

# **MySQL Enterprise Backup User's Guide (Version 4.0.0)**

---

## Abstract

This is the user manual for MySQL Enterprise Backup, a commercially licensed backup utility for MySQL databases. It explains the different kinds of backup and restore that can be performed with MySQL Enterprise Backup, and describes the commands for performing them. Strategies for optimizing backup and restore operations are also discussed.

For notes detailing the changes in each release, see the [MySQL Enterprise Backup 4.0 Release Notes](#).

For legal information, see the [Legal Notices](#).

For help with using MySQL, please visit either the [MySQL Forums](#) or [MySQL Mailing Lists](#), where you can discuss your issues with other MySQL users.

For additional documentation on MySQL products, including translations of the documentation into other languages, and downloadable versions in variety of formats, including HTML and PDF formats, see the [MySQL Documentation Library](#).

Document generated on: 2015-12-12 (revision: 6215)

---

---

# Table of Contents

Preface and Legal Notices .....	xi
I Getting Started with MySQL Enterprise Backup .....	1
1 Introduction to MySQL Enterprise Backup .....	5
1.1 The mysqlbackup Client .....	5
1.2 Overview of Backup Types .....	5
1.3 Files that Are Backed Up .....	7
1.3.1 Types of Files Contained in a Backup .....	7
1.3.2 Files Backed up for InnoDB Data .....	16
1.3.3 Files Backed up for Data Stored with MyISAM and Other Storage Engines .....	17
1.3.4 Files Generated by mysqlbackup .....	17
2 Installing MySQL Enterprise Backup .....	19
3 What's New in MySQL Enterprise Backup 4.0? .....	21
II Using MySQL Enterprise Backup .....	23
4 Backing Up a Database Server .....	27
4.1 Before the First Backup .....	27
4.1.1 Collect Database Information .....	27
4.1.2 Grant MySQL Privileges to Backup Administrator .....	29
4.1.3 Designate a Location for the Backup Directory .....	30
4.2 The Typical Backup / Verify / Restore Cycle .....	30
4.2.1 Backing Up an Entire MySQL Instance .....	30
4.2.2 Verifying a Backup .....	32
4.2.3 Restoring a Database .....	33
4.3 Backup Scenarios and Examples .....	36
4.3.1 Making a Single-File Backup .....	36
4.3.2 Making a Full Backup .....	40
4.3.3 Making a Differential or Incremental Backup .....	41
4.3.4 Making a Compressed Backup .....	45
4.3.5 Making a Partial Backup .....	46
4.3.6 Making an Optimistic Backup .....	50
4.3.7 Making a Back Up of In-Memory Database Data .....	52
4.3.8 Making Scheduled Backups .....	52
4.4 Making Backups with a Distributed File System (DFS) or Storage Access Network (SAN) .....	53
5 Recovering or Restoring a Database .....	55
5.1 Performing a Restore Operation .....	55
5.1.1 Restoring a Compressed Backup .....	56
5.1.2 Restoring an Encrypted Backup Image .....	56
5.1.3 Restoring an Incremental Backup .....	57
5.1.4 Restoring Backups Created with the <code>--use-tts</code> Option .....	58
5.1.5 Restoring a Backup from Cloud Storage to a MySQL Server .....	58
5.1.6 Advanced: Restoring a Directory Backup .....	59
5.2 Point-in-Time Recovery from a Hot Backup .....	60
5.3 Restoring a Backup with a Database Upgrade or Downgrade .....	61
6 Using MySQL Enterprise Backup with Replication .....	63
6.1 Setting Up a New Replication Slave .....	63
6.2 Backing up and Restoring a Slave Database .....	64
6.3 Restoring a Master Database .....	65
7 Performance Considerations for MySQL Enterprise Backup .....	67
7.1 Optimizing Backup Performance .....	67
7.2 Optimizing Restore Performance .....	70
8 Encryption for Backups .....	73

9 Using MySQL Enterprise Backup with Media Management Software (MMS) Products .....	75
9.1 Backing Up to Tape with Oracle Secure Backup .....	75
10 Troubleshooting for MySQL Enterprise Backup .....	79
10.1 Monitoring Backups with MySQL Enterprise Monitor .....	79
10.2 Error codes of MySQL Enterprise Backup .....	79
10.3 Working Around Corruption Problems .....	80
10.4 Using the MySQL Enterprise Backup Logs .....	80
10.5 Using the MySQL Enterprise Backup Manifest .....	82
III <code>mysqlbackup</code> Command Reference .....	85
11 <code>mysqlbackup</code> .....	89
12 <code>mysqlbackup</code> commands .....	91
12.1 Backup Operations .....	91
12.2 Update Operations .....	92
12.3 Restore Operations .....	93
12.4 Validation Operations .....	95
12.5 Other Single-File Backup Operations .....	96
13 <code>mysqlbackup</code> Command-Line Options .....	99
13.1 Standard Options .....	104
13.2 Connection Options .....	106
13.3 Server Repository Options .....	107
13.4 Backup Repository Options .....	109
13.5 Metadata Options .....	113
13.6 Compression Options .....	114
13.7 Incremental Backup Options .....	115
13.8 Partial Backup and Restore Options .....	117
13.9 Single-File Backup Options .....	123
13.10 Performance / Scalability / Capacity Options .....	125
13.11 Message Logging Options .....	132
13.12 Progress Report Options .....	133
13.13 Encryption Options .....	136
13.14 Cloud Storage Options .....	137
13.15 Options for Special Backup Types .....	139
14 Configuration Files and Parameters .....	141
IV Appendixes .....	143
A Frequently Asked Questions for MySQL Enterprise Backup .....	147
B MySQL Enterprise Backup Limitations .....	149
B.1 Limitations of MySQL Enterprise Backup .....	149
C Compatibility Information for MySQL Enterprise Backup .....	153
C.1 Cross-Platform Compatibility .....	153
C.2 Compatibility with MySQL Versions .....	153
C.3 Compatibility with Older MySQL Enterprise Backup .....	153
D Licenses for Third-Party Components .....	155
D.1 cURL (libcurl) License .....	155
D.2 GNU General Public License Version 3.0, 29 June 2007 and GCC Runtime Library Exception Version 3.1, 31 March 2009 .....	155
D.3 GNU Standard C++ Library (libstdc++) License .....	167
D.4 Google Controlling Master Thread I/O Rate Patch License .....	168
D.5 Google SMP Patch License .....	169
D.6 LZ4 License .....	170
D.7 OpenSSL v1.0 License .....	170
D.8 Percona Multiple I/O Threads Patch License .....	172
D.9 RE2 License .....	172
D.10 RegEX-Spencer Library License .....	173
D.11 RFC 3174 - US Secure Hash Algorithm 1 (SHA1) License .....	173

D.12 zlib License .....	174
MySQL Enterprise Backup Glossary .....	175
Index .....	189



---

## List of Tables

1.1 Types of Files in a Backup .....	7
4.1 Information Needed to Back Up a Database .....	27
13.1 List of All Options .....	99





