linux
- `yum` – for Red Hat Enterprise Linux and compatible Linux derivatives

To install the software from Percona repositories means to subscribe to them. Percona provides the [percona-release](#) repository management tool. It automatically enables the required repository for you and enables you to install and update both Percona Backup for MongoDB packages and required dependencies smoothly.

Important

Run the following commands as root or via the `sudo` command

1. Install `percona-release`

Install `percona-release` tool. If you have installed it before, [update](#) it to the latest version.

2. Enable the repository

Starting with of version 1.3.0, Percona Backup for MongoDB packages are stored in the `pbm` repository.

```
$ sudo percona-release enable pbm release
```

3. Install Percona Backup for MongoDB

On Debian and Ubuntu On Red Hat Enterprise Linux and derivatives

1. Reload the local package database:

```
$ sudo apt update
```

2. Install Percona Backup for MongoDB:

```
$ sudo apt install percona-backup-mongodb
```

Install Percona Backup for MongoDB:

```
$ sudo yum install percona-backup-mongodb
```

Next steps

Set up PBM

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)

 [Get a Percona Expert](#)

Last update: March 29, 2023

Created: March 29, 2023

5.1.3 Build from source code

Prerequisites

Building the project requires:

- Go 1.19 or above
- make
- git
- `krb5-devel` for Red Hat Enterprise Linux / CentOS or `libkrb5-dev` for Debian / Ubuntu. This package is required for Kerberos authentication in Percona Server for MongoDB.

See also

[Install and set up Go tools](#)

Procedure

1. CLONE THE REPOSITORY

```
$ git clone https://github.com/percona/percona-backup-mongodb
```

2. GO TO THE PROJECT DIRECTORY AND BUILD IT

```
$ cd percona-backup-mongodb
$ make build
```

After **make** completes, you can find `pbm` and `pbm-agent` binaries in the `./bin` directory:

```
$ cd bin
$ ./pbm version
```

To verify if Percona Backup for MongoDB has been built correctly and is ready for use, run

```
$ pbm version
```

Output

```
Version: [pbm version number]
Platform: linux/amd64
GitCommit: [commit hash]
GitBranch: main
BuildTime: [time when this version was produced in UTC format]
GoVersion: [Go version number]
```

Tip

Instead of specifying the path to `pbm` binaries, you can add it to the `PATH` environment variable:

```
$ export PATH=/percona-backup-mongodb/bin:$PATH
```

Post-install steps

On Debian and Ubuntu On Red Hat Enterprise Linux and derivatives

1. Create the environment file:

```
$ touch /etc/default/pbm-agent
```

2. Create the `pbm-agent.service` `systemd` unit file.

```
$ sudo vim /lib/systemd/system/pbm-agent.service
```

3. In the `pbm-agent.service` file, specify the following:

```
[Unit]
Description=pbm-agent
After=time-sync.target network.target

[Service]
EnvironmentFile=-/etc/default/pbm-agent
Type=simple
User=mongod
Group=mongod
PermissionsStartOnly=true
ExecStart=/usr/bin/pbm-agent

[Install]
WantedBy=multi-user.target
```

 **Note**

Make sure that the `ExecStart` directory includes the Percona Backup for MongoDB binaries. Otherwise, copy them from the `./bin` directory of your installation path.

4. Make `systemd` aware of the new service:

```
$ sudo systemctl daemon-reload
```

1. Create the environment file:

```
$ touch /etc/sysconfig/pbm-agent
```

2. Create the `pbm-agent.service` `systemd` unit file.

```
$ sudo vim /usr/lib/systemd/system/pbm-agent.service
```

3. In the `pbm-agent.service` file, specify the following:

```
[Unit]
Description=pbm-agent
After=time-sync.target network.target

[Service]
EnvironmentFile=-/etc/default/pbm-agent
Type=simple
User=mongod
Group=mongod
PermissionsStartOnly=true
ExecStart=/usr/bin/pbm-agent

[Install]
WantedBy=multi-user.target
```

Next steps

[Set up PBM](#)

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: March 29, 2023

Created: March 29, 2023

5.1.4 Download Percona Backup for MongoDB from Percona website

You can download Percona Backup for MongoDB from [Percona website](#) and install it:

- From binary tarballs.
- Manually, from the installation packages using `dpkg` (Debian and Ubuntu) or `rpm` (Red Hat Enterprise Linux and CentOS). However, you must make sure that all dependencies are satisfied.

Install from binary tarball

Find the link to the binary tarballs under the **Generic Linux** menu item on [Percona website](#).

1. FETCH THE BINARY TARBALL

Replace the `<version>` with the required version.

```
$ wget https://downloads.percona.com/downloads/percona-backup-mongodb/percona-backup-mongodb-<version>/binary/tarball/percona-backup-mongodb-<version>-x86_64.tar.gz
```

2. EXTRACT THE TARBALL

```
$ tar -xf percona-backup-mongodb-<version>-x86_64.tar.gz
```

3. EXPORT THE LOCATION OF THE BINARIES TO THE `PATH` VARIABLE

For example, if you've extracted the tarball to your `home` directory, the command would be the following:

```
$ export PATH=~/.percona-backup-mongodb-<version>/:$PATH
```

After Percona Backup for MongoDB is successfully installed on your system, you have `pbm` and `pbm-agent` programs available. See [Initial setup](#) for guidelines how to set up Percona Backup for MongoDB.

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: March 29, 2023

Created: March 29, 2023

5.1.5 Run Percona Backup for MongoDB in a Docker container

Docker images of Percona Backup for MongoDB are hosted publicly on [Docker Hub](#).

For more information about using Docker, see the [Docker Docs](#).

Note

Make sure that you are using the latest version of Docker. The ones provided via apt and yum may be outdated and cause errors.

By default, Docker will pull the image from Docker Hub if it is not available locally.

Prerequisites

- You need to deploy MongoDB or Percona Server for MongoDB. See [what MongoDB deployments are supported](#).
- [Create the pbm user](#) in your deployment. You will need this user credentials to start Percona Backup for MongoDB container.

Start Percona Backup for MongoDB

Start Percona Backup for MongoDB container with the following command:

```
$ docker run --name <container-name> -e PBM_MONGODB_URI="mongodb://  
<PBM_USER>:<PBM_USER_PASSWORD>@<HOST>:<PORT>" -d percona/percona-backup-mongodb:<tag>-multi
```

Where:

- `container-name` is the name you want to assign to your container.
- `PBM_MONGODB_URI` is a [MongoDB Connection URI](#) string used to connect to MongoDB nodes. See the [Initial setup](#) how to create the PBM user.
- `tag-multi` is the tag specifying the version you need. For example, `2.3.1-multi`. The `multi` part of the tag serves to identify the architecture (`x86_64` or `ARM64`) and pull the respective image. See the [full list of tags](#).

Note that every MongoDB node (including replica set secondary members and config server replica set nodes) requires a separate instance of Percona Backup for MongoDB. Thus, a typical, 3-node MongoDB replica set requires three instances of Percona Backup for MongoDB.

Set up Percona Backup for MongoDB

Percona Backup for MongoDB requires the remote storage where to store data. Use the following commands to configure it:

1. Start a Bash session:

```
$ docker exec -it --name <container-name> bash
```

2. Create a YAML configuration file:

```
$ vi /tmp/pbm_config.yaml
```

3. Specify remote storage parameters in the config file. The following example is for S3-compatible backup storage. Check what [other storages are supported](#) and [examples of storage configurations](#):

```
storage:
  type: s3
  s3:
    region: <your-region-here>
    bucket: <your-bucket-here>
    credentials:
      access-key-id: <your-access-key-id-here>
      secret-access-key: <your-secret-key-here>
```

4. Upload the config file:

```
$ pbm config --file /tmp/pbm_config.yaml
```

The command output displays your uploaded configuration.

Run Percona Backup for MongoDB

Percona Backup for MongoDB command line utility (`pbm`) provides the set of commands to control backups: create, restore, cancel backups, etc.

For example, to start a backup, use the following command:

```
$ docker exec -it --name <container-name> pbm backup
```

where `<container-name>` is the name you assigned to the container and `pbm backup` is the command to start a backup.

In the same way you can run other `pbm` commands. Find the full list of available commands in [Percona Backup for MongoDB reference](#).

Next steps

[List backups](#)

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)

 [Get a Percona Expert](#)

Last update: December 4, 2023

Created: September 25, 2023

5.1.6 Deploy Percona Backup for MongoDB on Kubernetes

Deploy and run Percona Backup for MongoDB on Kubernetes with Percona Operator for MongoDB. See [Quickstart guides](#) for the installation instructions on your favorite Kubernetes flavor.

Next steps

[Set up and configure](#)

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: March 29, 2023

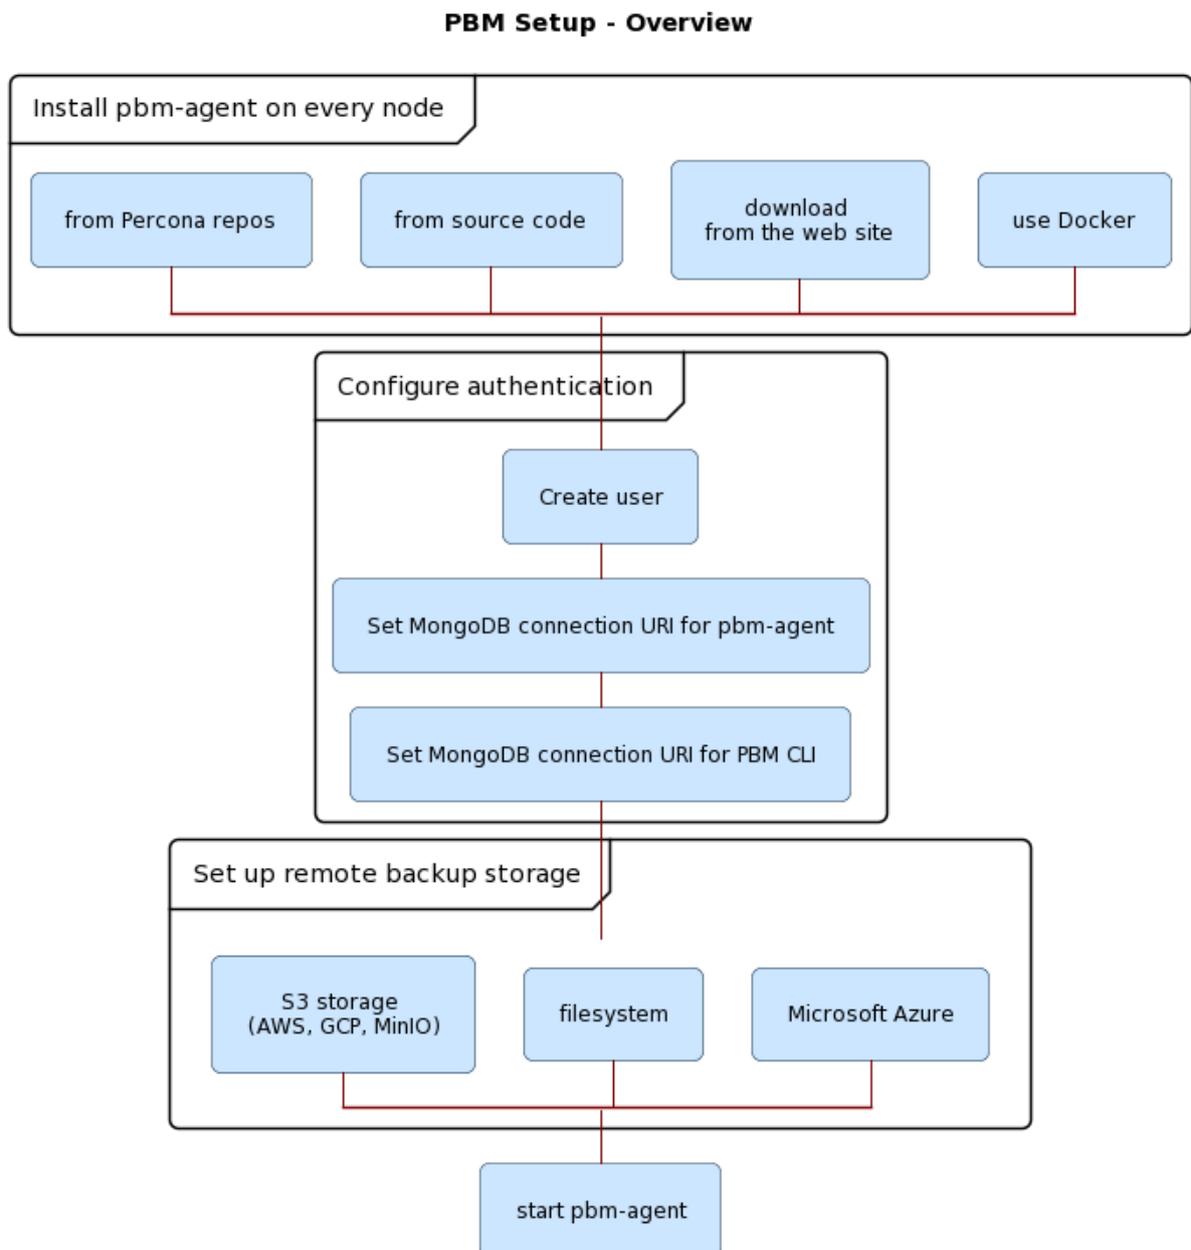
Created: March 29, 2023

5.2 Initial setup

After you install Percona Backup for MongoDB on every server with the `mongod` node that is not an arbiter node, the setup steps are the following:

1. Configure authentication in MongoDB.
2. Configure the remote backup storage.
3. Start `pbm-agent` process.

The following diagram outlines the installation and setup steps:



5.2.1 Configure authentication in MongoDB

Percona Backup for MongoDB uses the authentication and authorization subsystem of MongoDB. This means that to authenticate Percona Backup for MongoDB, you need to:

- Create a corresponding `pbm` user in the `admin` database
- Set a valid MongoDB connection URI string for `pbm-agent`
- Set a valid MongoDB connection URI string for `pbm CLI`

Create the `pbm` user

This step needs to be executed on a primary node of each replica set. In a sharded cluster, this means on every shard replica set and the config server replica set.

1. Create the role that allows any action on any resource.

```
db.getSiblingDB("admin").createRole({ "role": "pbmAnyAction",
  "privileges": [
    { "resource": { "anyResource": true },
      "actions": [ "anyAction" ]
    }
  ],
  "roles": []
});
```

2. Create the user and assign the role you created to it.

```
db.getSiblingDB("admin").createUser({user: "pbmuser",
  "pwd": "secretpwd",
  "roles" : [
    { "db" : "admin", "role" : "readWrite", "collection": "" },
    { "db" : "admin", "role" : "backup" },
    { "db" : "admin", "role" : "clusterMonitor" },
    { "db" : "admin", "role" : "restore" },
    { "db" : "admin", "role" : "pbmAnyAction" }
  ]
});
```

You can specify the `username` and `password` values and other options of the `createUser` command as you require so long as the roles shown above are granted.

Tip

To list all the host+port lists for the shard replica sets in a cluster, run the following command:

```
db.getSiblingDB("config").shards.find({}, {"host": true, "_id": false})
```

The replica set name at the *front* of these "host" strings will have to be placed as a `"/?replicaSet=xxxx"` argument in the parameters part of the connection URI (see below).

Set the MongoDB connection URI for `pbm-agent`

This step needs to be executed on each node where `pbm-agent` is installed.

A **`pbm-agent`** process connects to its localhost `mongod` node with a standalone type of connection.

To set the MongoDB URI connection string means to configure a service init script (`pbm-agent.service` systemd unit file) that runs a **`pbm-agent`**.

The `pbm-agent.service` systemd unit file includes the environment file. You set the MongoDB URI connection string for the `PBM_MONGODB_URI` variable within the environment file for every **`pbm-agent`**.

How to find the environment file

The path to the environment file is specified in the `pbm-agent.service` systemd unit file.

In Ubuntu and Debian, the `pbm-agent.service` systemd unit file is at the path `/lib/systemd/system/pbm-agent.service`.

In Red Hat and CentOS, the path to this file is `/usr/lib/systemd/system/pbm-agent.service`.

Example of `pbm-agent.service` systemd unit file

```
[Unit]
Description=pbm-agent
After=time-sync.target network.target

[Service]
EnvironmentFile=-/etc/default/pbm-agent
Type=simple
User=pbm
Group=pbm
PermissionsStartOnly=true
ExecStart=/usr/bin/pbm-agent

[Install]
WantedBy=multi-user.target
```

On Debian and Ubuntu Linux On Red Hat Enterprise Linux and derivatives

Edit the environment file `/etc/default/pbm-agent` and specify the MongoDB connection URI string for the `pbm` user to the local `mongod` node.

For example, if `mongod` node listens on port 27017, the MongoDB connection URI string will be the following:

```
PBM_MONGODB_URI="mongodb://pbmuser:secretpwd@localhost:27017/?authSource=admin"
```

Edit the environment file `/etc/sysconfig/pbm-agent` and specify the MongoDB connection URI string for the `pbm` user to the local `mongod` node.

For example, if `mongod` node listens on port 27017, the MongoDB connection URI string will be the following:

```
PBM_MONGODB_URI="mongodb://pbmuser:secretpwd@localhost:27017/?authSource=admin"
```

PASSWORDS WITH SPECIAL CHARACTERS

If the password includes special characters like #, @, / and so on, you must convert these characters using the [percent-encoding mechanism](#) when passing them to Percona Backup for MongoDB. For example, the password `secret#pwd` should be passed as follows in `PBM_MONGODB_URI`:

```
PBM_MONGODB_URI="mongodb://pbmuser:secret%23pwd@localhost:27017/?authSource=admin"
```

Set the MongoDB connection URI for `pbm` CLI

This step needs to be executed only on a host that you will use `pbm` CLI at.

Set the MongoDB URI connection string for `pbm` CLI in your shell. This allows you to call `pbm` commands without the `--mongodb-uri` flag.

Use the following command:

```
export PBM_MONGODB_URI="mongodb://pbmuser:secretpwd@localhost:27017/?authSource=admin&replSetName=xxxx"
```

For more information about what connection string to specify, refer to the [pbm connection string](#) section.

External authentication support in Percona Backup for MongoDB

In addition to SCRAM, Percona Backup for MongoDB supports other [authentication methods](#) that you use in MongoDB or Percona Server for MongoDB.

For external authentication, you create the `pbm` user in the format used by the authentication system and set the MongoDB connection URI string to include both the authentication method and authentication source.

For example, for [Kerberos authentication](#), create the `pbm` user in the `$external` database in the format `<username@KERBEROS_REALM>` (e.g. `pbm@PERCONATEST.COM`).

Specify the following string for MongoDB connection URI:

```
PBM_MONGODB_URI="mongodb://<username>%40<KERBEROS_REALM>@<hostname>:27018/?authMechanism=GSSAPI&authSource=%24external&replSetName=xxxx"
```

Note that you must first obtain the ticket for the `pbm` user with the `kinit` command before you start the **pbm-agent**:

```
$ sudo -u {USER} kinit pbm
```

Note that the `{USER}` is the user that you will run the `pbm-agent` process.

For [authentication and authorization via Native LDAP](#), you only create roles for LDAP groups in MongoDB as the users are stored and managed on the LDAP server. However, you still define the `$external` database as your authentication source:

```
PBM_MONGODB_URI="mongodb://<user>:<password>@<hostname>:27017/?authMechanism=PLAIN&authSource=%24external&replSetName=xxxx"
```

When using [AWS IAM authentication](#), create the `pbm` user in the `$external` database with the username that contains the ARN of the IAM user/role.

User authentication	Role authentication
<code>arn:aws:iam::<ARN>:user/<user_name></code>	
	<code>arn:aws:iam::<ARN>:role/<role_name></code>

The MongoDB connection URI string then looks like the following:

```
PBM_MONGODB_URI="mongodb://<aws_access_key_id>:<aws_secret_access_key>@<hostname>:27017/?  
authMechanism=MONGODB-AWS&authSource=%24external&replSetName=xxxx"
```

5.2.2 Configure remote backup storage

The easiest way to provide remote backup storage configuration is to specify it in a YAML config file and upload this file to Percona Backup for MongoDB using `pbm` CLI.

The storage configuration itself is out of scope of the present document. We assume that you have configured one of the supported remote backup storages.

1. Create a config file (e.g. `pbm_config.yaml`).
2. Specify the storage information within.

The following is the sample configuration for Amazon AWS:

```
storage:
  type: s3
  s3:
    region: us-west-2
    bucket: pbm-test-bucket
    prefix: data/pbm/backup
    credentials:
      access-key-id: <your-access-key-id-here>
      secret-access-key: <your-secret-key-here>
    serverSideEncryption:
      sseAlgorithm: aws:kms
      kmsKeyID: <your-kms-key-here>
```

Tip

If you are using AWS PrivateLink, the s3 endpoint needs to be specified explicitly. You can use the option `endpointUrl` for this scope, like in the following example:

```
...
s3:
  region: us-west-2
  bucket: pbm-test-bucket
  prefix: data/pbm/backup
  endpointUrl: https://your-endpoint-url-here
...
```

This is the sample configuration for Microsoft Azure Blob storage:

```
storage:
  type: azure
  azure:
    account: <your-account>
    container: <your-container>
    prefix: pbm
    credentials:
      key: <your-access-key>
```

This is the sample configuration for filesystem storage:

```
storage:
  type: filesystem
  filesystem:
    path: /data/local_backups
```

See more examples in [Configuration file examples](#).

3. Insert the config file

```
$ pbm config --file pbm_config.yaml
```

To learn more about Percona Backup for MongoDB configuration, see [Percona Backup for MongoDB configuration in a cluster \(or non-sharded replica set\)](#).

5.2.3 Start the pbm-agent process

Start `pbm-agent` on every server with the `mongod` node installed. It is best to use the packaged service scripts to run `pbm-agent`.

```
$ sudo systemctl start pbm-agent
$ sudo systemctl status pbm-agent
```

For example, imagine that you put `configsvr` nodes (listen port `27019`) collocated on the same servers as the first shard's `mongod` nodes (listen port `27018`, replica set name `sh1rs`). In this server there should be two `pbm-agent` processes, one connected to the shard (e.g. `"mongodb://username:password@localhost:27018/"`) and one to the `configsvr` node (e.g. `"mongodb://username:password@localhost:27019/"`).

For reference, the following is an example of starting `pbm-agent` manually. The output is redirected to a file and the process is backgrounded.

Important

Start the `pbm-agent` as the `mongod` user. The `pbm-agent` requires write access to the MongoDB data directory to make physical restores.

```
$ su mongod nohup pbm-agent --mongodb-uri "mongodb://username:password@localhost:27018/" > /data/mdb_node_xyz/pbm-agent.$(hostname -s).27018.log 2>&1 &
```

Replace `username` and `password` with those of your `pbm` user. `/data/mdb_node_xyz/` is the path where **pbm-agent** log files will be written. Make sure you have created this directory and granted write permissions to it for the `mongod` user.

Alternatively, you can run `pbm-agent` on a shell terminal temporarily if you want to observe and/or debug the startup from the log messages.

How to see the pbm-agent log

With the packaged `systemd` service, the log output to `stdout` is captured by `systemd`'s default redirection to `systemd-journald`. You can view it with the command below. See `man journalctl` for useful options such as `--lines`, `--follow`, etc.

```
$ journalctl -u pbm-agent.service
-- Logs begin at Tue 2019-10-22 09:31:34 JST. --
Jan 22 15:59:14 : Started pbm-agent.
Jan 22 15:59:14 pbm-agent[3579]: pbm agent is listening for the commands
...
...
```

If you started `pbm-agent` manually, see the file you redirected `stdout` and `stderr` to.

When a message `pbm agent is listening for the commands` is printed to the `pbm-agent` log file, `pbm-agent` confirms that it has connected to its `mongod` node successfully.

5.2.4 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: November 8, 2023

Created: August 18, 2022

5.3 Make a backup

5.3.1 Before you start

1. Install and set up Percona Backup for MongoDB
2. Check that `pbm agent` is running with the `pbm status` command

5.3.2 Make a backup

Logical Physical Selective Incremental Snapshot-based

To make a backup, run the following command:

```
$ pbm backup --type=logical
```

Logical backup is the default one so you can bypass the `--type` flag.

During *logical* backups, Percona Backup for MongoDB copies the actual data to the backup storage.

Starting with version 2.0.0, Percona Backup for MongoDB stores data in the new multi-file format where each collection has a separate file. The oplog is stored for all namespaces regardless whether this is a full or selective backup.

Multi-format is now the default data format since it allows [selective restore](#). Note, however, that you can make only full restores from backups made with earlier versions of Percona Backup for MongoDB.

 **Version added: 1.7.0**

```
$ pbm backup --type=physical
```

During a *physical* backup, Percona Backup for MongoDB stops [point-in-time recovery oplog slicing](#) if it's enabled, copies the contents of the `dbpath` directory (data and metadata files, indexes, journal and logs) from every shard and config server replica set to the backup storage.

 **Version added: 2.0.0**

Before you start, read about [selective backups known limitations](#).

To make a selective backup, run the `pbm backup` command and provide the value for the `--ns` flag in the format `<database.collection>`. The `--ns` flag value is case sensitive. For example, to back up the "Payments" collection, run the following command:

```
$ pbm backup --ns=staff.Payments
```

To back up the "Invoices" database and all collections that it includes, run the `pbm backup` command as follows:

```
$ pbm backup --ns=Invoices.*
```

During the backup process, Percona Backup for MongoDB stores data in the new multi-file format where each collection has a separate file. The oplog is stored for all namespaces regardless whether this is a full or selective backup.

Multi-format is now the default data format for both full and selective backups since it allows selective restore. Note, however, that you can make only full restores from backups made with earlier versions of Percona Backup for MongoDB.

 **Version added: 2.0.3**

Before you start, read more about [incremental backup](#).

To start incremental backups, first make a full incremental backup. It will serve as the base for subsequent incremental backups:

```
$ pbm backup --type incremental --base
```

The `pbm-agent` starts tracking the incremental backup history to be able to calculate and save the difference in data blocks. After that you can run regular incremental backups:

Compressed backups

By default, Percona Backup for MongoDB uses the `lz4` compression method when making a backup.

You can start a backup with a different compression method by passing the `--compression` flag to the **pbm backup** command.

For example, to start a backup with `gzip` compression, use the following command:

```
$ pbm backup --compression=gzip
```

Supported compression types are: `gzip`, `snappy`, `lz4`, `pgzip`, `zstd`. The `none` value means no compression is done during backup.

As of version 1.7.0, you can configure the compression level for backups. Specify the value for the `--compression-level` flag.

Default compression levels differ per compression method used.

The following table shows available compression levels per compression method:

Compression method	Supported compression levels	Default
<code>zstd</code>	1 - fastest speed, 2 - default, 3 - better compression, 4 - best compression	2
<code>snappy</code>	no levels	
<code>lz4</code>	From 1 (fastest) to 16	1
<code>gzip</code> and <code>pgzip</code>	-1 - default compression, 0 - no compression, 1 - best speed, 9 - best compression	-1

Note that the higher value you specify, the more time and computing resources it will take to compress the data.

5.3.3 Backups in sharded clusters

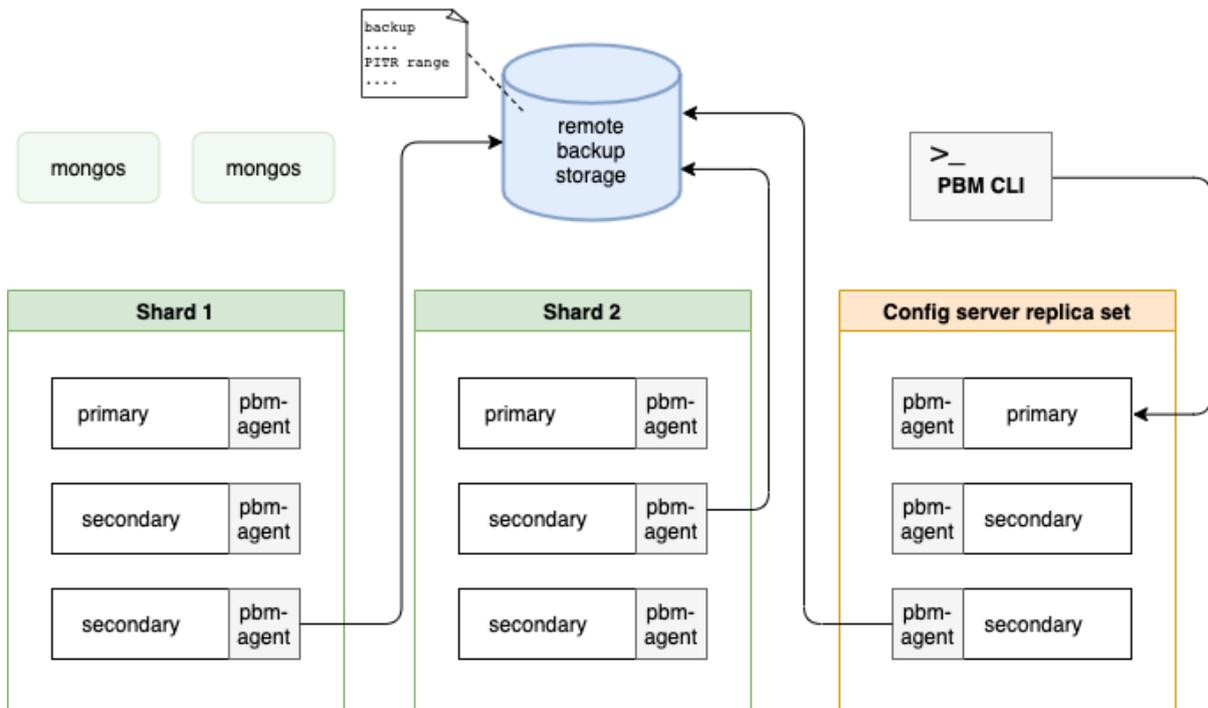
For PBM v1.0 (only)

Before running `pbm backup` on a cluster, stop the balancer.

In sharded clusters, one of the **pbm-agent** processes for every shard and the config server replica set writes backup snapshots into the remote backup storage directly. For logical backups, `pbm-agents` also write oplog slices. To learn more about oplog slicing, see Point-in-Time Recovery.

The `mongos` nodes are not involved in the backup process.

The following diagram illustrates the backup flow.



Important

If you reshard a collection in MongoDB 5.0 and higher versions, make a fresh backup to prevent data inconsistency and restore failure.

Adjust node priority for backups

In Percona Backup for MongoDB prior to version 1.5.0, the `pbm-agent` to do a backup is elected randomly among secondary nodes in a replica set. In sharded cluster deployments, the `pbm-agent` is elected among the secondary nodes in every shard and the config server replica sets. If no secondary node responds in a defined period, then the `pbm-agent` on the primary node is elected to do a backup.

As of version 1.5.0, you can influence the `pbm-agent` election by assigning a priority to `mongod` nodes in the Percona Backup for MongoDB [configuration file](#).

```
backup:
  priority:
    "localhost:28019": 2.5
    "localhost:27018": 2.5
    "localhost:27020": 2.0
    "localhost:27017": 0.1
```

The format of the priority array is `<hostname:port>:<priority>`.

To define priority in a sharded cluster, you can either list all nodes or specify priority for one node in each shard and config server replica set. The `hostname` and `port` uniquely identifies a node so that Percona Backup for MongoDB recognizes where it belongs to and grants the priority accordingly.

Note that if you listed only specific nodes, the remaining nodes will be automatically assigned priority `1.0`. For example, you assigned priority `2.5` to only one secondary node in every shard and config server replica set of the sharded cluster.

```

backup:
  priority:
    "localhost:27027": 2.5 # config server replica set
    "localhost:27018": 2.5 # shard 1
    "localhost:28018": 2.5 # shard 2

```

The remaining secondaries and the primary nodes in the cluster receive priority `1.0`.

The `mongod` node with the highest priority makes the backup. If this node is unavailable, the next priority node is selected. If there are several nodes with the same priority, one of them is randomly elected to make the backup.

If you haven't listed any nodes for the `priority` option in the config, the nodes have the default priority for making backups as follows:

- hidden nodes - priority 2.0
- secondary nodes - priority 1.0
- primary node - priority 0.5

Important

As soon as you adjust node priorities in the configuration file, it is assumed that you take manual control over them. The default rule to prefer secondary nodes over primary stops working.

This ability to adjust node priority helps you manage your backup strategy by selecting specific nodes or nodes from preferred data centers. In geographically distributed infrastructures, you can reduce network latency by making backups from nodes in geographically closest locations.

5.3.4 Next steps

- [List backups](#)
- [Make a restore](#)

5.3.5 Useful links

- [Backup and restore types](#)
- [Schedule backups](#)

5.3.6 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: July 13, 2023
Created: August 18, 2022

5.4 Restore a backup

To restore a backup, use the `pbm restore` command supplying the backup name from which you intend to restore. Percona Backup for MongoDB identifies the type of the backup (physical, logical or `incremental`) and restores the database up to the `restore_to_time` timestamp (available in `pbm list` output starting with version 1.4.0).

5.4.1 Considerations

Logical Physical Incremental

1. While the restore is running, prevent clients from accessing the database. The data will naturally be incomplete while the restore is in progress, and writes the clients make cause the final restored data to differ from the backed-up data.
2. If you enabled `Point-in-time recovery`, disable it before running `pbm restore`. This is because Point-in-Time Recovery oplog slicing and restore are incompatible operations and cannot be run together.
3. Backups made with Percona Backup for MongoDB prior to v1.5.0 are incompatible for restore with Percona Backup for MongoDB v1.5.0 and later. This is because processing of system collections `Users` and `Roles` has changed: in v1.5.0, `Users` and `Roles` are copied to temporary collection during backup and must be present in the backup during restore. In earlier versions of Percona Backup for MongoDB, `Users` and `Roles` are copied to a temporary collection during restore. Therefore, restoring from these backups with Percona Backup for MongoDB v1.5.0 isn't possible.

The recommended approach is to make a fresh backup after upgrading Percona Backup for MongoDB to version 1.5.0.

4. For versions earlier than 1.x, Percona Backup for MongoDB performs a full all-databases, all collections restore and does not offer an option to restore only a subset of collections in the backup, as MongoDB's `mongodump` tool does.
5. Starting with versions 1.x, Percona Backup for MongoDB replicates `mongodump`'s behavior to only drop collections in the backup. It does not drop collections that are created new after the time of the backup and before the restore. Run a `db.dropDatabase()` manually in all non-system databases (these are all databases except "local", "config" and "admin") before running `pbm restore` if you want to guarantee that the post-restore database only includes collections that are in the backup.