---

## List of Examples

4.1 Single-File Backup to Absolute Path .....	36
4.2 Single-File Backup to Relative Path .....	36
4.3 Single-File Backup to Standard Output .....	36
4.4 Convert Existing Backup Directory to Single Image .....	37
4.5 Extract Existing Image to Backup Directory .....	37
4.6 List Single-File Backup Contents .....	37
4.7 Validate a Single-File Backup .....	37
4.8 Extract Single-File Backup into Current Directory .....	37
4.9 Extract Single-File Backup into a Backup Directory .....	37
4.10 Selective Extract of Single File .....	37
4.11 Selective Extract of Single Directory .....	38
4.12 Dealing with Absolute Path Names .....	38
4.13 Single-File Backup to a Remote Host .....	38
4.14 Single-file Backup to a Remote MySQL Server .....	39
4.15 Stream a Backup Directory to a Remote MySQL Server .....	39
4.16 Creating a Cloud Backup on Amazon S3 .....	40
4.17 Creating a Cloud Backup on an OpenStack Object Storage .....	40
4.18 Extract an Existing Image from Amazon S3 Cloud Storage to Backup Directory .....	40
4.19 Making an Uncompressed Partial Backup of InnoDB Tables .....	49
4.20 Making a Compressed Partial Backup .....	50
4.21 Optimistic Backup Using the Option <code>optimistic-time=YYMMDDHHMMSS</code> .....	51
4.22 Optimistic Backup Using the Option <code>optimistic-time=now</code> .....	51
4.23 Optimistic Backup Using the <code>optimistic-busy-tables</code> Option .....	52
4.24 Optimistic and Partial Backup Using both the <code>optimistic-busy-tables</code> and <code>optimistic-time</code> Options .....	52
5.1 Restoring a Database .....	55
5.2 Restoring a Compressed Backup .....	56
5.3 Restoring an Encrypted Backup Image .....	56
5.4 Restoring an Incremental Backup Image .....	57
5.5 Restoring Selected Tables from a TTS Image Backup .....	58
5.6 Restoring and Renaming a Table from a TTS Backup .....	58
5.7 Restoring a Single-file Backup from Amazon S3 to a MySQL Server .....	59
5.8 Restoring a Single-file Backup from an OpenStack Object Storage to a MySQL Server .....	59
5.9 Restoring a Backup Directory using <code>copy-back-and-apply-log</code> .....	59
5.10 Applying the Log to a Backup .....	59
5.11 Applying the Log to a Compressed Backup .....	59
9.1 Sample <code>mysqlbackup</code> Commands Using MySQL Enterprise Backup with Oracle Secure Backup .....	77
12.1 Apply Log to Full Backup .....	93
14.1 Example <code>backup-my.cnf</code> file .....	141



---

# Preface and Legal Notices

This is the user manual for the MySQL Enterprise Backup product.

For license information, see the [Legal Notices](#). This product may contain third-party code. For license information on third-party code, see [Appendix D, Licenses for Third-Party Components](#).

## Legal Notices

Copyright © 2003, 2015, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

#### Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

#### Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

This documentation is NOT distributed under a GPL license. Use of this documentation is subject to the following terms:

You may create a printed copy of this documentation solely for your own personal use. Conversion to other formats is allowed as long as the actual content is not altered or edited in any way. You shall not publish or distribute this documentation in any form or on any media, except if you distribute the documentation in a manner similar to how Oracle disseminates it (that is, electronically for download on a Web site with the software) or on a CD-ROM or similar medium, provided however that the documentation is disseminated together with the software on the same medium. Any other use, such as any dissemination of printed copies or use of this documentation, in whole or in part, in another publication, requires the prior written consent from an authorized representative of Oracle. Oracle and/or its affiliates reserve any and all rights to this documentation not expressly granted above.

---

# **Part I Getting Started with MySQL Enterprise Backup**

---

---

---

---

## Table of Contents

1	Introduction to MySQL Enterprise Backup .....	5
1.1	The mysqlbackup Client .....	5
1.2	Overview of Backup Types .....	5
1.3	Files that Are Backed Up .....	7
1.3.1	Types of Files Contained in a Backup .....	7
1.3.2	Files Backed up for InnoDB Data .....	16
1.3.3	Files Backed up for Data Stored with MyISAM and Other Storage Engines .....	17
1.3.4	Files Generated by mysqlbackup .....	17
2	Installing MySQL Enterprise Backup .....	19
3	What's New in MySQL Enterprise Backup 4.0? .....	21





---

# Chapter 1 Introduction to MySQL Enterprise Backup

## Table of Contents

1.1 The <code>mysqlbackup</code> Client .....	5
1.2 Overview of Backup Types .....	5
1.3 Files that Are Backed Up .....	7
1.3.1 Types of Files Contained in a Backup .....	7
1.3.2 Files Backed up for InnoDB Data .....	16
1.3.3 Files Backed up for Data Stored with MyISAM and Other Storage Engines .....	17
1.3.4 Files Generated by <code>mysqlbackup</code> .....	17

MySQL Enterprise Backup is a backup utility for MySQL servers. It is a multi-platform, high-performance tool, offering rich features like “hot” (online) backup, incremental and differential backup, selective backup and restore, support for direct cloud storage backup, backup encryption and compression, and many other valuable features.

While optimized for use with InnoDB tables, MySQL Enterprise Backup is capable of backing up and restoring all kinds of tables created by any kinds of storage engines supported by MySQL. The parallelism of its read and write processes (performed in independent, multiple threads) and its block-level parallelism (different threads can read, process, or write different chunks within a single file) allow backup and restore processes to be completed with great speed, and often with a significant performance gain when compared to a [logical backup](#) using tools like `mysqldump`.

MySQL Enterprise Backup is a valuable tool for maintaining and safeguarding your MySQL data, and for quick and reliable recovery when accidents or disasters strike. It is part of the MySQL Enterprise Edition, available to subscribers under a commercial license.

Among other things, this manual explains:

- How to install MySQL Enterprise Backup.
- The different kinds of backups that can be performed with MySQL Enterprise Backup, how to perform them, and some tips on choosing the right kind of backups for your system.
- How to restore backups created by MySQL Enterprise Backup.
- How to use MySQL Enterprise Backup in special situations or for special purposes (for example, setting up replication, using Media Management Software (MMS) products, or using a Distributed File System (DFS)).
- The `mysqlbackup` client, its commands and command options.

## 1.1 The `mysqlbackup` Client

All MySQL Enterprise Backup functions are executed with the `mysqlbackup` client. It is used for performing different types of backup and restore operations, as well as other related tasks like backup compression, decompression, validation, and so on.

Use of the `mysqlbackup` client and the related commands are explained and illustrated throughout this manual. For detailed information on the `mysqlbackup` commands and command options, see [Part III, “mysqlbackup Command Reference”](#).

## 1.2 Overview of Backup Types

When it comes to formulating your backup strategy, performance and storage space are the key considerations. You want the backup to complete quickly, with as little CPU overhead on the database server as possible. You also want the backup data to be compact, so you can keep multiple backups on hand to restore at a moment's notice. Transferring the backup data to a different system should be quick and convenient. Under such considerations, various strategies for backing up your database often give you different advantages, for the different trade-offs you make when choosing a particular strategy. To choose the strategy that best fits your needs, you have to understand the nature of each kind of backups that MySQL Enterprise Backup can perform, for which this section is giving a brief overview.

## Kinds of backups according to the level of service disruption

Depending on how the database operations would be disrupted during a backup, the backup is classified as “hot,” “warm,” or “cold”:

- *Very Low to Low Level of Disruption:* A [hot backup](#) is a backup performed while the database is running. This type of backups does not block normal database operations. It captures even changes that occur while the backup is happening. Comparing to the other backup types, it causes the least disruption to the database server, and it is a desirable backup option when you want to avoid taking your application, web site, or web service offline. However, before a hot backup can be restored, there needs to be an extra process of preparing the backup to make it consistent (i.e., correctly reflecting the state of the state of the database at the time the backup was completed). See [Section 5.1.6, “Advanced: Restoring a Directory Backup”](#) for more explanations.

When connected to a running MySQL server, MySQL Enterprise Backup performs hot backup for InnoDB tables.

- *Medium to High Level of Disruption:* A [warm backup](#) is a backup performed with the database put under a read-only state. This type of backups blocks any write operations to the tables during the backup process, but still allow tables to be read.

When connected to a running MySQL server, MySQL Enterprise Backup backs up all MyISAM and other non-InnoDB tables using the [warm backup](#) technique after all InnoDB tables have already been backed up with the [hot backup](#) method. Therefore, to back up as much data as possible during the hot backup phase, you should designate InnoDB as the default storage engine for new tables (which has been the default setting since MySQL 5.5), or convert existing tables to use the InnoDB storage engine.

- *High to Very High Level of Disruption:* A [cold backup](#) is a backup created while the database is stopped. Not only is it most disruptive to the database service, it also renders a lot of MySQL Enterprise Backup features unusable (for example, the ability to retrieve information about the structure of the database automatically through a database connection, so they do not have to be supplied to MySQL Enterprise Backup by a configuration file or command-line options).

To avoid service disruption, you would typically perform a cold backup on a replication slave, which can be stopped without taking down the entire application or web site.

## Kinds of backups according to whether all data, or recent changes only are backed up

According to whether you want to include all data into your backup or only the recent changes, and according to recent changes since when, you can perform either a full backup, a differential backup, or an incremental backup. The three types of backups have different levels of requirements for CPU overhead and disk space, thus are suitable for different situations:

- A [full backup](#) includes the complete data from the database (except in cases where some tables are excluded with the [partial backup options](#)).

- A [differential backup](#) includes all changes to the data since the last full backup. It is faster than a full backup, saves storage space on the database server, and saves on network traffic when the backup is being transferred to a different server. However, it requires additional processing to make the backup ready for restore, which you can perform on a different system to minimize CPU overhead on the database server.
- An [Incremental backup](#) includes all changes to the data since the last backup. It offers similar advantages over a full backup as a differential backup does, and often to a even greater extent by further decreasing the backup size. But it might also require more preparations on a longer series of backups, before a restore can be performed.

## Compressed versus uncompressed backups

Backup compression saves you storage space and network traffic to transfer the backup data onto a different server. Compression does add some CPU overhead, but the overhead is algorithm dependent and it is fairly low for the default algorithm used by MySQL Enterprise Backup. Also, compression often reduces the IO overhead by a great deal, which might shorten the restore time, especially for slower IO devices. However, during the restore process, you need time for decompression and also storage space for both the compressed and decompressed data at the same time. So, take into account the additional storage space and the extra time needed during a restore when considering whether to create compressed backups.

When streaming backup data to another server, you might want to compress the backup either on the original server or the destination server, depending on which server has more spare CPU capacity and how much network traffic the compression could save.

For more on techniques and tradeoffs involving backup and restore performance, see [Chapter 7, Performance Considerations for MySQL Enterprise Backup](#).

## 1.3 Files that Are Backed Up

This section explains the various types of files contained in a backup.

### 1.3.1 Types of Files Contained in a Backup

The following table shows the different types of files that are included in a single-file backup image or a directory backup. In the case of a single-file backup, unpack the file into a [backup directory](#) structure using the [extract](#) or the [image-to-backup-dir](#) command to view the files.


**Table 1.1 Types of Files in a Backup**

File Name, Pattern, or Extension	Relation to Original Data Files	Notes
<code>ibdata*</code>	The InnoDB system tablespace, containing multiple InnoDB tables and associated indexes.	Because the original files might change while the backup is in progress, the <a href="#">apply-log</a> step applies the same changes to the corresponding backup files.
<code>*.ibd</code>	An InnoDB tablespace, which can be (a) a <a href="#">file-per-table tablespace</a> , containing a single InnoDB table and associated indexes, or (b) a file-per-table <a href="#">external tablespace</a> located outside of the server's data directory, containing a single InnoDB table	Because the original files might change while the backup is in progress, the apply-log step applies the same changes to the corresponding backup files.

File Name, Pattern, or Extension	Relation to Original Data Files	Notes
	and associated indexes, or (c) a <a href="#">general tablespace</a> , containing one or more tables and their indexes.	
*.ibz	Compressed form of InnoDB data files from the MySQL data directory.	Produced instead of <a href="#">.ibd</a> files in a compressed backup. The <a href="#">ibdata*</a> files representing the InnoDB system tablespace also receive this extension in a compressed backup.  The <a href="#">.ibz</a> files are uncompressed during the <a href="#">apply-log</a> or <a href="#">copy-back-and-apply-log</a> step.
*.frm	Hold metadata about all MySQL tables.	The database is put into a read-only state while these files are copied. These files are copied unmodified.
*.MYD	MyISAM table data.	The database is put into a read-only state while these files are copied. These files are copied unmodified.
*.MYI	MyISAM index data.	The database is put into a read-only state while these files are copied. These files are copied unmodified.
*.CSM	Metadata for CSV tables.	These files are copied unmodified. The <a href="#">backup_history</a> and <a href="#">backup_progress</a> tables created by <a href="#">mysqlbackup</a> use the CSV format, so the backup always includes some files with this extension.
*.CSV	Data for CSV tables.	These files are copied unmodified. The <a href="#">backup_history</a> and <a href="#">backup_progress</a> tables created by <a href="#">mysqlbackup</a> use the CSV format, so the backup always includes some files with this extension.
*.MRG	MERGE storage engine references to other tables.	The database is put into a read-only state while these files are copied. These files are copied unmodified.
*.TRG	Trigger parameters.	The database is put into a read-only state while these files are copied. These files are copied unmodified.