1. The Percona Server for MongoDB version for both backup and restore data must be within the same major release.
2. For PBM versions before 2.1.0, physical restores are not supported for deployments with arbiter nodes.
 1. The Percona Server for MongoDB version for both backup and restore data must be within the same major release.
2. Incremental backups made with PBM before PBM 2.1.0 are incompatible for restore with PBM 2.1.0 and onwards.
3. Physical restores are not supported for deployments with arbiter nodes.

5.4.2 Before you start

Logical Physical Selective Incremental

1. Stop the balancer.
2. Shut down all `mongos` nodes to stop clients from accessing the database while restore is in progress. This ensures that the final restored data doesn't differ from the backed-up data.
3. Disable point-in-time recovery if it is enabled. To learn more about point-in-time recovery, see [Point-in-time recovery](#).

1. Shut down all `mongos` nodes as the database won't be available while the restore is in progress.
2. Stop the arbiter nodes manually since there's no `pbm-agent` on these nodes to do that automatically.

You can restore a specific database or a collection either from a full or a selective backup. Read about [known limitations of selective restores](#).

Before you start, shut down all `mongos` nodes as the database won't be available while the restore is in progress.

5.4.3 Restore a database

Logical Physical Selective Incremental Snapshot-based

1. List the backups to restore from

```
$ pbm list
```

2. Restore from a desired backup. Replace the `<backup_name>` with the desired backup in the following command:

```
$ pbm restore <backup_name>
```

Note that you can restore a sharded backup only into a sharded environment. It can be your existing cluster or a new one. To learn how to restore a backup into a new environment, see [Restoring a backup into a new environment](#).

Post-restore steps

After a cluster's restore is complete, do the following:

1. Start the balancer and all `mongos` nodes to reload the sharding metadata.
2. We recommend to make a fresh backup to serve as the new base for future restores.

Adjust memory consumption

Starting with version 1.3.2, Percona Backup for MongoDB config includes the restore options to adjust the memory consumption by the `pbm-agent` in environments with tight memory bounds. This allows preventing out of memory errors during the restore operation.

```
restore:
  batchSize: 500
  numInsertionWorkers: 10
```

The default values were adjusted to fit the setups with the memory allocation of 1GB and less for the agent.

Note

The lower the values, the less memory is allocated for the restore. However, the performance decreases too.

Restore from a logical backup made on previous major version of Percona Server for MongoDB

In some cases you may need to restore from a backup made on previous major version of Percona Server for MongoDB. To make this happen, [Feature Compatibility Version \(FCV\)](#) values in both backup and the destination environment must match.

Starting with version 2.1.0, Percona Backup for MongoDB stores the FCV value in the backup metadata. If it doesn't match the FCV value on the destination environment, you see the warning in the `pbm status` output so that you can manually adjust it before the restore.

```
2023-04-10T10:48:54Z 302.80KB <logical> [ERROR: backup FCV "6.0" is incompatible with the
running mongo FCV "5.0"] [2023-04-10T10:49:14Z]
2023-04-10T08:40:10Z 172.25KB <logical> [ERROR: backup mongo version "6.0.5-4" is
incompatible with the running mongo version "5.0.15-13"] [2023-04-10T08:40:28Z]
```

The following example illustrates the restore from a backup made on Percona Server for MongoDB 4.4 on Percona Server for MongoDB 5.0.

1. Check the FCV value for the backup

```
$ pbm status
```

5.4.4 Restoring a backup into a new environment

To restore a backup from one environment to another, ensure the following:

1. Percona Backup for MongoDB configuration in the new environment must point to the same remote storage that is defined for the original environment, including the authentication credentials if it is an object store. Once you run `pbm list` and see the backups made from the original environment, then you can run the `pbm restore` command.
2. Don't run `pbm backup` from the new environment while Percona Backup for MongoDB configuration is pointing to the remote storage location of the original environment.

5.4.5 Restoring into a cluster / replica set with a different name

Starting with version 1.8.0, you can restore [logical backups](#) into a new environment that has the same or more number of shards and these shards have different replica set names. Starting with version 2.2.0, you can restore environments that have [custom shard names](#).

Starting with version 2.2.0, you can restore [physical](#) and [incremental physical](#) backups into a new environment with a different replica set names. Note that **the number of shards must be the same** as in the environment where the you made the backup.

To restore data to the environment with different replica set names, configure the name mapping between the source and target environments. You can either set the `PBM_REPLSET_REMAPPING` environment variable for `pbm` CLI or use the `--replset-remapping` flag for PBM commands. The mapping format is `<rsTarget>=<rsSource>`.

Important

Configure replica set name mapping for all shards in your cluster. Otherwise, Percona Backup for MongoDB attempts to restore the unspecified shard to the target shard with the same name. If there is no shard with such name or it is already mapped to another source shard, the restore fails.

Configure the replica set name mapping:

Using the environment variable for `pbm` CLI in your shell Using the command line

```
$ export PBM_REPLSET_REMAPPING="rsX=rsA,rsY=rsB"

$ pbm restore <timestamp> --replset-remapping="rsX=rsA,rsY=rsB"
```

The `--replset-remapping` flag is available for the following commands: `pbm restore`, `pbm list`, `pbm status`, `pbm oplog-replay`.

Note

Follow the post-restore steps on the new environment after the restore is complete.

This ability to restore data to clusters with different replica set names and the number of shards extends the set of environments compatible for the restore.

5.4.6 Next steps

[Point-in-time recovery](#)

5.4.7 Useful links

- [View restore progress](#)

5.4.8 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: July 13, 2023

Created: August 18, 2022

5.5 Make a point-in-time restore

5.5.1 Preconditions

Run `pbm status` or `pbm list` commands to check that the full backup snapshot exists and there are oplog slices.

5.5.2 Before you start

1. Disable point-in-time recovery. A restore and point-in-time recovery oplog slicing are incompatible operations and cannot be run simultaneously.

```
$ pbm config --set pitr.enabled=false
```

2. Stop the balancer and `mongos` nodes.
3. Make sure no writes are made to the database during restore.

5.5.3 Procedure

From logical backups From physical backups

Run `pbm restore` and specify the timestamp from the valid range:

```
$ pbm restore --time="2022-12-14T14:27:04"
```

The timestamp you specify for the restore must be within the time ranges in the PITR section of `pbm list` output. Percona Backup for MongoDB automatically selects the most recent backup among logical, physical and incremental in relation to the specified timestamp and uses that as the base for the restore.

To illustrate this behavior, let's use the following `pbm list` output as the example.

```
$ pbm list

2021-08-04T13:00:58Z [restore_to_time: 2021-08-04T13:01:23Z]
2021-08-05T13:00:47Z [restore_to_time: 2021-08-05T13:01:11Z]
2021-08-06T08:02:44Z [restore_to_time: 2021-08-06T08:03:09Z]
2021-08-06T08:03:43Z [restore_to_time: 2021-08-06T08:04:08Z]
2021-08-06T08:18:17Z [restore_to_time: 2021-08-06T08:18:41Z]

PITR <off>:
2021-08-04T13:01:24 - 2021-08-05T13:00:11
2021-08-06T08:03:10 - 2021-08-06T08:18:29
2021-08-06T08:18:42 - 2021-08-06T08:33:09
```

For timestamp `2021-08-06T08:10:10`, the backup snapshot `2021-08-06T08:02:44Z [restore_to_time: 2021-08-06T08:03:09]` is used as the base for the restore as it is the most recent one.

If you select a backup snapshot for the restore with the `--base-snapshot` option, the timestamp for the restore must also be later than the selected backup.

See also

[Restore a backup](#)

Post-restore steps

A restore operation changes the time line of oplog events. Therefore, all oplog slices made after the restore time stamp and before the last backup become invalid. After the restore is complete, do the following:

1. Make a new backup to serve as the starting point for oplog updates:

```
$ pbm backup
```

2. Re-enable point-in-time recovery to resume saving oplog slices:

```
$ pbm config --set pitr.enabled=true
```

Starting with version [2.2.0](#), you can recover your database from a full or an incremental physical backup in the same automated fashion as from a logical one. Percona Backup for MongoDB restores the backup snapshot and automatically replays the oplog events on top of it up to the specified time, guaranteeing data consistency. This helps you prevent data loss during a disaster and gives you the same user experience when managing backups and restores.

To restore a database from a physical backup, specify the time for the `pbm restore` command:

```
$ pbm restore --time <timestamp> -w
```

5.5.4 Select a backup snapshot for the restore

 **Version added: 1.6.0**

You can recover your database to the specific point in time using any backup snapshot, and not only the most recent one. Run the `pbm restore` command with the `--base-snapshot=<backup_name>` flag where you specify the desired backup snapshot.

To restore from any backup snapshot, Percona Backup for MongoDB requires continuous oplog. After the backup snapshot is made and point-in-time recovery is re-enabled, it copies the oplog saved with the backup snapshot and creates oplog slices from the end time of the latest slice to the new starting point thus making the oplog continuous.

5.5.5 Restore selected databases and collections

 **Important**

Supported only for replica sets. Available for logical backups.

1. Before you start:
 - a. Read [known limitations for selective backups and restores](#).
 - b. Check that you [have made a full backup](#) because it serves as the base for point-in-time recovery. Any selective backup is ignored.
2. To restore the desired database or a collection to a point in time, run the `pbm restore` command as follows:

```
$ pbm restore --base-snapshot <backup_name> --time <timestamp> \  
--ns <db.collection>
```

You can specify the selective backup as the base snapshot for the Point-in-time restore. In this case, Percona Backup for MongoDB restores only the namespace(s) included in this backup to the specified time.

Alternatively, you can use a full backup snapshot and restore the desired namespaces (databases or collections) up to the specific time from it. Specify them as the comma-separated list for the `pbm restore` command.

When point-in-time recovery is started, Percona Backup for MongoDB uses the provided base snapshot, restores the specified namespace(s) and replays oplog on top of it up to the specified time. If no base snapshot is provided, Percona Backup for MongoDB uses the most recent full backup snapshot.

5.5.6 Useful links

- [Restore a backup](#)
- [Replay oplog from arbitrary start time](#)

5.5.7 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: September 25, 2023

Created: March 29, 2023

6. Manage backups

6.1 List backups

Note

As of version 1.4.0, the `pbm list` command provides the information only about completed backups. To check for running backups, use the `pbm status`. For more information, see [Percona Backup for MongoDB status](#).

For Percona Backup for MongoDB version 1.3.4 and earlier, the `pbm list` command provides the running backup listed with an 'In progress' label. When that is absent, the backup is complete.

Use the `pbm list` command to view all completed backups.

```
$ pbm list
```

The output provides the following information:

- Backup name
- Backup type: [logical](#), [physical](#), [selective](#), [incremental](#). Available starting with version 1.7.0
- The time to which the sharded cluster / non-shared replica set will be returned to after the restore. Available starting with version 1.4.0.
- If [point-in-time recovery](#) is enabled, its status and the valid time ranges for the restore

Sample output

```
Backup snapshots:
2023-03-10T10:44:52Z <logical> [restore_to_time: 2023-03-10T10:44:56Z]
2023-03-10T10:49:20Z <physical> [restore_to_time: 2023-03-10T10:49:23Z]
2023-03-10T10:50:22Z <incremental> [restore_to_time: 2023-03-10T10:50:25Z]
2023-03-10T10:51:02Z <incremental> [restore_to_time: 2023-03-10T10:51:04Z]
2023-03-10T10:57:47Z <incremental> [restore_to_time: 2023-03-10T10:57:49Z]
2023-03-10T11:04:25Z <incremental> [restore_to_time: 2023-03-10T11:04:27Z]
2023-03-10T11:05:03Z <logical, selective> [restore_to_time: 2023-03-10T11:05:07Z]
```

6.1.1 Restore to time

In logical backups, the completion time almost coincides with the backup finish time. To define the completion time, Percona Backup for MongoDB waits for the backup snapshot to finish on all cluster nodes. Then it captures the oplog from the backup start time up to that time.

In physical backups, the completion time is only a few seconds after the backup start time. By holding the `$backupCursor` open guarantees that the checkpoint data won't change during the backup, and Percona Backup for MongoDB can define the completion time ahead.

6.1.2 Useful links

[View detailed information about a backup](#)

6.1.3 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: March 29, 2023

Created: August 18, 2022

6.2 View detailed information about a backup

To view a detailed information about a backup, run the following command:

```
$ pbm describe-backup <backup-name>
```

The output provides the backup name, type, status, size and the information about the cluster topology it was taken in. For [selective backups](#), it also shows the namespaces that were backed up.

Output

```
name: "2022-08-17T10:49:03Z"
type: logical
last_write_ts: 1662039300,2
last_transition_ts: "1662039304"
namespaces:
- Invoices.*
mongodb_version: 5.0.10-9
pbm_version: 2.0.0
status: done
size: 10234670
error: ""
replsets:
- name: rs1
  status: done
  iscs: false
  last_write_ts: 1662039300,2
  last_transition_ts: "1662039304"
  error: ""
```

Version added: 2.3.0

You can view the list of collections included in the *logical* or *selective* backup. This simplifies troubleshooting as it helps identify the backup contents for environments where databases are frequently created or dropped.

To view the backup contents, use the `--with-collections` flag:

```
$ pbm describe-backup <backup-name> --with-collections
```

 **Output**

```

name: "2023-09-14T14:44:33Z"
opid: 65031c51e6a16fa0e3deeb5f
type: logical
last_write_time: "2023-09-14T14:44:39Z"
last_transition_time: "2023-09-14T14:44:57Z"
mongodb_version: 6.0.9-7
fcv: "6.0"
pbm_version: 2.2.1
status: done
size_h: 89.3 KiB
replsets:
- name: rs0
  status: done
  node: rs00:30000
  last_write_time: "2023-09-14T14:44:38Z"
  last_transition_time: "2023-09-14T14:44:56Z"
  collections:
  - admin.pbmRRoles
  - admin.pbmRUsers
  - admin.system.roles
  - admin.system.users
  - admin.system.version
  - db0.c0
  - db0.c1
  - db1.c0
- name: rs1
  status: done
  node: rs10:30100
  last_write_time: "2023-09-14T14:44:38Z"
  last_transition_time: "2023-09-14T14:44:49Z"
  collections:
  - admin.pbmRRoles
  - admin.pbmRUsers
  - admin.system.roles
  - admin.system.users
  - admin.system.version
  - db0.c0
  - db1.c0
  - db1.c1
- name: cfg
  status: done
  node: cfg0:27000
  last_write_time: "2023-09-14T14:44:39Z"
  last_transition_time: "2023-09-14T14:44:42Z"
  configsvr: true
  collections:
  - admin.pbmAgents
  - admin.pbmBackups
  - admin.pbmCmd
  - admin.pbmConfig
  - admin.pbmLock
  - admin.pbmLockOp
  - admin.pbmLog
  - admin.pbmOpLog
  - admin.pbmPITRChunks
  - admin.pbmRRoles
  - admin.pbmRUsers
  - admin.system.roles
  - admin.system.users
  - admin.system.version
  - config.chunks
  - config.collections
  - config.databases
  - config.settings
  - config.shards

```

```
- config.tags  
- config.version
```

6.2.1 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: September 25, 2023

Created: March 29, 2023

6.3 Schedule backups

We recommend using `crond` or similar services to schedule backup snapshots.

Important

Before configuring `crond`, make sure that you have [installed](#) and [configured](#) Percona Backup for MongoDB to make backups in your database. Start a backup manually to verify this:

```
$ pbm backup
```

The recommended approach is to create a `crontab` file in the `/etc/cron.d` directory and specify the command in it. This simplifies server administration especially if multiple users have access to it.

`pbm` CLI requires a [valid MongoDB URI connection string](#) to authenticate in MongoDB. Instead of specifying the MongoDB URI connection string as a command line argument, which is a potential security risk, we recommend creating an environment file and specify the `export PBM_MONGODB_URI=$PBM_MONGODB_URI` statement within.

As an example, let's configure to run backup snapshots on 23:30 every Sunday. The steps are the following:

1. Create an environment file. Let's name it `pbm-cron`.

Debian and Ubuntu Red Hat Enterprise Linux and derivatives

```
$ vim /etc/default/pbm-cron
```

```
$ vim /etc/sysconfig/pbm-cron
```

2. Specify the environment variable in `pbm-cron`:

```
$ export PBM_MONGODB_URI="mongodb://pbmuser:secretpwd@localhost:27017/?replSetName=xxxx"
```

3. Grant access to the `pbm-cron` file for the user that will execute the `cron` task.
4. Create a `crontab` file. Let's name it `pbm-backup`.

```
$ touch pbm-backup
```

5. Specify the command in the file:

```
30 23 * * sun <user-to-execute-cron-task> . /etc/default/pbm-cron; /usr/bin/pbm backup
```

Note the dot `.` before the environment file. It sources (includes) the environment file for the rest of the shell commands.

6. Verify that backups are running in `/var/log/cron` or `/var/log/syslog` logs:

```
$ grep CRON /var/log/syslog
```

6.3.1 Schedule backups with point-in-time recovery running

It is convenient to automate making backups on a schedule using `crond` if you enabled [point-in-time recovery](#).