File Name, Pattern, or Extension	Relation to Original Data Files	Notes
*.TRN	Trigger namespace information.	The database is put into a read-only state while these files are copied. These files are copied unmodified.
*.opt	Database configuration information.	The database is put into a read-only state while these files are copied. These files are copied unmodified.
*.par	Definitions for partitioned tables.	The database is put into a read-only state while these files are copied. These files are copied unmodified.
*.ARM	Archive storage engine metadata.	The database is put into a read-only state while these files are copied. These files are copied unmodified.
*.ARZ	Archive storage engine data.	The database is put into a read-only state while these files are copied. These files are copied unmodified.
backup-my.cnf	Records the configuration parameters that specify the layout of the MySQL data files.	Used in restore operations to reproduce the same layout as when the backup was taken.
ibbackup_ibd_files	Records names of the .ibd files and their space IDs during an incremental backup.	This file is created during an incremental backup. During a restore, the information in the file is used to delete the tables from the full backup that has been removed between the time of the full backup and the time of the incremental backup.
ibbackup_logfile	A condensed version of the <code>ib_logfile*</code> files from the MySQL data directory.	The InnoDB log files ( <code>ib_logfile*</code> ) are fixed-size files that are continuously updated during the database's operation. For backup purposes, only the changes that are committed while the backup is in progress are needed. These changes are recorded in <code>ibbackup_logfile</code> , and used to re-create the <code>ib_logfile*</code> files during the apply-log phase.
ibbackup_redo_log_only	Created instead of the <code>ibbackup_logfile</code> for incremental backups taken with the <code>--incremental-with-redo-log-only</code> option.	

File Name, Pattern, or Extension	Relation to Original Data Files	Notes
<code>ib_logfile*</code>	Created in the backup directory by <code>mysqlbackup</code> during the <code>apply-log</code> phase after the initial backup.	These files are not copied from the original data directory, but rather re-created in the backup directory during the <code>apply-log</code> phase after the initial backup, using the changes recorded in the <code>ibbackup_logfile</code> file.
<code>*.bl</code>	Renamed version of each <code>.isl</code> file from the backed-up server.	<p>A <code>.isl</code> file is created when you specify the location of an InnoDB table using the syntax <code>CREATE TABLE ... DATA DIRECTORY = ...</code> (see <a href="#">Creating a File-Per-Table Tablespace Outside the Data Directory</a> for details) or when a table is put into a <a href="#">general tablespace</a> (only for MySQL 5.7 and after; see <a href="#">general tablespace</a> for details). The <code>.isl</code> file acts like a symbolic link pointing to the tablespace file. The <code>.bl</code> files might or might not be turned back into <code>.isl</code> files during the <code>copy-back</code> or <code>copy-back-and-apply-log</code> operation. If the specified directory does not exist on the server where the backup is restored, <code>mysqlbackup</code> attempts to create it. If the directory cannot be created, the restore operation fails. Thus, if you would like to use the <code>DATA DIRECTORY</code> clause to put tables at different locations or to restore to a server with a different file structure where the corresponding directories cannot be created, edit the <code>.bl</code> files before restoring to point to directories that do exist on the destination server.</p> <p>When restoring a backup created with <a href="#">transportable tablespaces (TTS)</a>, if the directory on the target server pointed to by a <code>.bl</code> file already contains <code>.ibd</code> files, the <code>--force [105]</code> option is required when you restore the backup.</p>
Timestamped directory, such as <code>2011-05-26_13-42-02</code>	Created by the <code>--with-timestamp</code> option. All the backup files go inside this subdirectory.	Use the <code>--with-timestamp</code> option to easily keep more than one set of backup data under the same main backup directory.

File Name, Pattern, or Extension	Relation to Original Data Files	Notes
<code>datadir</code> directory	A subdirectory that stores the data files and database subdirectories from the original MySQL instance.	Created under the backup directory by <code>mysqlbackup</code> .
binary log files	Binary log files from the server, which are included in a backup by default (except when the backup is created with the <code>--use-tts</code> option). They allow a snapshot of the server to be taken, so a server can be cloned to its exact state. Using a full backup as a basis, the binary log files that are included with an incremental backup can be used for a point-in-time recovery (PITR), which restores a database to its state at a certain point in time after the last full backup. See <a href="#">Section 5.2, “Point-in-Time Recovery from a Hot Backup”</a> for details.	<p>Saved under the <code>datadir</code> directory under the backup directory. Use the <code>--skip-binlog</code> option to exclude binary logs in the backup. For MySQL 5.5 and earlier, as well as all offline backups, use the <code>--log-bin-index</code> option to specify the absolute path of the index file on the MySQL server that lists all the used binary log files, if it is different from the default value of the option, for <code>mysqlbackup</code> to find the binary log files and include them in the backups. The index file itself, with the locations of the binary log files properly updated to point to the files' locations in the backup directory, is included into the backup under the <code>datadir</code> directory.</p> <p>The binary log files are compressed and saved with the <code>.bz</code> extension when being included in a compressed backup.</p> <p>The binary log files and the index file, when included in a backup, are always copied into the restored server's data directory during a restore operation. Use the <code>--skip-binlog</code> option to skip the restoring of the binary log.</p> <div>  <div> <b>Note</b> <p>If any binary log files are missing on the server you are backing up, you should use the <code>--skip-binlog</code> option</p> </div> </div>





File Name, Pattern, or Extension	Relation to Original Data Files	Notes
		to avoid <code>mysqlbackup</code> throwing an error for the missing files.
relay log files	Relay log files from a slave server, which are included in a backup of a slave server by default (except when the backup is created with the <code>--use-tts</code> option). Their inclusion saves the time and resources required for fetching the relay logs from the master when the slave is being restored.	<p>Saved under the <code>datadir</code> directory under the backup directory. Use the <code>--skip-relaylog</code> option to exclude relay logs in the backup. For offline backup, use the <code>--relay-log-index</code> option to specify the absolute path of the index file on the MySQL server that lists all the used relay log files, if it is different from the default value of the option, for <code>mysqlbackup</code> to find the relay log files and include them in the backups. The index file itself, with the locations of the relay log files properly updated to point to the files' locations in the backup directory, is included into the backup under the <code>datadir</code> directory.</p> <p>The relay log files are compressed and saved with the <code>.bz</code> extension when being included in a compressed backup.</p> <p>The relay log files and the index file, when included in a backup, are always copied into the restored server's data directory during a restore operation. Use the <code>--skip-relaylog</code> option to skip the restoring of the relay log.</p>
*.bz	Compressed binary log or relay log files.	The binary log and relay log files are compressed and saved with the <code>.bz</code> extension when being included in a compressed backup. They are decompressed during a restore.
slave status log files	Usually named <code>master.info</code> and <code>relay-log.info</code> , they are included by default in a backup of a slave database in a replication setup. See <a href="#">Slave Status Logs</a> , for details.	Saved under the <code>datadir</code> directory under the backup directory. For an offline backup, use the <code>--master-info-file</code> and <code>--relaylog-info-file</code> options to specify the absolute



File Name, Pattern, or Extension	Relation to Original Data Files	Notes
		<p>paths of the information files, if they are different from the default values of the options, for <code>mysqlbackup</code> to find those files and include them in the backups.</p> <p>The copying of these files are skipped during a backup or a restore when the <code>--skip-relay-log</code> option is used.</p>
Backup image file	A single-file backup produced by the <code>backup-to-image</code> option, with a name specified by the <code>--backup-image</code> option.	You can move the image file without losing or damaging the contents inside it, then unpack it with <code>mysqlbackup</code> using the <code>extract</code> command and specifying the same image name with the <code>--backup-image</code> option. Although some extra files such as <code>backup-my.cnf</code> and the <code>meta</code> subdirectory are present in the backup directory, these files are also included in the image file and do not need to be moved along with it.
Any other files in the database subdirectories under the <code>datadir</code> directory (that is, under <code>backup-dir/datadir/data-base-name</code> )	Copied from the database subdirectories under the MySQL data directory.	By default, any unrecognized files in the database subdirectories under the MySQL data directory are copied to the backup. To omit such files, specify the <code>--only-known-file-types</code> option.
<code>meta</code> directory	A subdirectory that stores files with metadata about the backup.	Created under the backup directory by <code>mysqlbackup</code> . All files listed below go inside the <code>meta</code> subdirectory.
<code>backup_variables.txt</code>	Holds important information about the backup. For use by <code>mysqlbackup</code> only.	<code>mysqlbackup</code> consults and possibly updates this file during operations after the initial backup, such as the apply-log phase or the restore phase.
<code>image_files.xml</code>	Contains the list of all the files (except itself) that are present in the single-file backup produced by the <code>backup-to-image</code> or <code>backup-dir-to-image</code> options. For details about this file, see <a href="#">Section 10.5, “Using the MySQL Enterprise Backup Manifest”</a> .	This file is not modified at any stage once generated.
<code>backup_create.xml</code>	Lists the command line arguments and environment in which the backup was created. For details	This file is not modified once it is created. You can prevent this file from being generated

File Name, Pattern, or Extension	Relation to Original Data Files	Notes
	about this file, see <a href="#">Section 10.5, “Using the MySQL Enterprise Backup Manifest”</a> .	by specifying the <code>--disable-manifest</code> option.
<code>backup_content.xml</code>	Essential metadata for the files and database definitions of the backup data. It also contains details of all the plugins defined on the backed-up server, by which users should make sure the same plugins are defined in the same manner on the target server for restoration. For details about this file, see <a href="#">Section 10.5, “Using the MySQL Enterprise Backup Manifest”</a> .	This file is not modified once created. You can prevent this file from being generated by specifying the <code>--disable-manifest</code> option.
<code>comments.txt</code>	Produced by the <code>--comments</code> or <code>--comments-file</code> option.	The comments are specified by you to document the purpose or special considerations for this backup job.
<code>gtid_executed.sql</code>	Signifies the backup came from a server with GTIDs enabled.	GTIDs are a replication feature in MySQL 5.6 and higher. See <a href="#">Replication with Global Transaction Identifiers</a> for details. When you back up a server with GTIDs enabled using <code>mysqlbackup</code> , the file named <code>backup_gtid_executed.sql</code> is created in the <code>meta</code> folder under the backup directory. Edit and execute this file after restoring the backup data on a slave server; see <a href="#">Section 6.1, “Setting Up a New Replication Slave”</a> for details.
<code>server-my.cnf</code>	Contains values of the backed-up server's global variables that are set to non-default values. Use this file or <code>server-all.cnf</code> to start the target server for restoration.	During a <code>copy-back</code> or <code>copy-back-and-apply-log</code> operation, the <a href="#">server repository options</a> values (e.g., <code>--datadir</code> , <code>--innodb_data_home_dir</code> , etc.) in the file are modified if the command makes changes to them through the command options. However, during an <code>apply-incremental-backup</code> operation, the values already saved in the file take precedence and they are not modified by the option values supplied through the command.

File Name, Pattern, or Extension	Relation to Original Data Files	Notes
		 <div> <b>Warning</b>  When using the file to restart the target server, change parameters like <code>--tmpdir</code>, <code>--general-log</code>, etc., and any global variable that uses an absolute filepath to avoid the accidental usage of the wrong file locations by the target server. </div>
<code>server-all.cnf</code>	Contains values of all the global variables of the backed-up server. Use this file or <code>server-my.cnf</code> to start the target server for restoration.	<p>During a <code>copy-back</code> or <code>copy-back-and-apply-log</code> operation, the <code>server repository options</code> values (e.g., <code>--datadir</code>, <code>--innodb_data_home_dir</code>, etc.) in the file are modified if the command makes changes to them through the command options. However, during an <code>apply-incremental-backup</code> operation, the values already saved in the file take precedence and they are not modified by the option values supplied through the command.</p> <div>  <b>Warning</b>  When using the file </div>

File Name, Pattern, or Extension	Relation to Original Data Files	Notes
		to restart the target server, change parameters like <code>--tmpdir</code> , <code>--general-log</code> , etc., and any global variable that uses an absolute filepath to avoid the accidental usage of the wrong file locations by the target server.
<code>ib_buffer_pool</code>	<p>The file produced on the server when <code>innodb_buffer_pool_dump_at_shutdown</code> (enabled by default on MySQL 5.7.7 and after) or <code>innodb_buffer_pool_dump_now</code> is enabled. It holds the list of tablespace IDs and page IDs of the server's <a href="#">buffer pool</a>.</p> <p>The actual file name might be different, as it can be configured by the server's system variable <code>innodb_buffer_pool_filename</code>.</p>	<p>With the default setting on MySQL server 5.7.7 and after <code>innodb_buffer_pool_load_at_startup</code> the target server, during start up, is going to restore the buffer pool state of the backed up server using this file. See <a href="#">Preloading the InnoDB Buffer Pool for Faster Restart</a> for details.</p>

### 1.3.2 Files Backed up for InnoDB Data

The InnoDB-related data files that are backed up include the [ibdata\\* files](#) (which represent the [system tablespace](#) and possibly the data for some user tables), any [.ibd files](#) (which contains data from user tables created with the [file-per-table](#) setting enabled), and the data extracted from the [ib\\_logfile\\* files](#) (the [redo log](#) information representing changes that occur while the backup is running), which is stored in a new backup file [ibbackup\\_logfile](#).

If you use the compressed backup feature, the [.ibd files](#) are renamed in their compressed form to [.ibz files](#).

The files, as they are originally copied, form a [raw backup](#) that requires further processing before it is ready to be restored. You then run the [apply](#) step (either as part of a [copy-back-and-apply-log](#) command or a [backup-and-apply-log](#) command, or as a separate [apply-log](#) command), which updates the backup files based on the changes recorded in the [ibbackup\\_logfile](#) file, producing a [prepared backup](#). At this point, the backup data corresponds to a single point in time. The files are now ready to be restored, or for some other uses such as testing, reporting, or deployment as a replication slave.