You can configure point-in-time recovery and `crond` in any order. Note, however, that point-in-time recovery will start running only after at least one full backup has been made.

- Make a fresh backup manually. It will serve as the starting point for incremental backups.
- Enable point-in-time recovery.
- Configure `crond` to run backup snapshots on a schedule.

When it is time for another backup snapshot, Percona Backup for MongoDB automatically disables point-in-time recovery and re-enables it once the backup is complete.

6.3.2 Backup storage cleanup

Previous backups are not automatically removed from the backup storage. You need to remove the oldest ones periodically to limit the amount of space used in the backup storage.

 **Version added: 2.1.0**

Starting with version 2.1.0, you can use the `pbm cleanup --older-than` command to delete outdated backup snapshots and point-in-time recovery oplog slices. You can configure a `cron` task to automate storage cleanup by specifying the following command in the `crontab` file:

```
$ $ /usr/bin/pbm cleanup -y --older-than 30d --wait
```

This command deletes backups and oplog slices that are older than 30 days. You can change the period by specifying a desired interval for the `--older-than` flag.

For PBM version 2.0.5 and earlier, use the `pbm delete backup --older-than <timestamp>` command. You can configure a `cron` task to automate backup deletion by specifying the following command in the `crontab` file:

```
$ /usr/bin/pbm delete-backup -f --older-than $(date -d '-1 month' +%Y-%m-%d)
```

This command deletes backups that are older than 30 days. You can change the period by specifying a desired interval for the `date` function.

6.3.3 Get expert help

If you need assistance, visit the [community forum](#) for comprehensive and free database knowledge, or contact our [Percona Database Experts](#) for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

6.4 Cancel a backup

You can cancel a running backup if, for example, you want to do another maintenance of a server and don't want to wait for the large backup to finish first.

To cancel the backup, use the `pbm cancel-backup` command.

```
$ pbm cancel-backup
Backup cancellation has started
```

After the command execution, the backup is marked as canceled in the `pbm status` output:

```
$ pbm status
```

Output:

```
2020-04-30T18:05:26Z Canceled at 2020-04-30T18:05:37Z
```

6.4.1 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: March 29, 2023

Created: August 18, 2022

6.5 Delete backups

Use `pbm delete-backup` to delete backup snapshots and `pbm delete-pitr` to delete point-in-time recovery oplog slices. Use the `pbm cleanup --older-than` command to [automate backup storage cleanup](#).

6.5.1 Delete outdated data

 **Version added: 2.1.0**

You can use the `pbm cleanup --older-than` command to delete both outdated backup snapshots and point-in-time recovery oplog slices. This simplifies the [automation of the backup rotation](#).

The timestamp you specify for the `--older-than` flag must be in the following format:

- `%Y-%M-%DT%H:%M:%S` (for example, `2023-04-20T13:13:20`) or
- `%Y-%M-%D` (`2023-04-20`)
- `XXd` (1d or 30d). Only days are supported.

During the cleanup, you see the backups and oplog slices to be deleted and are asked to confirm the action. To bypass it, add the `--yes` flag:

```
$ pbm cleanup --older-than=`%Y-%M-%D` --yes
```

Behavior

The timestamp you specify is considered to be the time to which you would wish to restore. Therefore, PBM doesn't delete all backup snapshots and oplog slices that could be used to restore to this time.

Here's how the cleanup works:

- **Physical and selective backups** are deleted up to the specified time.
- **Incremental physical backups** are deleted up to the specified time if the timestamp doesn't fall within the backup chain. If it does, PBM checks for the most recent base incremental backup in relation to the specified timestamp. PBM keeps this backup and the whole chain deriving from it to ensure the potential restore.

For example, you have the following list of backups:

```

Snapshots:
  2023-04-14T19:34:52Z 520.86MB <incremental> [restore_to_time: 2023-04-14T19:34:54Z]
  2023-04-14T08:12:50Z 576.63MB <incremental, base> [restore_to_time:
2023-04-14T08:12:52Z]
  2023-04-12T03:02:08Z 498.50MB <incremental> [restore_to_time: 2023-04-12T03:02:10Z]
  2023-04-11T19:30:14Z 552.77MB <incremental, base> [restore_to_time:
2023-04-11T19:30:16Z]
  2023-04-11T14:25:51Z 572.41MB <physical> [restore_to_time: 2023-04-11T14:25:54Z]

```

You wish to delete all backups that are older than 2023-04-14T15:00:00

```
$ pbm cleanup --older-than="2023-04-14T15:00:00"
```

This timestamp falls inside the backup chain that starts with the 2023-04-14T08:12:50Z backup. That's why PBM keeps this backup and the incremental backup chain deriving from it and deletes all data that is older than this backup.

Output:

```

S3 us-east-1 s3://http://192.168.56.1:9000/bcp/pbme2etest
Snapshots:
  2023-04-14T19:34:52Z 520.86MB <incremental> [restore_to_time: 2023-04-14T19:34:54Z]
  2023-04-14T08:12:50Z 576.63MB <incremental, base> [restore_to_time:
2023-04-14T08:12:52Z]

```

- **Logical backup** cleanup also depends on the point-in-time recovery settings.
 - By default, PBM looks for the most recent backup in relation to the specified timestamp and deletes all logical backups and oplog slices up to the backup's `restore_to_time` value.

To illustrate, let's say you have the following backup list:

```

Snapshots:
  2023-04-13T13:26:58Z 147.29MB <logical> [restore_to_time: 2023-04-13T13:27:15Z]
  2023-04-13T10:12:08Z 147.29MB <logical> [restore_to_time: 2023-04-13T10:12:27Z]
  2023-04-13T08:48:32Z 147.28MB <logical> [restore_to_time: 2023-04-13T08:48:51Z]
PITR chunks [2.11MB]:
  2023-04-13T08:48:52Z - 2023-04-13T13:27:15Z

```

You wish to delete all data up to 2023-04-13T12:00:00.

The most recent backup in relation to this timestamp is 2023-04-13T10:12:08Z 147.29MB. So PBM deletes all backups that are older than this backup. It also deletes all oplog slices up to the backup's `restore_to_time: 2023-04-13T10:12:27Z`. The output after the cleanup looks like this:

```

Snapshots:
  2023-04-13T13:26:58Z 147.29MB <logical> [restore_to_time: 2023-04-13T13:27:15Z]
  2023-04-13T10:12:08Z 147.29MB <logical> [restore_to_time: 2023-04-13T10:12:27Z]

```

```
PITR chunks [157.94KB]:
2023-04-13T10:12:28Z - 2023-04-13T13:27:46Z
```

- When point-in-time recovery is enabled and you specify the timestamp greater than the `restore_to_time` for the most recent logical backup, PBM keeps this backup and all oplog slices deriving from it to ensure point-in-time recovery.
- When the specified timestamp equals to the `restore_to_time` value for any full logical, physical and base incremental backups, PBM deletes all logical backup snapshots and oplog slices up to this backup's `restore_to_time`.

6.5.2 Delete backup snapshots

Considerations

1. You can only delete a backup that is not running (has the “done” or the “error” state). To check the backup state, run the `pbm status` command.
2. To ensure oplog continuity for [point-in-time restore](#), the `pbm delete-backup` command deletes any backup(s) except the following:
 - A backup snapshot that can serve as the base for any point-in-time recovery and has point-in-time recovery time ranges deriving from it. To delete such a backup, first [delete the oplog slices](#) that are created after the `restore-to time` value for this backup.
 - The most recent backup if point-in-time recovery is enabled and there are no oplog slices following this backup yet.

To illustrate this, let's take the following `pbm list` output:

```
Backup snapshots:
2022-10-05T14:13:50Z <logical> [restore_to_time: 2022-10-05T14:13:55Z]
2022-10-06T14:52:42Z <logical> [restore_to_time: 2022-10-06T14:52:47Z]
2022-10-07T14:57:17Z <logical> [restore_to_time: 2022-10-07T14:57:22Z]

PITR <on>:
2022-10-05T14:13:56Z - 2022-10-05T18:52:21Z
```

You can delete a backup `2022-10-06T14:52:42Z` since it has no point-in-time oplog slices. You cannot delete the following backups:

- `2022-10-05T14:13:50Z` because it is the base for recovery to any point in time from the PITR time range `2022-10-05T14:13:56Z - 2022-10-05T18:52:21Z`
- `2022-10-07T14:57:17Z` because PITR is enabled and there are no oplog slices following it yet.

Behavior

You can delete either a specified backup snapshot or all backup snapshots older than the specified time. Starting with version 2.0.0, you can also delete [selective backups](#).

A specific backup Backups older than the specified time

To delete a backup, specify the `<backup_name>` as an argument.

```
$ pbm delete-backup <backup_name>
```

To delete backups that were created before the specified time, pass the `--older-than` flag to the `pbm delete-backup` command. Specify the timestamp as an argument for `pbm delete-backup` in the following format:

- `%Y-%M-%DT%H:%M:%S` (for example, `2021-04-20T13:13:20Z`) or
- `%Y-%M-%D` (`2021-04-20`).

EXAMPLE

View backups:

```
$ pbm list
```

Output:

```
Backup snapshots:
 2021-04-20T20:55:42Z
 2021-04-20T23:47:34Z
 2021-04-20T23:53:20Z
 2021-04-21T02:16:33Z
```

Delete backups created before the specified timestamp

```
pbm delete-backup -f --older-than 2021-04-21
```

Output:

```
Backup snapshots:
 2021-04-21T02:16:33Z
```

By default, the `pbm delete-backup` command asks for your confirmation to proceed with the deletion. To bypass it, add the `-f` or `--force` flag.

```
$ pbm delete-backup --force 2021-04-20T13:45:59Z
```

For Percona Backup for MongoDB 1.5.0 and earlier versions, when you delete a backup, all oplog slices that relate to this backup are deleted too. For example, you delete a backup snapshot `2020-07-24T18:13:09` while there is another snapshot `2020-08-05T04:27:55` created after it. The **pbm-agent** deletes only oplog slices that relate to `2020-07-24T18:13:09`.

The same applies if you delete backups older than the specified time.

Note that when point-in-time recovery is enabled, the most recent backup snapshot and oplog slices that relate to it are not deleted.

6.5.3 Delete oplog slices

 **Version added: 1.6.0**

You can delete oplog slices saved before the specified time or all slices altogether. By deleting old and/or unnecessary slices, you can save storage space.

Behavior

To view oplog slices, run the `pbm list` command. If you have [deleted the snapshot](#) and want to delete the respective oplog slices, run the `pbm list --unbacked` command to view them.

Delete all oplog slices Earlier than the specified timestamp

Run the `pbm delete-pitr` and pass the `--all` flag:

```
$ pbm delete-pitr --all
```

To delete slices that are made earlier than the specified time, run the `pbm delete-pitr` command with the `--older-than` flag and pass the timestamp for it. The timestamp must be in the following format:

- `%Y-%M-%DT%H:%M:%S` (for example, `2021-07-20T10:01:18`) or
- `%Y-%M-%D` (`2021-07-20`).

```
$ pbm delete-pitr --older-than 2021-07-20T10:01:18
```

To enable [point-in-time recovery](#) from the most recent backup snapshot, Percona Backup for MongoDB does not delete slices that were made after that snapshot. For example, if the most recent snapshot is `2021-07-20T07:05:23Z` [`restore_to_time: 2021-07-21T07:05:44`] and you specify the timestamp `2021-07-20T07:05:44`, Percona Backup for MongoDB deletes only slices that were made before `2021-07-20T07:05:23Z`.

6.5.4 Get expert help

If you need assistance, visit the [community forum](#) for comprehensive and free database knowledge, or contact our [Percona Database Experts](#) for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: April 18, 2023

Created: August 18, 2022

6.6 View restore progress

 **Version added: 2.0.0**

You can track the status of both physical and logical restores. This gives you a clear understanding of the restore progress so that you can react accordingly.

To view the restore status, run the `pbm describe-restore` command and specify the restore name. To track the progress of a physical restore, also specify the path to the Percona Backup for MongoDB configuration file. Since `mongod` nodes are shut down during a physical restore, Percona Backup for MongoDB uses the configuration file to read the restore status on storage.

```
$ pbm describe-restore 2022-08-15T11:14:55.683148162Z -c pbm_config.yaml
```

The output provides the following information:

- Restore name
- The name of the backup from which the database was restored
- Type
- Status
- opID
- The time of the restore start
- Last transition time – the time when the restore process changed its status
- The name of every replica set, its restore status and the last transition time

For physical backups only, the following additional information is provided:

- The node name
- Restore status on the node
- Last transition time

For version 1.8.1 and earlier, tracking restore progress during physical restores is not available. To check the restore status, the options are:

- Check the `stderr` logs of the leader `pbm-agent`. The leader ID is printed once the restore has started.
- Check the status in the metadata file created on the remote storage for the restore. This file is in the root of the storage path and has the format `.pbm.restore/<restore_timestamp>.json`.

6.6.1 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: March 29, 2023

Created: September 21, 2022

6.7 Replay oplog from arbitrary start time

 **Version added: 1.7.0**

You can replay the [oplog](#) for a specific period on top of any backup: logical, physical, storage level snapshot (like [EBS-snapshot](#)). Starting with version [1.8.0](#), you can save oplog slices without the mandatory base backup snapshot. This behavior is controlled by the `pitr.oplogOnly` configuration parameter:

```
pitr:
  oplogOnly: true
```

By replaying these oplog slices on top of the backup snapshot with the `pbm oplog-replay` command, you can manually restore sharded clusters and non-sharded replica sets to a specific point in time from a backup made by any tool and not only by Percona Backup for MongoDB. Plus, you reduce time, storage space, and administration efforts on making the redundant base backup snapshot.

Warning

Use the oplog replay functionality with caution, only when you are sure about the starting time from which to replay oplog. The oplog replay does not guarantee data consistency when restoring from any backup. However, it is less error-prone for backups made with Percona Backup for MongoDB.

6.7.1 Oplog replay for physical backups

Starting with version [2.2.0](#), oplog replay on top of a physical backups made with Percona Backup for MongoDB is done automatically as part of [point-in-time recovery](#).

This section describes how to manually replay oplog on top of physical backups with Percona Backup for MongoDB version [2.1.0](#) and earlier.

After you [restored a physical backup](#), do the following:

1. Stop point-in-time recovery, if enabled, to release the lock.
2. Run `pbm status` or `pbm list` commands to find oplog chunks available for replay.
3. Run the `pbm oplog-replay` command and specify the `--start` and `--end` flags with the timestamps.

```
$ pbm oplog-replay --start="2022-01-02T15:00:00" --end="2022-01-03T15:00:00"
```

4. After the oplog replay, make a fresh backup and enable the point-in-time recovery oplog slicing.

6.7.2 Oplog replay for storage level snapshots

When making a backup, Percona Backup for MongoDB stops the point-in-time recovery. This is done to maintain data consistency after the restore.

Storage-level snapshots are saved with point-in-time recovery enabled. Thus, after the database restore from such a backup, point-in-time recovery is automatically enabled and starts oplog slicing. These new

oplog slices might conflict with the existing oplogs saved during the backup. To replay the oplog in such a case, do the following after the restore:

1. Disable point-in-time recovery.
2. Delete the oplog slices that might have been created.
3. Re-sync the data from the storage.
4. Run the `pbm oplog-replay` command and specify the `--start` and `--end` flags with the timestamps.

```
$ pbm oplog-replay --start="2022-01-02T15:00:00" --end="2022-01-03T15:00:00"
```

5. After the oplog replay, make a fresh backup and enable the point-in-time recovery oplog slicing.

Known limitations

The oplog replay fails if you rename the entire database or a collection.

6.7.3 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: July 13, 2023

Created: August 18, 2022

6.8 View backup logs

 **Version added: 1.4.0**

You can see the logs from all `pbm-agents` in your MongoDB environment using `pbm CLI`. This reduces time for finding required information when troubleshooting issues.

 **Note**

The log information about restores from physical backups is not available in `pbm` logs.

To view `pbm-agent` logs, run the `pbm logs` command and pass one or several flags to narrow down the search.

The following flags are available:

- `-t, --tail` - Show the last N rows of the log
- `-e, --event` - Filter logs by all backups or a specific backup
- `-n, --node` - Filter logs by a specific node or a replica set
- `-s, --severity` - Filter logs by severity level. The following values are supported (from low to high):
 - `D` - Debug
 - `I` - Info
 - `W` - Warning
 - `E` - Error
 - `F` - Fatal
- `-o, --output` - Show log information as text (default) or in JSON format.
- `-i, --opid` - Filter logs by the operation ID

6.8.1 Examples

The following are some examples of filtering logs:

Show logs for all backups

```
$ pbm logs --event=backup
```

Show the last 100 lines of the log about a specific backup 2020-10-15T17:42:54Z

```
$ pbm logs --tail=100 --event=backup/2020-10-15T17:42:54Z
```

Include only errors from the specific replica set

```
$ pbm logs -n rs1 -s E
```

The output includes log messages of the specified severity type and all higher levels. Thus, when `ERROR` is specified, both `ERROR` and `FATAL` messages are shown in the output.

6.8.2 Implementation details

`pbm-agents` write log information into the `pbmLog` collection in the [PBM Control collections](#). Every `pbm-agent` also writes log information to `stderr` so that you can retrieve it when there is no healthy `mongod` node in your cluster or replica set. For how to view an individual `pbm-agent` log, see [How to see the `pbm-agent` log](#).