To restore InnoDB tables to their original state, you must also have the corresponding [.frm files](#) along with the backup data. Otherwise, the table definitions could be missing or outdated if someone has run [ALTER TABLE](#) or [DROP TABLE](#) statements since the backup. By default, [mysqlbackup](#) automatically copies the [.frm](#) files during a backup operation and restores the files during a restore operation.

**Note**

To avoid concurrency issues during backups of busy databases, you can use the [--only-innodb](#) or [--only-innodb-with-frm](#) option to back up only InnoDB tables and the associated data.

### 1.3.3 Files Backed up for Data Stored with MyISAM and Other Storage Engines

[mysqlbackup](#) also backs up the [.MYD files](#), [.MYI files](#), and the [.frm](#) files associated with the MyISAM tables. Files with other extensions that are backed up are shown in [Table 1.1, “Types of Files in a Backup”](#).

**Note**

While MySQL Enterprise Backup can back up non-InnoDB data (like MYISAM tables), the MySQL server to be backed up must support InnoDB (i.e., the backup process will fail if the server was started up with the [--innodb=OFF](#) or [--skip-innodb](#) option), and the server must contain at least one InnoDB table.

MyISAM tables and these other types of files cannot be backed up in the same non-blocking way as InnoDB tables can be. They can be backed up using the [warm backup](#) technique: changes to these tables are prevented while they are being backed up, possibly making the database unresponsive for a time, but no shutdown is required during the backup.

### 1.3.4 Files Generated by [mysqlbackup](#)

Inside the image backup file created by the [backup-to-image](#) command of [mysqlbackup](#) are some new files that are produced during the backup process. These files are used to control later tasks such as verifying and restoring the backup data. The files generated during the backup process include:

- [backup-my.cnf](#): Records the crucial configuration parameters that apply to the backup. These parameter values are used during a restore operation, so that the original values are used regardless of changes to your server's [my.cnf](#) file since the backup.
- [meta/backup\\_create.xml](#): Lists the command line arguments and environment in which the backup was created.
- [meta/backup\\_content.xml](#): Essential metadata for the files and database definitions of the backup data.
- [server-my.cnf](#): Contains values of the backed-up server's global variables that are set to non-default values.
- [server-all.cnf](#): Contains values of all the global variables of the backed-up server.

For details about these and other files contained in the backup, see [Table 1.1, “Types of Files in a Backup”](#).



---

## Chapter 2 Installing MySQL Enterprise Backup

Install MySQL Enterprise Backup on each database server whose contents you intend to back up. Typically, you perform all backup and restore operations locally, by running `mysqlbackup` on the same server as the MySQL instance.

MySQL Enterprise Backup is packaged as either an archive file (`.tgz`, archived with `tar` and compressed with `gzip`) or as a platform-specific installer.

### Installing on Unix and Linux Systems

For all Linux and Unix systems, the product is available as a `.tgz` file. Unpack this file as follows:

```
tar xvzf package.tgz
```

`mysqlbackup` is unpacked into a subdirectory. You can either copy them into a system directory (preserving their execute permission bits), or add to your `$PATH` setting the directory where you unpacked it.



#### Note

On FreeBSD 10.1, there is a dependency for `mysqlbackup` on the shared library `libz.so.5`, which can be installed on your system by installing, for example, the `compat8x` libraries.

For certain Linux distributions, the product is also available as an RPM archive. When you install the RPM using the command `sudo rpm -i package_name.rpm`, the `mysqlbackup` client is installed in the directory `/opt/mysql/meb-4.0`. You must add this directory to your `$PATH` setting.

Installation packages for Debian and Ubuntu platforms are also available. Install the package with the following command `sudo dpkg -i package_name.deb`.

### Installing on Windows Systems

The product can be installed together with other MySQL products with the MySQL Installer for Windows. It can also be installed separately with either an individual `.msi` installer or `.zip` file.

When installing with a `.msi` installer, specify the installation location, preferably under the same directory where other MySQL products have been installed. Choose the option **Include directory in Windows PATH**, so that you can run `mysqlbackup` from any directory.

When installing with a `.zip` file, simply unzip the file and put `mysqlbackup.exe` at the desired installation location. You can add that location to the `%PATH%` variable, so that you can run the `mysqlbackup` client from any directory.

Verify the installation by selecting the menu item **Start > Programs > MySQL Enterprise Backup 4.0 > MySQL Enterprise Backup Command Line**. The menu item displays version information and opens a command prompt for running the `mysqlbackup` command.





---

## Chapter 3 What's New in MySQL Enterprise Backup 4.0?

This chapter highlights the new features in MySQL Enterprise Backup 4.0, as well as any significant changes made to MySQL Enterprise Backup with the release of this series.

- **Supports MySQL 5.7.** MySQL Enterprise Backup 4.0 supports MySQL 5.7 (starting with 5.7.9; for any earlier versions of MySQL 5.7 or 5.6, please use MySQL Enterprise Backup 3.12 instead), on which:
  - Tables in [general tablespaces](#) can be backed up and restored. MySQL Enterprise Backup treats these general table spaces similarly as it treats the tablespaces created outside of the database's data directory using the `CREATE TABLE ... DATA DIRECTORY =` statement; for details, see the description for the `.b1` files in [Table 1.1, "Types of Files in a Backup"](#).
  - InnoDB tables with [page sizes](#) of 32 and 64KB can be backed up and restored.
  - Partitioned tables can be selectively backed up or restored using the `--use-tts`, `--include-tables`, and `--exclude-tables` options.  
Some of the MySQL 5.7 features are not supported by MySQL Enterprise Backup, and some limitations apply for the supported features. See [Appendix B, MySQL Enterprise Backup Limitations](#) for details.
- **Store Tape Backup Information on the MySQL Server.** Each time an online backup is made to a tape using the System Backup to Tape (SBT) API, an entry is made to the `mysql.backup_sbt_history` table on the backed-up MySQL instance. That facilitates the management of tape backups by allowing easy look-ups for information on them. See [Section 9.1, "Backing Up to Tape with Oracle Secure Backup"](#) for details.



---

## **Part II Using MySQL Enterprise Backup**

---

---

---

---

## Table of Contents

4	Backing Up a Database Server .....	27
4.1	Before the First Backup .....	27
4.1.1	Collect Database Information .....	27
4.1.2	Grant MySQL Privileges to Backup Administrator .....	29
4.1.3	Designate a Location for the Backup Directory .....	30
4.2	The Typical Backup / Verify / Restore Cycle .....	30
4.2.1	Backing Up an Entire MySQL Instance .....	30
4.2.2	Verifying a Backup .....	32
4.2.3	Restoring a Database .....	33
4.3	Backup Scenarios and Examples .....	36
4.3.1	Making a Single-File Backup .....	36
4.3.2	Making a Full Backup .....	40
4.3.3	Making a Differential or Incremental Backup .....	41
4.3.4	Making a Compressed Backup .....	45
4.3.5	Making a Partial Backup .....	46
4.3.6	Making an Optimistic Backup .....	50
4.3.7	Making a Back Up of In-Memory Database Data .....	52
4.3.8	Making Scheduled Backups .....	52
4.4	Making Backups with a Distributed File System (DFS) or Storage Access Network (SAN) .....	53
5	Recovering or Restoring a Database .....	55
5.1	Performing a Restore Operation .....	55
5.1.1	Restoring a Compressed Backup .....	56
5.1.2	Restoring an Encrypted Backup Image .....	56
5.1.3	Restoring an Incremental Backup .....	57
5.1.4	Restoring Backups Created with the <code>--use-tts</code> Option .....	58
5.1.5	Restoring a Backup from Cloud Storage to a MySQL Server .....	58
5.1.6	Advanced: Restoring a Directory Backup .....	59
5.2	Point-in-Time Recovery from a Hot Backup .....	60
5.3	Restoring a Backup with a Database Upgrade or Downgrade .....	61
6	Using MySQL Enterprise Backup with Replication .....	63
6.1	Setting Up a New Replication Slave .....	63
6.2	Backing up and Restoring a Slave Database .....	64
6.3	Restoring a Master Database .....	65
7	Performance Considerations for MySQL Enterprise Backup .....	67
7.1	Optimizing Backup Performance .....	67
7.2	Optimizing Restore Performance .....	70
8	Encryption for Backups .....	73
9	Using MySQL Enterprise Backup with Media Management Software (MMS) Products .....	75
9.1	Backing Up to Tape with Oracle Secure Backup .....	75
10	Troubleshooting for MySQL Enterprise Backup .....	79
10.1	Monitoring Backups with MySQL Enterprise Monitor .....	79
10.2	Error codes of MySQL Enterprise Backup .....	79
10.3	Working Around Corruption Problems .....	80
10.4	Using the MySQL Enterprise Backup Logs .....	80
10.5	Using the MySQL Enterprise Backup Manifest .....	82



---

# Chapter 4 Backing Up a Database Server

## Table of Contents

4.1 Before the First Backup .....	27
4.1.1 Collect Database Information .....	27
4.1.2 Grant MySQL Privileges to Backup Administrator .....	29
4.1.3 Designate a Location for the Backup Directory .....	30
4.2 The Typical Backup / Verify / Restore Cycle .....	30
4.2.1 Backing Up an Entire MySQL Instance .....	30
4.2.2 Verifying a Backup .....	32
4.2.3 Restoring a Database .....	33
4.3 Backup Scenarios and Examples .....	36
4.3.1 Making a Single-File Backup .....	36
4.3.2 Making a Full Backup .....	40
4.3.3 Making a Differential or Incremental Backup .....	41
4.3.4 Making a Compressed Backup .....	45
4.3.5 Making a Partial Backup .....	46
4.3.6 Making an Optimistic Backup .....	50
4.3.7 Making a Back Up of In-Memory Database Data .....	52
4.3.8 Making Scheduled Backups .....	52
4.4 Making Backups with a Distributed File System (DFS) or Storage Access Network (SAN) .....	53

This section explains the preparations you need for creating backups with MySQL Enterprise Backup, the typical backup-verify-restore cycle, and the different backup scenarios for using MySQL Enterprise Backup. It also includes sample commands and outputs, showing you how to use the `mysqlbackup` client in different situations.

## 4.1 Before the First Backup

This section outlines some of the preparations needed before you can start working with MySQL Enterprise Backup.

### 4.1.1 Collect Database Information

Before backing up a particular database server for the first time, gather some information and use it to make some planning decisions, as outlined in the following table.

**Table 4.1 Information Needed to Back Up a Database**

Information to Gather	Where to Find It	How to Use It
Path to MySQL configuration file	Default system locations, hardcoded application default locations, or from the <code>--defaults-file</code> option in the <code>mysqld</code> startup script.	The preferred way to convey database configuration information to <code>mysqlbackup</code> is to use the <code>--defaults-file</code> option. When connection and data layout information is available from the configuration file, you no longer need to supply separately most of the information listed below.
MySQL port	MySQL configuration file or <code>mysqld</code> startup script.	Used to connect to the database instance during backup operations.

Information to Gather	Where to Find It	How to Use It
		Specified via the <code>--port</code> option of <code>mysqlbackup</code> . The specification is not needed if the information is available from the MySQL configuration file. Also not needed when doing a cold (offline) backup, which works directly on the files using OS-level file operations.
Path to MySQL data directory	MySQL configuration file or <code>mysqld</code> startup script.	Used to retrieve files from the database instance during backup operations, and to copy files back to the database instance during restore operations. Automatically retrieved from database connection for hot and warm backups, and taken from the MySQL configuration file for cold backups.
ID and password of privileged MySQL user	You record this during installation of your own databases, or get it from the DBA when backing up databases you do not own. Not needed when doing an offline (cold) backup, which works directly on the files using OS-level file permissions. For cold backups, you log in as an administrative user.	Specified via the <code>--password</code> option of the <code>mysqlbackup</code> . Prompted from the terminal if the <code>--password</code> option is present without the password argument.
Path under which to store backup data or metadata, temporarily or permanently	You choose this. See <a href="#">Section 4.1.3, “Designate a Location for the Backup Directory”</a> for details.	In general, this directory must be empty for <code>mysqlbackup</code> to write data into it.
Owner and permission information for backed-up files (for Linux, Unix, and OS X systems)	In the MySQL data directory.	If you do the backup using a different OS user ID or a different <code>umask</code> setting than applies to the original files, you might need to run commands such as <code>chown</code> and <code>chmod</code> on the backup data. See <a href="#">Section B.1, “Limitations of MySQL Enterprise Backup”</a> for details.
Size of InnoDB redo log files	Calculated from the values of the <code>innodb_log_file_size</code> and <code>innodb_log_files_in_group</code> configuration variables. Use the technique explained for the <code>--incremental-with-redo-log-only</code> option.	Only needed if you perform incremental backups using the <code>--incremental-with-redo-log-only</code> option rather than the <code>--incremental</code> option. The size of the InnoDB redo log and the rate of generation for redo data dictate how often you must perform incremental backups.



Information to Gather	Where to Find It	How to Use It
Rate at which redo data is generated	Calculated from the values of the InnoDB <a href="#">logical sequence number</a> at different points in time. Use the technique explained for the <a href="#">--incremental-with-redo-log-only</a> option.	Only needed if you perform incremental backups using the <a href="#">--incremental-with-redo-log-only</a> option rather than the <a href="#">--incremental</a> option. The size of the InnoDB redo log and the rate of generation for redo data dictate how often you must perform incremental backups.

## 4.1.2 Grant MySQL Privileges to Backup Administrator

For most backup operations, the [mysqlbackup](#) command connects to the MySQL server using the credentials supplied with the [--user](#) and [--password](#) options. The specified [user](#) needs certain privileges. You can either create a new user with a minimal set of privileges, or use an administrative account such as root.