Note

Log information from `pbmLog` collection is shown in the UTC timezone and from the `stderr` - in the server's time zone.

6.8.3 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: March 29, 2023

Created: August 18, 2022

7. Administer PBM

7.1 Percona Backup for MongoDB configuration via pipelines

 **Version added: 2.0.1**

To apply or update the configuration, Percona Backup for MongoDB reads the configuration file on the filesystem. When you run PBM remotely (in a cloud as Docker containers or pods in Kubernetes), you must upload the configuration file to the remote host's filesystem every time you update it.

To simplify the configuration management, you can manage the configuration file locally and use the UNIX pipeline to pass the file's contents to Percona Backup for MongoDB on a remote host/running in a container.

Here's how to do it:

1. Create/update the configuration file (for example, `/etc/pbm_config.yaml`)
2. Create an environment variable for the path to the configuration file

```
$ export CONFIG_PATH="/etc/pbm_config.yaml"
```

3. Pass the configuration file contents to Percona Backup for MongoDB. For example, if you run Percona Backup for MongoDB in Docker, use one of the following commands:

- Connect to the existing container and pass the configuration:

```
$ cat "$CONFIG_PATH" | docker compose exec -T $SERVICE_NAME pbm config --file="-"
```

Replace the `$SERVICE_NAME` with your [service name](#).

- Create a new container to pass the configuration and exit:

```
$ cat "$CONFIG_PATH" | docker run -i --env PBM_MONGODB_URI="mongodb://  
<PBM_USER>:<PBM_USER_PASSWORD>@<HOST>:<PORT>" --network=$NET_ID $CONTAINER_ID pbm  
config --file="-"
```

Specify the valid `PBM_MONGODB_URI` connection string, the ID of the network the container will connect to and the container ID.

As a result, your DBAs spend less time on administering Percona Backup for MongoDB and can focus on other activities instead.

7.1.1 Get expert help

If you need assistance, visit the [community forum](#) for comprehensive and free database knowledge, or contact our [Percona Database Experts](#) for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: March 29, 2023

Created: October 12, 2022

7.2 Automate access to S3 buckets for Percona Backup for MongoDB

When you run MongoDB and Percona Backup for MongoDB using AWS resources (on EC2 instances or using EKS), you can automate access to AWS S3 buckets for Percona Backup for MongoDB. Percona Backup for MongoDB uses the AWS environment variables and metadata to access S3 buckets so that you don't have to explicitly specify the S3 credentials in the PBM configuration file. Thereby you control the access to your cloud infrastructure from a single place.

7.2.1 IAM instance profile

 **Version added: 1.6.0**

IAM (Identity Access Management) is the AWS service that allows you to securely control access to AWS resources.

Using the IAM instance profile, you can automate access to S3 buckets for Percona Backup for MongoDB running on EC2 instance. The steps are the following:

1. Create the [IAM instance profile](#) and the permission policy within where you specify the access level that grants the access to S3 buckets.
2. Attach the IAM profile to an EC2 instance.
3. Configure an S3 storage bucket and verify the connection from the EC2 instance to it.
4. Provide the [remote storage information for PBM in a config file](#). Leave the `s3.credentials` array empty

```
storage:
  type: s3
  s3:
    region: <your-S3-region>
    bucket: <bucket-name>
```

 **Note**

If you specify S3 credentials, they override the EC2 instance environment variables and metadata, and are used for authentication instead.

5. Start the `pbm-agent` process

 **See also**

AWS documentation: [How can I grant my Amazon EC2 instance access to an Amazon S3 bucket?](#)

7.2.2 IAM Roles for Service Accounts (IRSA)

 **Version added: 2.0.3**

IRSA is the native way for AWS EKS (Amazon Elastic Kubernetes Service) to allow applications running in EKS pods to access the AWS API using permissions configured in AWS IAM roles.

To benefit from using the AWS IRSA credentials with PBM, the high-level steps are the following:

1. [Create a cluster](#) with `eksctl` and OIDC provider setup enabled. This feature works with EKS clusters version 1.13 and above.
2. Create an IAM role and specify the policy that defines the access to an S3 bucket.
3. Create a service account and annotate it with the IAM role.
4. Configure your pod by using the service account created in the previous step and assume the IAM role.
5. Provide the remote storage information for PBM in a config file. Leave the `s3.credentials` array empty, since PBM uses the `AWS_ROLE_ARN / AWS_WEB_IDENTITY_TOKEN_FILE` environment variables which are either automatically provided (i.e. injected by Kubernetes mutating admission controller in EKS) or which you can define manually (if you don't want to the admission controller to modify your pods)

Note

If IRSA-related credentials are defined, they have the priority over any IAM instance profile. However, if you intentionally specify S3 credentials in PBM configuration file, they override any IRSA/IAM instance profile related credentials and are used for authentication instead.

See also

AWS documentation:

- [Introducing fine-grained IAM roles for service accounts](#)
- [How do I use the IAM roles for service accounts \(IRSA\) feature with Amazon EKS to restrict access to an Amazon S3 bucket?](#)

7.2.3 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: March 29, 2023

Created: August 18, 2022

7.3 Percona Backup for MongoDB diagnostics tools

Percona Backup for MongoDB provides diagnostics tools to operate data backups.

7.3.1 pbm-speed-test

pbm-speed-test allows field-testing compression and backup upload speed of logical backups. You can use it:

- To check performance before starting a backup
- To find out what slows down the running backup

By default, **pbm-speed-test** operates with fake semi random data documents. To run **pbm-speed-test** on a real collection, provide a valid MongoDB connection URI string for the `--mongodb-uri` flag.

Run **pbm-speed-test** for the full set of available commands.

Compression test

```
$ pbm-speed-test compression --compression=s2 --size-gb 10
```

Output:

```
Test started ....
10.00GB sent in 8s.
Avg upload rate = 1217.13MB/s.
```

pbm-speed-test compression uses the compression library from the config file and sends a fake semi random data document (1 GB by default) to the black hole storage. (Use the `pbm config` command to change the compression library).

To test compression on a real collection, pass the `--sample-collection` flag with the `<my_db.my_collection>` value.

Run `pbm-speed-test compression --help` for the full set of supported flags:

```
$ pbm-speed-test compression --help
usage: pbm-speed-test compression

Run compression test

Flags:
  --help                Show context-sensitive help (also try
                        --help-long and --help-man).
  --mongodb-uri=MONGODB-URI  MongoDB connection string
  -c, --sample-collection=SAMPLE-COLLECTION
                        Set collection as the data source
  -s, --size-gb=SIZE-GB    Set data size in GB. Default 1
  --compression=s2        Compression type
                        <none>/<gzip>/<snappy>/<lz4>/<s2>/<pgzip>/<zstd>
  --compression-level=COMPRESSION-LEVEL
                        Compression level (specific to the compression type)
                        <none>/<gzip>/<snappy>/<lz4>/<s2>/<pgzip>/<zstd>
```

Upload speed test

```
$ pbm-speed-test storage --compression=s2
```

Output

```
Test started
1.00GB sent in 1s.
Avg upload rate = 1744.43MB/s.
```

`pbm-speed-test storage` sends the semi random data (1 GB by default) to the remote storage defined in the config file. Pass the `--size-gb` flag to change the data size.

To run the test with the real collection's data instead of the semi random data, pass the `--sample-collection` flag with the `<my_db.my_collection>` value.

Run `pbm-speed-test storage --help` for the full set of available flags:

```
$ pbm-speed-test storage --help
usage: pbm-speed-test storage

Run storage test

Flags:
  --help                Show context-sensitive help (also try --help-long and --
help-man).
  --mongodb-uri=MONGODB-URI  MongoDB connection string
  -c, --sample-collection=SAMPLE-COLLECTION
                        Set collection as the data source
  -s, --size-gb=SIZE-GB    Set data size in GB. Default 1
  --compression=s2        Compression type <none>/<gzip>/<snappy>/<lz4>/<s2>/<pgzip>/
<zstd>
  --compression-level=COMPRESSION-LEVEL
                        Compression level (specific to the compression type)
```

7.3.2 Backup progress tracking

If you have a large logical backup, you can track the backup progress in the logs of the `pbm-agent` that makes it. A line is appended every minute showing bytes copied vs. total size for the current collection.

Start a backup:

```
$ pbm backup
```

Check backup progress:

1. Check what `pbm-agent` makes the backup:

```
pbm logs
```

2. Connect to the `mongod` server where the `pbm-agent` is running and check its logs

```
$ journalctl -u pbm-agent.service
```

Sample output:

```
2020/05/06 21:31:12 Backup 2020-05-06T18:31:12Z started on node rs2/localhost:28018
2020-05-06T21:31:14.797+0300 writing admin.system.users to archive on stdout
2020-05-06T21:31:14.799+0300 done dumping admin.system.users (2 documents)
2020-05-06T21:31:14.800+0300 writing admin.system.roles to archive on stdout
2020-05-06T21:31:14.807+0300 done dumping admin.system.roles (1 document)
2020-05-06T21:31:14.807+0300 writing admin.system.version to archive on stdout
2020-05-06T21:31:14.815+0300 done dumping admin.system.version (3 documents)
2020-05-06T21:31:14.816+0300 writing test.testt to archive on stdout
2020-05-06T21:31:14.829+0300 writing test.testt2 to archive on stdout
2020-05-06T21:31:14.829+0300 writing config.cache.chunks.config.system.sessions to
archive on stdout
2020-05-06T21:31:14.832+0300 done dumping config.cache.chunks.config.system.sessions
(1 document)
2020-05-06T21:31:14.834+0300 writing config.cache.collections to archive on stdout
2020-05-06T21:31:14.835+0300 done dumping config.cache.collections (1 document)
2020/05/06 21:31:24 [##.....] test.testt 130841/1073901 (12.2%)
2020/05/06 21:31:24 [#####.....] test.testt2 131370/300000 (43.8%)
2020/05/06 21:31:24
2020/05/06 21:31:34 [#####.....] test.testt 249603/1073901 (23.2%)
2020/05/06 21:31:34 [#####.....] test.testt2 249603/300000 (83.2%)
2020/05/06 21:31:34
2020/05/06 21:31:37 [#####.....] test.testt2 300000/300000 (100.0%)
```

7.3.3 pbm-agent logs

 **Version added: 1.4.0**

To troubleshoot issues with specific events or node(s), use the `pbm logs` command. It provides logs of all `pbm-agent` processes in your environment.

`pbm logs` has the set of filters to refine logs for specific events like `backup`, `restore`, `pitr` or for a specific node, and to manage log verbosity level. For example, to view logs about a specific backup with the Debug verbosity level, run the `pbm logs` command as follows:

```
$ pbm logs --severity=D --event=backup/2020-10-15T17:42:54Z
```

To learn more about available filters and usage examples, refer to [Viewing backup logs](#).

7.3.4 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: March 29, 2023

Created: August 18, 2022

7.4 Percona Backup for MongoDB status

 **Version added: 1.4.0**

You can check the status of Percona Backup for MongoDB running in your MongoDB environment using the `pbm status` command.

```
$ pbm status
```

The output provides the information about:

- Your MongoDB deployment and `pbm-agents` running in it: to what `mongod` node each agent is connected, the Percona Backup for MongoDB version it runs and the agent's state
- The currently running backups / restores, if any
- Backups stored in the remote backup storage: backup name, completion time, size and status (complete, canceled, failed)
- [Point-in-time recovery](#) status (enabled or disabled)
- Valid time ranges for point-in-time recovery and the data size

This simplifies troubleshooting since the whole information is provided in one place.

Sample output

```
pbm status

Cluster:
=====
config:
- config/localhost:27027: pbm-agent v1.3.2 OK
- config/localhost:27028: pbm-agent v1.3.2 OK
- config/localhost:27029: pbm-agent v1.3.2 OK
rs1:
- rs1/localhost:27018: pbm-agent v1.3.2 OK
- rs1/localhost:27019: pbm-agent v1.3.2 OK
- rs1/localhost:27020: pbm-agent v1.3.2 OK
rs2:
- rs2/localhost:28018: pbm-agent v1.3.2 OK
- rs2/localhost:28019: pbm-agent v1.3.2 OK
- rs2/localhost:28020: pbm-agent v1.3.2 OK

PITR incremental backup:
=====
Status [OFF]

Currently running:
=====
(none)

Backups:
=====
S3 us-east-1 https://storage.googleapis.com/backup-test
  Snapshots:
    2020-12-16T10:36:52Z 491.98KB [restore_to_time: 2020-12-16T10:37:13Z]
    2020-12-15T12:59:47Z 284.06KB [restore_to_time: 2020-12-15T13:00:08Z]
    2020-12-15T11:40:46Z 0.00B [canceled: 2020-12-15T11:41:07Z]
```

```
2020-12-11T16:23:55Z 284.82KB [restore_to_time: 2020-12-11T16:24:16Z]
2020-12-11T16:22:35Z 284.04KB [restore_to_time: 2020-12-11T16:22:56Z]
2020-12-11T16:21:15Z 283.36KB [restore_to_time: 2020-12-11T16:21:36Z]
2020-12-11T16:19:54Z 281.73KB [restore_to_time: 2020-12-11T16:20:15Z]
2020-12-11T16:19:00Z 281.73KB [restore_to_time: 2020-12-11T16:19:21Z]
2020-12-11T15:30:38Z 287.07KB [restore_to_time: 2020-12-11T15:30:59Z]
PITR chunks:
2020-12-16T10:37:13 - 2020-12-16T10:43:26 44.17KB
```

7.4.1 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: March 29, 2023

Created: March 29, 2023

7.5 Upgrade Percona Backup for MongoDB

The recommended and most convenient way to upgrade PBM is from Percona repositories.

7.5.1 Important notes

1. Backward compatibility between data backup and restore is supported for upgrades within one major version only (for example, from 1.1.x to 1.2.y). When you upgrade Percona Backup for MongoDB over several major versions (for example, from 1.0.x to 1.2.y), we recommend to make a backup right after the upgrade.
2. Percona Backup for MongoDB v1.5.0 and later is incompatible with Percona Backup for MongoDB v1.4.1 and earlier due to different processing of system collections `Users` and `Roles` during backup / restore operations. After the upgrade to Percona Backup for MongoDB v1.5.0 and later, make sure to make a fresh backup.
3. Starting from v1.7.0, the user running the `pbm-agent` process is changed from `pbm` to `mongod`. This is done for the following reasons:
 - To make physical backups and restores, the user running the `pbm-agent` process must have the read / write permissions to the MongoDB `dataDir`.
 - To use the filesystem-based backup storage, the user running the `pbm-agent` process must also have the read / write permissions to the backup directory.
4. Starting from version 1.3.0, Percona Backup for MongoDB packages are stored in the `pbm` repository and the `tools` repository for backward compatibility.
5. Upgrade Percona Backup for MongoDB on all nodes where it is installed.

7.5.2 Prerequisites

Run all commands as root or via `sudo`.

1. Install `percona-release` tool. If you have installed it before, [update](#) it to the latest version.
2. Enable the repository

```
$ sudo percona-release enable pbm release
```

Note

For `apt`-based systems, run `sudo apt update` to update the local cache.

7.5.3 Upgrade to the latest version

On Debian and Ubuntu Linux

On Red Hat Enterprise Linux and derivatives

1. Stop `pbm-agent`

```
$ sudo systemctl stop pbm-agent
```

2. Install new packages

```
$ sudo apt install percona-backup-mongodb
```

3. Reload the `systemd` process

Starting from v1.7.0, reload the `systemd` process to update the unit file with the following command:

```
$ sudo systemctl daemon-reload
```

4. Update permissions

For a *filesystem-based backup storage*, grant read / write permissions to the backup directory to the `mongod` user.

5. Start `pbm-agent`

```
$ sudo systemctl start pbm-agent
```

1. Stop `pbm-agent`

```
$ sudo systemctl stop pbm-agent
```

2. Install new packages

```
$ sudo yum install percona-backup-mongodb
```

3. Reload the `systemd` process

Starting from v1.7.0, reload the `systemd` process to update the unit file with the following command:

```
$ sudo systemctl daemon-reload
```

4. Update permissions

For a *filesystem-based backup storage*, grant read / write permissions to the backup directory to the `mongod` user.

5. Start `pbm-agent`

```
$ sudo systemctl start pbm-agent
```

7.5.4 Upgrade to a specific version

On Debian and Ubuntu Linux

On Red Hat Enterprise Linux and derivatives

1. List available versions

```
$ sudo apt-cache madison percona-backup-mongodb
```

Output:

```
percona-backup-mongodb | 1.8.1-1.stretch | http://repo.percona.com/tools/apt stretch/main amd64 Packages
percona-backup-mongodb | 1.8.0-1.stretch | http://repo.percona.com/tools/apt stretch/main amd64 Packages
percona-backup-mongodb | 1.7.0-1.stretch | http://repo.percona.com/tools/apt stretch/main amd64 Packages
percona-backup-mongodb | 1.6.1-1.stretch | http://repo.percona.com/tools/apt stretch/main amd64 Packages
percona-backup-mongodb | 1.6.0-1.stretch | http://repo.percona.com/tools/apt stretch/main amd64 Packages
percona-backup-mongodb | 1.5.0-1.stretch | http://repo.percona.com/tools/apt stretch/main amd64 Packages
```

2. Stop pbm-agent

```
$ sudo systemctl stop pbm-agent
```

3. Install packages

Install a specific version packages. For example, to upgrade to Percona Backup for MongoDB 1.7.0, run the following command:

```
$ sudo apt install percona-backup-mongodb=1.7.0-1.stretch
```

4. Update permissions

For a *filesystem-based backup storage*, grant read / write permissions to the backup directory to the `mongod` user.

5. Start pbm-agent

```
$ sudo systemctl start pbm-agent
```

1. List available versions

```
$ sudo yum list percona-backup-mongodb --showduplicates
```

Output:

```
Available Packages
percona-backup-mongodb.x86_64      1.8-1.el7          pbm-release-x86_64
percona-backup-mongodb.x86_64      1.8.0-1.el7        pbm-release-x86_64
percona-backup-mongodb.x86_64      1.7.0-1.el7        pbm-release-x86_64
percona-backup-mongodb.x86_64      1.6.1-1.el7        pbm-release-x86_64
percona-backup-mongodb.x86_64      1.6.0-1.el7        pbm-release-x86_64
percona-backup-mongodb.x86_64      1.5.0-1.el7        pbm-release-x86_64
```

2. Stop pbm-agent

```
$ sudo systemctl stop pbm-agent
```

 **Note**

If MongoDB runs under a *different user than* `mongod` (the default configuration for Percona Server for MongoDB), use the same user to run the `pbm-agent`. For filesystem-based storage, grant the read / write permissions to the backup directory for this user.

7.5.5 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: March 29, 2023

Created: August 18, 2022

7.6 Uninstall Percona Backup for MongoDB

To uninstall Percona Backup for MongoDB, do the following steps:

1. Check that no backups are currently in progress in the output of `pbm list`.
2. Before the next 2 steps, make sure you know where the remote backup storage is, so you can delete backups made by Percona Backup for MongoDB. If it is an S3-compatible object storage, you will need to use another tool such as Amazon AWS's "aws s3", Minio's `mc`, the web AWS Management Console, etc. to do that once Percona Backup for MongoDB is uninstalled. Don't forget to note the connection credentials before they are deleted too.
3. Uninstall the **pbm-agent** and `pbm` executables. If you installed using a package manager, see [Install Percona Backup for MongoDB](#) for relevant package names and commands for your OS distribution.
4. Drop the PBM control collections.
5. Drop the PBM database user. If this is a cluster, the `dropUser` command will need to be run on each shard as well as in the config server replica set.
6. (Optional) Delete the backups from the remote backup storage.

7.6.1 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: March 29, 2023

Created: August 18, 2022

8. FAQ

8.1 What's the difference between PBM and `mongodump` ?

`mongodump` is a “logical” backup solution only while Percona Backup for MongoDB supports both logical and physical backups. Both solutions have equal performance for non-sharded replica sets. However, as opposed to `mongodump`, Percona Backup for MongoDB allows you to achieve the following goals:

- Make consistent backups and restores in sharded clusters.
- Backup / restore both the whole data set and specific namespaces -- databases and collections. (See [Selective backup and restore](#) for more information.)
- Restore your database to a specific point in time.
- Run backups / restores on each replica set in parallel while `mongodump` runs in one process on `mongos` node.

8.2 Why does Percona Backup for MongoDB use UTC timezone instead of server local timezone?

`pbm-agents` use UTC time zone by design. The reason behind this is to avoid user misunderstandings when replica set / cluster nodes are distributed geographically in different time zones.

Starting with version 2.0.1, you can change the time zone for `pbm logs` output.

8.3 Can I restore a single collection with Percona Backup for MongoDB?

Yes. Starting with version 2.0.0, you can restore a single collection with Percona Backup for MongoDB. This functionality is available for logical backups and restores only. To learn more, see [Selective backup and restore](#).

8.4 Can I back up specific shards in a cluster?

No, since this would result in backups with inconsistent timestamps across the cluster. Such backups would be invalid for restore.

Percona Backup for MongoDB backs up the whole state of a sharded cluster, and this guarantees data consistency during the restore.

8.5 Do I need to stop the balancer for PITR restore?

Yes. The preconditions for both Point-in-Time Recovery restore and regular restore are the same:

1. In a sharded cluster, stop the balancer.
2. Make sure no writes are made to the database during restore. This ensures data consistency.
3. Disable Point-in-Time Recovery if it is enabled. This is because oplog slicing and restore are exclusive operations and cannot be run together. Note that oplog slices made after the restore and before the next backup snapshot become invalid. Make a fresh backup and re-enable Point-in-Time Recovery.

8.6 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: March 29, 2023

Created: August 18, 2022

9. Reference

9.1 pbm commands

`pbm` CLI is the command line utility to control the backup system. This page describes `pbm` commands available in Percona Backup for MongoDB.

For how to get started with Percona Backup for MongoDB, see [Initial setup](#).

9.1.1 pbm backup

Creates a backup snapshot and saves it in the remote backup storage.

The command has the following syntax:

```
$ pbm backup [<flags>]
```

For more information about using `pbm backup`, see [Starting a backup](#)

The command accepts the following flags:

Flag	Description
<code>-t, --type</code>	The type of backup. Supported values: physical, logical (default), incremental, external . When not specified, Percona Backup for MongoDB makes a logical backup.
<code>--base</code>	For incremental backups only. Set the backup as the base and start tracking the incremental backup history to calculate and save the difference in data blocks for subsequent incremental backups.
<code>--compression</code>	Create a backup with compression. Supported compression methods: <code>gzip</code> , <code>snappy</code> , <code>lz4</code> , <code>s2</code> , <code>pgzip</code> , <code>zstd</code> . Default: <code>s2</code> The <code>none</code> value means no compression is done during backup.
<code>--compression-level</code>	Configure the compression level from 0 to 10. The default value depends on the compression method used.
<code>-o, --out=text</code>	Shows the output format as either plain text or a JSON object. Supported values: <code>text</code> , <code>json</code>
<code>--wait</code>	Wait for the backup to finish. The flag blocks the shell session.
<code>-l, --list-files</code>	For external backups only. Shows the list of files per node to copy.
<code>--ns="database.collection"</code>	Makes a logical backup of the specified namespace - the database and collection(s). To back up all collections in the database, specify the value in the <code>--ns="database.*"</code> format. In version 2.0.0, only a single namespace is supported for the backup.

JS output

```
{
  "name": "<backup_name>",
  "storage": "<my-backup-dir>"
}
```

9.1.2 pbm backup-finish

Closes the `backupCursor` and finishes the external backup. Must be run after running `pbm backup -t external`. To learn more, refer to [API for snapshot-based physical backups](#).

The command has the following syntax:

```
$ pbm backup-finish [backup-name]
```

9.1.3 pbm cancel-backup

Cancels a running backup. The backup is marked as canceled in the backup list.

The command accepts the following flags:

Flag	Description
<code>-o, --out=text</code>	Shows the output format as either plain text or a JSON object. Supported values: <code>text</code> , <code>json</code>

JSON output

```
{
  "msg": "Backup cancellation has started"
}
```

9.1.4 pbm cleanup

Deletes outdated backup snapshots and point-in-time recovery oplog slices.

The command has the following syntax:

```
pbm delete-pitr [<flags>]
```

The command accepts the following flags:

Flag	Description
<code>--older-than=TIMESTAMP</code>	Deletes backups older than date / time specified in the format: <ul style="list-style-type: none"> <code>-%Y-%M-%DT%H:%M:%S</code> (e.g. 2020-04-20T13:13:20), <code>-%Y-%M-%D</code> (e.g. 2020-04-20), <code>-%Xd</code> (e.g. 30d). Only days are supported
<code>-w, --wait</code>	Wait for the cleanup to finish. The flag blocks the shell session
<code>-y, --yes</code>	Cleans up the data storage without asking for a user's confirmation
<code>--dry-run</code>	Checks for the old data to be deleted without deleting it. Allows to verify what data to delete

9.1.5 pbm config

Sets, changes or lists Percona Backup for MongoDB configuration.

The command has the following syntax:

```
$ pbm config [<flags>] [<key>]
```

The command accepts the following flags:

Flag	Description
<code>--force-resync</code>	Resync backup list with the current storage
<code>--list</code>	List current settings
<code>--file=FILE</code>	Upload the config information from a YAML file
<code>--set=SET</code>	Set a new config option value. Specify the option in the <code><key.name=value></code> format.
<code>-o, --out=text</code>	Shows the output format as either plain text or a JSON object. Supported values: text, json

PE configuration output

```
{
  "pitr": {
    "enabled": false,
    "oplogSpanMin": 0
  },
  "storage": {
    "type": "filesystem",
    "s3": {
      "region": "",
      "endpointUrl": "",
      "bucket": ""
    },
    "azure": {},
    "filesystem": {
      "path": "<my-backup-dir>"
    }
  },
  "restore": {
    "batchSize": 500,
    "numInsertionWorkers": 10
  },
  "backup": {}
}
```

Setting a config value

```
[
  {
    "key": "pitr.enabled",
    "value": "true"
  }
]
```

9.1.6 pbm delete-backup

Deletes the specified backup snapshot or all backup snapshots that are older than the specified time. The command deletes backups that are not running regardless of the remote backup storage being used.

The following is the command syntax:

```
$ pbm delete-backup [<flags>] [<name>]
```

The command accepts the following flags:

Flag	Description
<code>--older-than=TIMESTAMP</code>	Deletes backups older than date / time specified in the format: <ul style="list-style-type: none"> - %Y-%M-%DT%H:%M:%S (e.g. 2020-04-20T13:13:20) or - %Y-%M-%D (e.g. 2020-04-20)
<code>--force</code>	Forcibly deletes backups without asking for user's confirmation
<code>--yes</code>	Deletes backups without asking for user's confirmation

9.1.7 pbm delete-pitr

Deletes oplog slices produced for Point-in-Time Recovery.

The command has the following syntax:

```
pbm delete-pitr [<flags>]
```

The command accepts the following flags:

Flag	Description
<code>-a, --all</code>	Deletes all oplog
<code>--older-than=TIMESTAMP</code>	Deletes oplog slices older than date / time specified in the format: <ul style="list-style-type: none"> - %Y-%M-%DT%H:%M:%S (e.g. 2020-04-20T13:13:20) or - %Y-%M-%D (e.g. 2020-04-20) <p>When you specify a timestamp, Percona Backup for MongoDB rounds it down to align with the completion time of the closest backup snapshot and deletes oplog slices that precede this time. Thus, extra slices remain. This is done to ensure oplog continuity. To illustrate, the PITR time range is <code>2021-08-11T11:16:21 - 2021-08-12T08:55:25</code> and backup snapshots are:</p> <pre>2021-08-12T08:49:46Z 13.49MB [restore_to_time: 2021-08-12T08:50:06] 2021-08-11T11:36:17Z 7.37MB [restore_to_time: 2021-08-11T11:36:38]</pre> <p>Say you specify the timestamp <code>2021-08-11T19:16:21</code>. The closest backup is <code>2021-08-11T11:36:17Z 7.37KB [restore_to_time: 2021-08-11T11:36:38]</code>. PBM rounds down the timestamp to <code>2021-08-11T11:36:38</code> and deletes all slices that precede this time. As a result, your PITR time range is <code>2021-08-11T11:36:38 - 2021-08-12T09:00:25</code>.</p> <p>NOTE: Percona Backup for MongoDB doesn't delete the oplog slices that follow the most recent backup. This is done to ensure point in time recovery from that backup snapshot. For example, if the snapshot is <code>2021-07-20T07:05:23Z [restore_to_time: 2021-07-21T07:05:44]</code> and you specify the timestamp <code>2021-07-20T07:05:45</code>, Percona Backup for MongoDB deletes only slices that were made before <code>2021-07-20T07:05:23Z</code>.</p>
<code>--force</code>	Forcibly deletes oplog slices without asking a user's confirmation
<code>-o, --out=json</code>	Shows the output as either the plain text (default) or a JSON object. Supported values: <code>text, json</code> .
<code>--yes</code>	Deletes backups without asking for user's confirmation

9.1.8 pbm describe-backup

Provides the detailed information about a backup:

- backup name
- type
- status
- namespaces - what was backed up during a selective backup
- size
- error message for failed backup
- last write timestamp
- last write time - human-readable indication of the last write
- last transition time - the timestamp when a backup changed its status
- cluster information: the replica set name, the backup status on this replica set, whether it is used as a config server replica set, last write timestamp
- replica set info: name, backup status, last write timestamp and last transition time, `mongod` security options, if encryption is configured.
- for snapshot-based backups, provides the list of files being copied
- for logical and selective backups, provides the list of collections included in the backup. Available with version 2.3.0.

The command has the following syntax:

```
$ pbm describe-backup [<backup-name>] [<flags>]
```

Flag	Description
<code>-o, --out=text</code>	Shows the status as either plain text or a JSON object. Supported values: <code>text</code> , <code>json</code>
<code>-l, --list-files</code>	Shows the list of files being copied for snapshot-based backups
<code>--with-collections</code>	Shows the collections included in the backup. For logical and selective backups only. Available with version 2.3.0.

JSON output

```
{
  "name": "<backup_name>",
  "opid": "<string>",
  "type": "logical",
  "last_write_ts": Timestamp,
  "last_transition_ts": Timestamp,
  "last_write_time": "2022-09-30T14:02:49Z",
  "last_transition_time": "2022-09-30T14:02:54Z",
  "namespaces": [
    "flight.booking"
  ],
  "mongodb_version": "<version>",
  "pbm_version": "<version>",
  "status": "done",
  "size": 470805945,
  "size_h": "449.0 MiB",
  "replsets": [
    {
      "name": "<name>",
      "status": "done",
      "last_write_ts": Timestamp,
      "last_transition_ts": Timestamp,
      "last_write_time": "2022-09-30T14:02:49Z",
      "last_transition_time": "2022-09-30T14:02:53Z"
    }
  ]
}
```

9.1.9 pbm describe-restore

Shows the detailed information about the restore.

The command has the following syntax:

```
$ pbm describe-restore [<restore-timestamp>] [<flags>]
```

The command accepts the following flags:

Flag	Description
<code>-c, --config=CONFIG</code>	Only for physical restores . Points Percona Backup for MongoDB to a configuration file so it can read the restore status from the remote storage. For example, <code>pbm describe-restore -c /etc/pbm/conf.yaml <restore-name></code> .
<code>-o, --out=TEXT</code>	Shows the output as either the plain text (default) or a JSON object. Supported values: <code>text, json</code> .

Selective restore status

```

{
  "name": "<restore_name>",
  "opid": "string",
  "backup": "<backup_name>",
  "type": "logical",
  "status": "done",
  "ts_to_restore": Timestamp,
  "time_to_restore": "Time",
  "namespaces": [
    "<database.*>"
  ]
  "replsets": [
    {
      "name": "rs1",
      "status": "done",
      "last_transition_ts": Timestamp,
      "last_transition_time": "Time"
    },
    {
      "name": "rs0",
      "last_transition_ts": Timestamp,
      "last_transition_time": "Time"
    },
    {
      "name": "cfg",
      "status": "done",
      "last_transition_ts": Timestamp,
      "last_transition_time": "Time"
    }
  ],
}

```

Physical restore status

```

{
  "name": "<restore_name>",
  "opid": "string",
  "backup": "<backup_name>",
  "type": "physical",
  "status": "done",
  "last_transition_ts": Timestamp,
  "last_transition_time": "Time",
  "replsets": [
    {
      "name": "rs1",
      "status": "done",
      "last_transition_ts": Timestamp,
      "last_transition_time": "Timestamp",
      "nodes": [
        {
          "name": "IP:port",
          "status": "done",
          "last_transition_ts": Timestamp,
          "last_transition_time": "Timestamp"
        }
      ]
    }
  ],
}

```

9.1.10 pbm help

Returns the help information about `pbm` commands.

9.1.11 pbm list

Provides the list of backups. In versions 1.3.4 and earlier, the command lists all backups and their states. Backup states are the following:

- In progress - A backup is running
- Canceled - A backup was canceled
- Error - A backup was finished with an error
- No status means a backup is complete

As of version 1.4.0, only successfully completed backups are listed. To view currently information about a running or a failed backup, run `pbm status`.

When Point-in-Time Recovery is enabled, the `pbm list` also provides the list of valid time ranges for recovery and point-in-time recovery status.

The command has the following syntax:

```
$ pbm list [<flags>]
```

The command accepts the following flags:

Flag	Description
<code>--restore</code>	Shows last N restores. Starting with version 2.0, the output shows restore names instead of backup names, as multiple restores can be done from a single backup.
<code>--size=0</code>	Shows last N backups. It also provides the information whether the restore is a selective one.
<code>-o, --out=text</code>	Shows the output format as either plain text or a JSON object. Supported values: <code>text</code> , <code>json</code>
<code>--unbacked</code>	Shows Point-in-Time Recovery oplog slices that were saved without the base backup snapshot. Available starting with version 1.8.0.
<code>--replset-remapping</code>	Maps the replica set names for the data restore / oplog replay. The value format is <code>to_name_1=from_name_1,to_name_2=from_name_2</code>

List of backups

```
{
  "snapshots": [
    {
      "name": "<backup_name>",
      "status": "done",
      "completeTS": Timestamp,
      "pbmVersion": "1.6.0"
    }
  ],
  "pitr": {
    "on": false,
    "ranges": [
      {
        "range": {
          "start": Timestamp,
          "end": Timestamp
        }
      },
      {
        "range": {
          "start": Timestamp,
          "end": Timestamp
        }
      },
      {
        "range": {
          "start": Timestamp,
          "end": Timestamp (no base snapshot)
        }
      }
    ]
  }
}
```

Restore history

Full restore

```
{
  "start": Timestamp,
  "status": "done",
  "type": "snapshot",
  "snapshot": "<backup_name>",
  "name": "<restore_name>"
}
```

Selective restore

```
{
  "start": Timestamp,
  "status": "done",
  "type": "snapshot",
  "snapshot": "<backup_name>",
  "name": "<restore_name>",
  "namespaces": [
    "<database.collection>"
  ]
}
```

Point-in-time restore

```
{
  "start": Timestamp,
  "status": "done",
  "type": "pitr",
  "snapshot": "<backup_name>",
  "point-in-time": Timestamp,
  "name": "<restore_name>"
}
```

Selective point-in-time restore

```
{
  "start": Timestamp,
  "status": "done",
  "type": "pitr",
  "snapshot": "<backup_name>",
  "point-in-time": Timestamp,
  "name": "<restore_name>",
  "namespaces": [
    "<database.collection>"
  ]
}
```

9.1.12 pbm logs

Shows log information from all `pbm-agent` processes.

The command has the following syntax:

```
$ pbm logs [<flags>]
```

The command accepts the following flags:

Flag	Description
<code>-t, --tail=20</code>	Shows last N entries. By default, the output shows last 20 entries. 0 means to show all log messages.
<code>-e, --event=EVENT</code>	Shows logs filtered by a specified event. Supported events: - backup - restore - resyncBcpList - pitr - pitrestore - delete
<code>-o, --out=text</code>	Shows log information as text (default) or in JSON format. Supported values: <code>text</code> , <code>json</code>
<code>-n, --node=NODE</code>	Shows logs for a specified node or a replica set. Specify the node in the format <code>replset[/host:port]</code>
<code>-f, --follow</code>	Follow log output. Allow to view the logs dynamically
<code>-s, --severity=I</code>	Shows logs filtered by severity level. Supported levels are (from low to high): D - Debug, I - Info (default), W - Warning, E - Error, F - Fatal. The output includes both the specified severity level and all higher ones
<code>-- timezone =TIMEZONE</code>	Timezone of the log output. Supported values: <code>UTC</code> (default), <code>local</code> or the timezone in the IANA timezone format (e.g. <code>America/New_York</code>)
<code>-i, --opid=OPID</code>	Show logs for an operation in progress. The operation is identified by the OpID
<code>-x, --extra</code>	Show extra data in the text format

Find the usage examples in [Viewing backup logs](#).

Log output

```
[
  {
    "t": "",
    "s": 3,
    "rs": "rs0",
    "node": "example.mongodb.com:27017",
    "e": "",
    "eobj": "",
    "ep": {
      "T": 0,
      "I": 0
    },
    "msg": "listening for the commands"
  },
  ....
]
```

9.1.13 pbm oplog-replay

Allows to replay the oplog on top of any backup: logical, physical, storage level snapshot (like EBS-snapshot) and restore it to a specific point in time.

To learn more about the usage, refer to Point-in-Time Recovery oplog replay.

The command has the following syntax:

```
$ pbm oplog-replay [<flags>]
```

The command accepts the following flags:

Flag	Description
start=timestamp	The start time for the oplog replay.
end=timestamp	The end time for the oplog replay.
--replset-remapping	Maps the replica set names for the oplog replay. The value format is to_name_1=from_name_1,to_name_2=from_name_2.

9.1.14 pbm restore

Restores database from a specified backup / to a specified point in time. Depending on the backup type, makes either logical, physical, or a snapshot-based restore.

The command has the following syntax:

```
$ pbm restore [<flags>] [<backup_name>]
```

For more information about using `pbm restore`, see [Restoring a backup](#).

The command accepts the following flags:

Flag	Description
<code>--external</code>	Indicates the backup as the one made outside PBM (for example, snapshot-based)
<code>--time=TIME</code>	Restores the database to the specified point in time. Available for logical restores and if Point-in-time recovery is enabled.
<code>-w</code>	Wait for the restore to finish. The flag blocks the shell session.
<code>-o, --out=text</code>	Shows the output format as either plain text or a JSON object. Supported values: <code>text</code> , <code>json</code>
<code>--base-snapshot</code>	Restores the database from a specified backup to the specified point in time. Without this flag, the most recent backup preceding the timestamp is used for point in recovery. Available in Percona Backup for MongoDB starting from version 1.6.0. In version 2.3.0, this flag is optional for point-in-time recovery from physical backups . In version 2.2.0, this flag is mandatory for making a point-in-time recovery from physical backups . Without it, PBM looks for a logical backup to restore from.
<code>--replset-remapping</code>	Maps the replica set names for the data restore / oplog replay. The value format is <code>to_name_1=from_name_1,to_name_2=from_name_2</code>
<code>--ns="database.collection"</code>	Restores the specified namespace(s) - databases and collections. To restore all collections in the database, specify the values as <code>--ns="database.*"</code> . The <code>--ns</code> flag accepts several namespaces as the comma-separated list. For example, <code>ns="db1.*,db2.coll2,db3.coll1,db3.collX"</code>
<code>-c, --config</code>	The path to the <code>mongod.conf</code> file

Restore output

```
{
  "name": "<restore_name>"
  "snapshot": "<backup_name>"
}
```

Point-in-time restore

```
{
  "name": "<restore_name>",
  "point-in-time": "<backup_name>"
}
```

9.1.15 pbm restore-finish

Instructs PBM to complete the snapshot-based physical restore. Must be run after running `pbm restore --external`. To learn more, refer to [API for snapshot-based physical backups](#).

The command has the following syntax:

```
$ pbm restore-finish <restore_name> [flags]
```

The command accepts the following flags:

Flag	Description
<code>-c</code>	The path to the PBM configuration file. Required to complete the restore.

9.1.16 pbm status

Shows the status of Percona Backup for MongoDB. The output provides the following information:

- `pbm-agent` processes version and state
- Currently running backups or restores
- Backups stored in the remote storage
- Point-in-Time Recovery status
- Valid time ranges for point-in-time recovery and the data size

The command accepts the following flags:

Flag	Description
<code>-o, --out=text</code>	Shows the status as either plain text or a JSON object. Supported values: <code>text</code> , <code>json</code>
<code>-s, --sections=SECTIONS</code>	Shows the status for the specified section. You can pass several flags to view the status for multiple sections. Supported values: <code>cluster</code> , <code>pitr</code> , <code>running</code> , <code>backups</code> .
<code>--replset-remapping</code>	Maps the replica set names for the data restore / oplog replay. The value format is <code>to_name_1=from_name_1,to_name_2=from_name_2</code>

Status information

```

{
  "backups": {
    "type": "FS",
    "path": "<my-backup-dir>",
    "snapshot": [
      ...
      {
        "name": "<backup_name>",
        "size": 3143396168,
        "status": "done",
        "completeTS": Timestamp,
        "pbmVersion": "1.6.0"
      },
    ],
  },
  "pitrChunks": {
    "pitrChunks": [
      ...
      {
        "range": {
          "start": Timestamp,
          "end": Timestamp
        }
      },
      {
        "range": {
          "start": Timestamp,
          "end": Timestamp (no base snapshot) !!! no backup found
        }
      },
    ],
  },
  "size": 677901884
},
"cluster": [
  {
    "rs": "<replSet_name>",
    "nodes": [
      {
        "host": "<replSet_name>/example.mongodb:27017",
        "agent": "<version>",
        "ok": true
      }
    ]
  }
],
"pitr": {
  "conf": true,
  "run": false,
  "error": "Timestamp.000+0000 E [<replSet_name>/example.mongodb:27017] [pitr] <error_message>"
},
"running": {
  "type": "backup",
  "name": "<backup_name>",
  "startTS": Timestamp,
  "status": "oplog backup",
  "opID": "6113b631ea9ba5b815fee7c6"
}
}

```

9.1.17 pbm version

Shows the version of Percona Backup for MongoDB.

The command accepts the following flags:

Flag	Description
<code>--short</code>	Shows only version info
<code>--commit</code>	Shows only git commit info
<code>-o, --out=text</code>	Shows the output as either plain text or a JSON object. Supported values: <code>text</code> , <code>json</code>

Version information

```
{
  "Version": "1.6.0",
  "Platform": "linux/amd64",
  "GitCommit": "f9b9948bb8201ba1a6400f6558496934a0685efd",
  "GitBranch": "main",
  "BuildTime": "2021-07-28_15:24_UTC",
  "GoVersion": "go1.16.6"
}
```

9.1.18 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: September 25, 2023

Created: August 18, 2022

9.2 Configuration file options

9.2.1 Percona Backup for MongoDB configuration in a cluster (or non-sharded replica set)

The configuration information is stored in a single document of the `admin.pbmConfig` collection. That single copy is shared by all the `pbm-agent` processes in a cluster (or non-sharded replica set), and can be read or updated using the `pbm` CLI tool.

You can see the whole config by running

```
> db.getSiblingDB("admin").pbmConfig.findOne()
```

But you don't have to use the `mongo` shell; the `pbm` CLI has a "config" subcommand to read and update it.

Percona Backup for MongoDB config contains the following settings:

- [Remote backup storage configuration](#) is available starting with version 1.0 or 1.1.
- [Point-in-time recovery configuration](#) is available starting with version 1.3.0.
- [Restore options](#) are available as starting with version 1.3.2.

Run `pbm config --list` to see the whole config. (Sensitive fields such as keys will be redacted.)

Insert the whole Percona Backup for MongoDB config from a YAML file

If you are initializing a cluster or a non-sharded replica set for the first time, it is simplest to write the whole config as YAML file and use the `pbm config --file` command to upload all the values in one command.

Find the config file examples for the remote backup storage (required) in the [Example config files](#) section. For more information about available config file options, see [Configuration file options](#).

Use the following command to upload the config file. For example, config file name is `pbm_config.yaml`:

```
$ pbm config --file pbm_config.yaml
```

Execute the command while connecting to config server replica set if it is a cluster. Otherwise just connect to the non-sharded replica set as normal. (See [MongoDB connection strings - A Reminder \(or Primer\)](#) if you are not familiar with MongoDB connection strings yet.)

Accessing or updating single config values

You can set a single value at a time. For nested values, use dot-concatenated key names as shown in the following example:

```
pbm config --set storage.s3.bucket="operator-testing"
```

To list a single value, you can specify just the key name by itself. If set, the command returns the value.

Success No value

```
pbm config storage.s3.bucket  
operator-testing
```

```
pbm config storage.s3.INVALID-KEY  
Error: unable to get config key: invalid config key
```

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: October 6, 2023

Created: August 18, 2022

9.2.2 Remote backup storage options

Common options

STORAGE.TYPE

Type: string

Required: YES

Remote backup storage type. Supported values: `s3`, `filesystem`, `azure`.

S3 type storage options

```
storage:
  type: s3
  s3:
    region: <string>
    bucket: <string>
    prefix: <string>
    endpointUrl: <string>
    credentials:
      access-key-id: <your-access-key-id-here>
      secret-access-key: <your-secret-key-here>
      session-token: <string>
    uploadPartSize: <int>
    maxUploadParts: <int>
    storageClass: <string>
    serverSideEncryption:
      sseAlgorithm: aws:kms
      kmsKeyID: <your-kms-key-here>
      sseCustomerAlgorithm: AES256
      sseCustomerKey: <your_encryption_key>
    retryer:
      numMaxRetries: 3
      minRetryDelay: 30
      maxRetryDelay: 5
```

STORAGE.S3.PROVIDER

Type: string

Required: NO

The storage provider's name.

Supported values: `aws`, `gcs`

STORAGE.S3.BUCKET

Type: string

Required: YES

The name of the storage bucket. See the [AWS Bucket naming rules](#) and [GCS bucket naming guidelines](#) for bucket name requirements.

STORAGE.S3.REGION

Type: string

Required: YES (for AWS and GCS)

The location of the storage bucket. Use the [AWS region list](#) and [GCS region list](#) to define the bucket region

STORAGE.S3.PREFIX

Type: string
Required: NO

The path to the data directory on the bucket. If undefined, backups are stored in the bucket root directory

STORAGE.S3.ENDPOINTURL

Type: string
Required: YES (for MinIO and GCS)

The URL to access the bucket. The default value for GCS is `https://storage.googleapis.com`

STORAGE.S3.FORCEPATHSTYLE

Type: boolean
Required: NO

By default, PBM uses the [path-style URLs](#) for accessing the S3 bucket. Setting this option to `false` instructs PBM to send virtual-hosted-style requests instead.

STORAGE.S3.CREDENTIALS.ACCESS-KEY-ID

Type: string
Required: YES

Your access key to the storage bucket. This option can be omitted when you run Percona Backup for MongoDB using an EC2 instance profile. To learn more, refer to [Automate access to S3 buckets for Percona Backup for MongoDB](#)

STORAGE.S3.CREDENTIALS.SECRET-ACCESS-KEY

Type: string
Required: YES

The key to sign your programmatic requests to the storage bucket. This option can be omitted when you run Percona Backup for MongoDB using an EC2 instance profile. To learn more, refer to [Automate access to S3 buckets for Percona Backup for MongoDB](#)

STORAGE.S3.CREDENTIALS.SESSION-TOKEN

Type: string
Required: NO

The AWS session token used to validate the [temporary security credentials](#) for accessing the S3 storage.

STORAGE.S3.UPLOADPARTSIZE

Type: int
Required: NO

The size of data chunks in bytes to be uploaded to the storage bucket. Default: 10MB

Percona Backup for MongoDB automatically increases the `uploadPartSize` value if the size of the file to be uploaded exceeds the max allowed file size. (The max allowed file size is calculated with the default values of `uploadPartSize` * `maxUploadParts` and is appr. 97,6 GB).

The `uploadPartSize` value is printed in the `pbm-agent` log.

By setting this option, you can manually adjust the size of data chunks if Percona Backup for MongoDB failed to do it for some reason. The defined `uploadPartSize` value overrides the default value and is used for calculating the max allowed file size

STORAGE.S3.MAXUPLOADPARTS

Type: int
Required: NO
Default: 10,000

The maximum number of data chunks to be uploaded to the storage bucket. Default: 10,000

By setting this option, you can override the value defined in the [AWS SDK](#).

It can be useful when using an S3 provider that supports a smaller number of chunks for multipart uploads.

The `maxUploadParts` value is printed in the `pbm-agent` log.

STORAGE.S3.STORAGECLASS

Type: string
Required: NO

The [storage class](#) assigned to objects stored in the S3 bucket. If not provided, the `STANDARD` storage class will be used. This option is available in Percona Backup for MongoDB as of v1.7.0.

STORAGE.S3.DEBUGLOGLEVELS

Type: string
Required: NO

Enables S3 debug logging for different types of S3 requests. S3 log messages are printed in the `pbm Logs` output.

Supported values are: `LogDebug`, `Signing`, `HTTPBody`, `RequestRetries`, `RequestErrors`, `EventStreamBody`.

To specify several event types, separate them by comma. To learn more about the event types, see [the documentation](#)

When undefined, no S3 debug logging is performed.

STORAGE.S3.INSECURESKIPTLSVERIFY

Type: bool
Required: NO
Default: False

Disables the TLS verification of the S3 storage. This allows Percona Backup for MongoDB to upload data to S3-like storages that use self-issued TLS certificates. Available in Percona Backup for MongoDB as of version 1.7.0.

 **Warning**

Use this option with caution as it might leave a hole for man-in-the-middle attacks.

Server-side encryption options**SERVERSIDEENCRYPTION.SSEALGORITHM**

Type: string
Required: NO

The key management mode used for server-side encryption with the encryption keys stored in AWS KMS.

Supported value: `aws:kms`

SERVERSIDEENCRYPTION.KMSKEYID

Type: string
Required: NO

Your customer-managed key stored in the AWS KMS.

SERVERSIDEENCRYPTION.SSECUSTOMERALGORITHM

Type: string
Required: NO

The key management mode for [server-side encryption with customer-provided keys \(SSE-C\)](#).

Supported value: AES256

SERVERSIDEENCRYPTION.SSECUSTOMERKEY

Type: string
Required: NO

Your custom encryption key. This key is not stored on the S3 storage side. Thus, it is your responsibility to track what data is encrypted with what key and for storing the key.

Upload retry options

RETRYER.NUMMAXRETRIES

Type: int
Required: NO
Default: 3

The maximum number of retries to upload data to S3 storage. A zero value means no retries will be performed. Available in Percona Backup for MongoDB as of 1.7.0.

RETRYER.MINRETRYDELAY

Type: time.Duration
Required: NO
Default: 30

The minimum time (in ms) to wait till the next retry. Available in Percona Backup for MongoDB as of 1.7.0.

RETRYER.MAXRETRYDELAY

Type: time.Duration
Required: NO
Default: 5

The maximum time (in minutes) to wait till the next retry. Available in Percona Backup for MongoDB as of 1.7.0.

Filesystem storage options

```
storage:
  type: filesystem
  filesystem:
    path: <string>
```

STORAGE.FILESYSTEM.PATH

Type: string
Required: YES

The path to the backup directory

Microsoft Azure Blob storage options

```
storage:
  type: azure
  azure:
    account: <string>
    container: <string>
    endpointUrl: <string>
    prefix: <string>
    credentials:
      key: <your-access-key>
```

STORAGE.AZURE.ACCOUNT

Type: string
Required: YES

The name of your storage account.

STORAGE.AZURE.CONTAINER

Type: string
Required: YES

The name of the storage container. See the [Container names](#) for naming conventions.

STORAGE.AZURE.ENDPOINTURL

Type: string
Required: NO

The URL to access the data in Microsoft Azure Blob Storage. The default value is `https://<storage-account>.blob.core.windows.net`.

STORAGE.AZURE.PREFIX

Type: string
Required: NO

The path (sub-folder) to the backups inside the container. If undefined, backups are stored in the container root directory.

STORAGE.AZURE.CREDENTIALS.KEY

Type: string
Required: YES

Your access key to authorize access to data in your storage account.

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)

 [Get a Percona Expert](#)

Last update: December 4, 2023

Created: August 18, 2022

9.2.3 Point-in-time recovery options

```

pitr:
  enabled: <boolean>
  oplogSpanMin: <float64>
  compression: <string>
  compressionLevel: <int>

```

PITR.ENABLED

Type: boolean

Default: False

Enables point-in-time recovery

PITR.OPLOGSPANMIN

Type: float64

Default: 10

The duration of an oplog span in minutes. If set when the `pbm-agent` is making an oplog slice, the slice's span is updated right away.

If the new duration is smaller than the previous one, the `pbm-agent` is triggered to save a new slice with the updated span. If the duration is larger, then the next slice is saved with the updated span in scheduled time.

PITR.COMPRESSION

Type: string

Default: s2

The compression method for Point-in-Time Recovery oplog slices. Available in Percona Backup for MongoDB as of version 1.7.0.

Supported values: `gzip`, `snappy`, `lz4`, `s2`, `pgzip`, `zstd`. *Default:* `s2`.

PITR.COMPRESSIONLEVEL

Type: int

The compression level is from 0 till 10. The default value depends on the compression method used.

The following table shows available compression levels per compression method:

Compression method	Compression levels	Default
<code>zstd</code>	1 - fastest speed, 2 - default, 3 - better compression, 4 - best compression	2
<code>snappy</code>	no levels	
<code>lz4</code>	From 1 (fastest) to 16	1
<code>gzip</code> and <code>pgzip</code>	-1 - default compression, 0 - no compression, 1 - best speed, 9 - best compression	-1

Note that the greater value you specify, the more time and computing resources it will take to compress the data.

PITR.OPLOGONLY

Type: boolean

Default: False

Required: NO

Controls whether the base backup is required to start [Point-in-Time recovery oplog slicing](#). When set to true, Percona Backup for MongoDB saves oplog chunks without the base backup snapshot.

Available in Percona Backup for MongoDB starting with version 1.8.0. To learn more about the usage, see [Point-in-Time Recovery oplog replay](#).

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: June 30, 2023

Created: August 18, 2022

9.2.4 Backup options

```

backup:
  priority:
    "localhost:28019": 2.5
    "localhost:27018": 2.5
    "localhost:27020": 2.0
    "localhost:27017": 0.1
  compression: <string>
  compressionLevel: <int>
  timeouts:
    startingStatus: 60

```

PRIORITY

Type: array of strings

The list of `mongod` nodes and their priority for making backups. The node with the highest priority is elected for making a backup. If several nodes have the same priority, the one among them is randomly elected to make a backup.

If not set, the replica set nodes have the default priority as follows:

- hidden nodes - 2.0
- secondary nodes - 1.0
- primary node - 0.5

BACKUP.COMPRESSION

Type: string

Default: s2

The compression method for backup snapshots. Available in Percona Backup for MongoDB as of version 1.8.0.

When `none` is specified, backups are made without compression.

Supported values: `gzip`, `snappy`, `lz4`, `s2`, `pgzip`, `zstd`. Default: `s2`.

BACKUP.COMPRESSIONLEVEL

Type: int

The compression level. The default value depends on the compression method used.

The following table shows available compression levels per compression method:

Compression method	Supported compression levels	Default
<code>zstd</code>	1 - fastest speed, 2 - default, 3 - better compression, 4 - best compression	2
<code>snappy</code>	no levels	
<code>lz4</code>	From 1 (fastest) to 16	1
<code>gzip</code> and <code>pgzip</code>	-1 - default compression, 0 - no compression, 1 - best speed, 9 - best compression	-1

Note that the greater value you specify, the more time and computing resources it will take to compress the data.

BACKUP.TIMEOUTS.STARTINGSTATUS

Type: unit32

Default: 33

The wait time (in seconds) for PBM to start physical backups on all shards. Increasing this value is useful when it takes longer than usual to open the `$backupCursor`.

The 0 (zero) value resets the timeout to the default 33 seconds.

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: July 27, 2023

Created: August 18, 2022

9.2.5 Restore options

```
restore:
  batchSize: <int>
  numInsertionWorkers: <int>
  numDownloadWorkers: <int>
  maxDownloadBufferMb: <int>
  downloadChunkMb: <int>
  mongodLocation: <string>
  mongodLocationMap:
    "node01:2017": <string>
    "node03:27017": <string>
```

BATCHSIZE

Type: int
Default: 500

The number of documents to buffer.

NUMINSERTIONWORKERS

Type: int
Default: 10

The number of workers that add the documents to buffer.

NUMDOWNLOADWORKERS

Type: int
Default: number of CPU cores

The number of workers that request data chunks from the storage during the restore. The default value equals to the number of CPU cores.

MAXDOWNLOADBUFFERMB

Type: int

The maximum size of the in-memory buffer that is used to download files from the S3 storage. When unspecified or set to 0, the size cannot exceed the value calculated as `numDownloadWorkers * downloadChunkMb * 16` MB. By default, the number of CPU cores * 32 * 16 MB.

DOWNLOADCHUNKMB

Type: int
Default: 32

The size of the data chunk in MB to download from the S3 storage.

MONGODLOCATION

Type: string

The custom path to `mongod` binaries. When undefined, Percona Backup for MongoDB uses the default path to make database restarts during physical restore.

MONGODLOCATIONMAP

Type: array of strings

The list of custom paths to `mongod` binaries on every node. Percona Backup for MongoDB uses the values to make restarts of the database during physical restore.

Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: April 18, 2023

Created: August 18, 2022

9.3 Glossary

9.3.1 ACID

Set of properties that guarantee database transactions are processed reliably. Stands for [Atomicity](#), [Consistency](#), [Isolation](#), [Durability](#).

9.3.2 Amazon S3

Amazon S3 (Simple Storage Service) is an object storage service provided through a web service interface offered by Amazon Web Services.

9.3.3 Atomicity

Atomicity means that database operations are applied following an “all or nothing” rule. A transaction is either fully applied or not at all.

9.3.4 Blob

A blob stands for Binary Large Object, which includes objects such as images and multimedia files. In other words these are various data files that you store in Microsoft’s data storage platform. Blobs are organized in [containers](#) which are kept in Azure Blob storage under your storage account.

9.3.5 Bucket

A bucket is a container on the s3 remote storage that stores backups.

9.3.6 Collection

A collection is the way data is organized in MongoDB. It is analogous to a table in relational databases.

9.3.7 Completion time

Starting with version 2.0.0, the completion time is renamed “restore_to_time”

The completion time is the time to which the sharded cluster / non-shared replica set will be returned to after the restore. It is reflected in the “complete” section of the `pbm list` / `pbm status` command outputs.

In `logical` backups, the completion time almost coincides with the backup finish time. To define the completion time, Percona Backup for MongoDB waits for the backup snapshot to finish on all cluster nodes. Then it captures the oplog from the backup start time up to that time.

In `physical` backups, the completion time is only a few seconds after the backup start time. By holding the `$backupCursor` open, Percona Backup for MongoDB guarantees that the checkpoint data won’t change during the backup. In such a way Percona Backup for MongoDB can define the completion time ahead.

9.3.8 Consistency

In the context of backup and restore, consistency means that the data restored will be consistent in a given point in time. Partial or incomplete writes to disk of atomic operations (for example, to table and index data structures separately) won’t be served to the client after the restore. The same applies to multi-document transactions that started but didn’t complete by the time the backup was finished.

9.3.9 Container

A container is like a directory in Azure Blob storage that contains a set of [blobs](#).

9.3.10 Durability

Once a transaction is committed, it will remain so.

9.3.11 EBS-snapshot

An EBS (Amazon Elastic Block Storage) snapshot is the point-in-time copy of your data, and can be used to enable disaster recovery, migrate data across regions and accounts, and improve backup compliance.

9.3.12 GCP

GCP (Google Cloud Platform) is the set of services, including storage service, that runs on Google Cloud infrastructure.

9.3.13 Isolation

The Isolation requirement means that no transaction can interfere with another.

9.3.14 Jenkins

[Jenkins](#) is a continuous integration system that we use to help ensure the continued quality of the software we produce. It helps us achieve the aims of:

- No failed tests in trunk on any platform
- Aid developers in ensuring merge requests build and test on all platforms,
- No known performance regressions (without a damn good explanation).

9.3.15 MinIO

MinIO is a cloud storage server compatible with [Amazon S3](#), released under Apache License v2.

9.3.16 Oplog

Oplog (operations log) is a fixed-size collection that keeps a rolling record of all operations that modify data in the database.

9.3.17 Oplog slice

A compressed bundle of [oplog](#) entries stored in the Oplog Store database in MongoDB. The oplog size captures an approximately 10-minute frame. For a snapshot, the oplog size is defined by the time that the slowest replica set member requires to perform mongodump.

9.3.18 OpID

A unique identifier of an operation such as backup, restore, resync. When a pbm-agent starts processing an operation, it acquires a lock and an opID. This prevents processing the same operation twice (for example, if there are network issues in distributed systems). Using opID as a log filter allows viewing logs for an operation in progress.

9.3.19 pbm-agent

A `pbm-agent` is a PBM process running on the `mongod` node for backup and restore operations. A `pbm-agent` instance is required for every `mongod` node (including replica set secondary members and config server replica set nodes).

9.3.20 pbm CLI

Command-line interface for controlling the backup system. PBM CLI can connect to several clusters so that a user can manage backups on many clusters.

9.3.21 PBM Control collections

PBM Control collections are [collections](#) with config, authentication data and backup states. They are stored in the `admin` db in the cluster or non-sharded replica set and serve as the communication channel between `pbm-agent` and `pbm CLI`. `pbm CLI` creates a new `pbmCmd` document for a new operation. `pbm-agents` monitor it and update as they process the operation.

9.3.22 Percona Backup for MongoDB

Percona Backup for MongoDB (PBM) is a low-impact backup solution for MongoDB non-sharded replica sets and clusters. It supports both [Percona Server for MongoDB](#) and MongoDB Community Edition.

9.3.23 Percona Server for MongoDB

Percona Server for MongoDB is a drop-in replacement for MongoDB Community Edition with enterprise-grade features.

9.3.24 Point-in-Time Recovery

Point-in-Time Recovery is restoring the database up to a specific moment in time. The data is restored from the backup snapshot and then events that occurred to the data are replayed from `oplog`.

9.3.25 Replica set

A replica set is a group of `mongod` nodes that host the same data set.

9.3.26 S3 compatible storage

This is the storage that is built on the [S3](#) API.

9.3.27 Server-side encryption

Server-side encryption is the encryption of data by the remote storage server as it receives it. The data is encrypted when it is written to S3 bucket and decrypted when you access the data.

9.3.28 Technical preview feature

Technical preview features are not yet ready for enterprise use and are not included in support via SLA. They are included in this release so that users can provide feedback prior to the full release of the feature in a future GA release (or removal of the feature if it is deemed not useful). This functionality can change (APIs, CLIs, etc.) from tech preview to GA.

9.3.29 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: March 29, 2023

Created: August 18, 2022

10. Contribution guidelines

We're glad that you would like to become a Percona community member and participate in keeping open source open.

There are many ways how you can contribute:

- Submit bug reports or feature requests.
- Submit a code patch.
- Contribute to documentation.

11. Submit bug reports or feature requests

If you find a bug in Percona Backup for MongoDB, you can submit a report to the [JIRA issue tracker](#) for Percona Backup for MongoDB.

Start by searching the open tickets for a similar report. If you find that someone else has already reported your problem, then you can upvote that report to increase its visibility.

If there is no existing report, submit a report following these steps:

1. Sign in to [JIRA issue tracker](#). You will need to create an account if you do not have one.
2. In the *Summary*, *Description*, *Steps To Reproduce*, *Affects Version* fields describe the problem you have detected. For PBM the important diagnostic information is: log files from the pbm-agents; a dump of the PBM control collections.

As a general rule of thumb, try to create bug reports that are:

- *Reproducible*: Describe the steps to reproduce the problem.
- *Specific*: Include the version of Percona Backup for MongoDB, your environment, and so on.
- *Unique*: Check if there already exists a JIRA ticket to describe the problem.
- *Scoped to a Single Bug*: Report only one bug in one JIRA ticket.

11.1 Submit a code patch

If you'd like to submit a code patch, follow the [Contributing guide in PBM code repository](#).

11.2 Contribute to documentation

11.2.1 Edit documentation online via GitHub

1. Click the **Edit this page** icon next to the page title. The source `.md` file of the page opens in GitHub editor in your browser. If you haven't worked with the repository before, GitHub creates a [fork](#) of it for you.
2. Edit the page. You can check your changes on the **Preview** tab.
3. Commit your changes.
 - In the *Commit changes* section, describe your changes.
 - Select the **Create a new branch for this commit and start a pull request** option
 - Click **Propose changes**.

4. GitHub creates a branch and a commit for your changes. It loads a new page on which you can open a pull request to Percona. The page shows the base branch – the one you offer your changes for, your commit message and a diff – a visual representation of your changes against the original page. This allows you to make a last-minute review. When you are ready, click the **Create pull request** button.
5. Someone from our team reviews the pull request and if everything is correct, merges it into the documentation. Then it gets published on the site.

11.2.2 Edit documentation locally

If you wish to have full control over the documentation process, follow the [Contributing to documentation guide](#).

11.3 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: March 29, 2023

Created: August 18, 2022

12. Copyright and licensing information

12.1 Documentation licensing

Percona Backup for MongoDB documentation is (C)2016–2023 Percona LLC and/or its affiliates and is distributed under the [Creative Commons Attribution 4.0 International Public License](#) license.

12.2 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: June 28, 2023

Created: August 18, 2022

13. Trademark Policy

This [Trademark Policy](#) is to ensure that users of Percona-branded products or services know that what they receive has really been developed, approved, tested and maintained by Percona. Trademarks help to prevent confusion in the marketplace, by distinguishing one company's or person's products and services from another's.

Percona owns a number of marks, including but not limited to Percona, XtraDB, Percona XtraDB, XtraBackup, Percona XtraBackup, Percona Server, and Percona Live, plus the distinctive visual icons and logos associated with these marks. Both the unregistered and registered marks of Percona are protected.

Use of any Percona trademark in the name, URL, or other identifying characteristic of any product, service, website, or other use is not permitted without Percona's written permission with the following three limited exceptions.

First, you may use the appropriate Percona mark when making a nominative fair use reference to a bona fide Percona product.

Second, when Percona has released a product under a version of the GNU General Public License ("GPL"), you may use the appropriate Percona mark when distributing a verbatim copy of that product in accordance with the terms and conditions of the GPL.

Third, you may use the appropriate Percona mark to refer to a distribution of GPL-released Percona software that has been modified with minor changes for the sole purpose of allowing the software to operate on an operating system or hardware platform for which Percona has not yet released the software, provided that those third party changes do not affect the behavior, functionality, features, design or performance of the software. Users who acquire this Percona-branded software receive substantially exact implementations of the Percona software.

Percona reserves the right to revoke this authorization at any time in its sole discretion. For example, if Percona believes that your modification is beyond the scope of the limited license granted in this Policy or that your use of the Percona mark is detrimental to Percona, Percona will revoke this authorization. Upon revocation, you must immediately cease using the applicable Percona mark. If you do not immediately cease using the Percona mark upon revocation, Percona may take action to protect its rights and interests in the Percona mark. Percona does not grant any license to use any Percona mark for any other modified versions of Percona software; such use will require our prior written permission.

Neither trademark law nor any of the exceptions set forth in this Trademark Policy permit you to truncate, modify or otherwise use any Percona mark as part of your own brand. For example, if XYZ creates a modified version of the Percona Server, XYZ may not brand that modification as "XYZ Percona Server" or "Percona XYZ Server", even if that modification otherwise complies with the third exception noted above.

In all cases, you must comply with applicable law, the underlying license, and this Trademark Policy, as amended from time to time. For instance, any mention of Percona trademarks should include the full trademarked name, with proper spelling and capitalization, along with attribution of ownership to Percona Inc. For example, the full proper name for XtraBackup is Percona XtraBackup. However, it is acceptable to omit the word "Percona" for brevity on the second and subsequent uses, where such omission does not cause confusion.

In the event of doubt as to any of the conditions or exceptions outlined in this Trademark Policy, please contact trademarks@percona.com for assistance and we will do our very best to be helpful.

13.1 Get expert help

If you need assistance, visit the community forum for comprehensive and free database knowledge, or contact our Percona Database Experts for professional support and services.

 [Community Forum](#)  [Get a Percona Expert](#)

Last update: June 28, 2023
Created: November 25, 2021