The minimum privileges for the MySQL user with which [mysqlbackup](#) connects to the server are:

- [RELOAD](#) on all databases and tables.
- [CREATE](#), [INSERT](#), [DROP](#), and [UPDATE](#) on the tables [mysql.backup\\_progress](#) and [mysql.backup\\_history](#), and also [SELECT](#) on [mysql.backup\\_history](#).
- [SUPER](#), to enable and disable logging, and to optimize locking in order to minimize disruption to database processing.
- [REPLICATION CLIENT](#), to retrieve the [binary log](#) position, which is stored with the backup.
- [PROCESS](#), to process DDL statements with the [ALGORITHM = INPLACE](#) clause.

To set these privileges for a MySQL user ([mysqlbackup](#) in this example) connecting from localhost, issue statements like the following from the [mysql](#) client program:

```
GRANT RELOAD ON *.* TO 'mysqlbackup'@'localhost';
GRANT CREATE, INSERT, DROP, UPDATE ON mysql.backup_progress TO 'mysqlbackup'@'localhost';
GRANT CREATE, INSERT, SELECT, DROP, UPDATE ON mysql.backup_history TO 'mysqlbackup'@'localhost';
GRANT REPLICATION CLIENT ON *.* TO 'mysqlbackup'@'localhost';
GRANT SUPER ON *.* TO 'mysqlbackup'@'localhost';
GRANT PROCESS ON *.* TO 'mysqlbackup'@'localhost';
```

The following additional privileges are required for using [transportable tablespaces \(TTS\)](#) to back up and restore InnoDB tables:

- [LOCK TABLES](#) and [SELECT](#) for backing up tables
- [CREATE](#) and [ALTER](#) for restoring tables

The following additional privileges are required for creating tape backups using the System Backup to Tape (SBT) API ( see [Section 9.1, “Backing Up to Tape with Oracle Secure Backup”](#) for details):

- [CREATE](#), [INSERT](#), [DROP](#), and [UPDATE](#) on the [mysql.backup\\_sbt\\_history](#) table.

To set these additional privileges, issue a statement like the following from the [mysql](#) client program:

```
GRANT LOCK TABLES, SELECT, CREATE, ALTER ON *.* TO 'mysqlbackup'@'localhost';
GRANT CREATE, INSERT, DROP, UPDATE ON mysql.backup_sbt_history TO 'mysqlbackup'@'localhost';
```

### 4.1.3 Designate a Location for the Backup Directory

Most `mysqlbackup` operations, including those on single-file backups, write data or metadata to a designated directory referred to as the “**backup directory**” in this manual. See the description for `--backup-dir` for details on its usage for different operations.

Choose in advance for this directory a location on a file system with sufficient storage; it could even be remotely mounted from a different server. You specify the path to this directory with the `--backup-dir` option for many `mysqlbackup` commands.

If you use the backup directory as a location to store your backups, it is preferable to keep each backup within a timestamped subdirectory underneath the main backup directory. To make `mysqlbackup` create these subdirectories automatically, specify the `--with-timestamp` option each time you run `mysqlbackup`.

## 4.2 The Typical Backup / Verify / Restore Cycle

To illustrate the basic steps in creating and making use of a backup, the following example shows how to perform a full backup, verify it, and then restore it to a server.

### 4.2.1 Backing Up an Entire MySQL Instance

In the following example, we back up an entire MySQL instance to a single file using the `backup-to-image` command, which appears at the end of the sample command. We specify some of the connection information for the database using the `--user` and `--host` options (and, with the `--password` option, tell the MySQL server to prompt for a user password). The location and filename for the single-file backup is specified using the `--backup-image` option, and the location for an empty folder to store temporary files is supplied with the `--backup-dir` option.

The output echoes all the parameters used by the backup operation, including several that are retrieved automatically using the database connection. The unique ID for this backup job is recorded in special tables that `mysqlbackup` creates inside the MySQL instance, allowing you to monitor long-running backups and view information on previous backups. The final output section repeats the location of the backup data and provides the **LSN** values that you might use when you perform an **incremental backup** next time over the **full backup** that has just been made.

```
$ ./mysqlbackup --user=root --password --host=127.0.0.1 --backup-image=/home/admin/backups/my.mbi \
  --backup-dir=/home/admin/backup-tmp backup-to-image

MySQL Enterprise Backup version 4.0.0 Linux-3.2.0-25-generic-pae-i686 [2015/09/29]
Copyright (c) 2003, 2015, Oracle and/or its affiliates. All Rights Reserved.

mysqlbackup: INFO: Starting with following command line ...
./mysqlbackup --user=root --password --host=127.0.0.1
--backup-image=/home/admin/backups/my.mbi
--backup-dir=/home/admin/backup-tmp backup-to-image

mysqlbackup: INFO:
Enter password:
mysqlbackup: INFO: MySQL server version is '5.7.9'.
mysqlbackup: INFO: Got some server configuration information from running server.

IMPORTANT: Please check that mysqlbackup run completes successfully.
At the end of a successful 'backup-to-image' run mysqlbackup
prints "mysqlbackup completed OK!".
```

```
150930 19:19:26 mysqlbackup: INFO: MEB logfile created at /home/admin/backup-tmp/meta/MEB_2015-09-30.19-19
```

```
-----
Server Repository Options:
-----
```

```
datadir = /var/lib/mysql/
innodb_data_home_dir =
innodb_data_file_path = ibdata1:12M:autoextend
innodb_log_group_home_dir = /var/lib/mysql/
innodb_log_files_in_group = 2
innodb_log_file_size = 50331648
innodb_page_size = 16384
innodb_checksum_algorithm = crc32
innodb_undo_directory = /var/lib/mysql/
innodb_undo_tablespace = 0
innodb_undo_logs = 128
```

```
-----
Backup Config Options:
-----
```

```
datadir = /home/admin/backup-tmp/datadir
innodb_data_home_dir = /home/admin/backup-tmp/datadir
innodb_data_file_path = ibdata1:12M:autoextend
innodb_log_group_home_dir = /home/admin/backup-tmp/datadir
innodb_log_files_in_group = 2
innodb_log_file_size = 50331648
innodb_page_size = 16384
innodb_checksum_algorithm = crc32
innodb_undo_directory = /home/admin/backup-tmp/datadir
innodb_undo_tablespace = 0
innodb_undo_logs = 128
```

```
Backup Image Path = /home/admin/backups/my.mbi
```

```
mysqlbackup: INFO: Unique generated backup id for this is 14436551662849660
```

```
mysqlbackup: INFO: Creating 14 buffers each of size 16777216.
```

```
150930 19:19:28 mysqlbackup: INFO: Full Image Backup operation starts with following threads
1 read-threads 6 process-threads 1 write-threads
```

```
150930 19:19:28 mysqlbackup: INFO: System tablespace file format is Barracuda.
```

```
150930 19:19:28 mysqlbackup: INFO: Starting to copy all innodb files...
```

```
mysqlbackup: INFO: Could not find binlog index file. binlogs will not be copied for this backup.
Point-In-Time-Recovery will not be possible.
```

```
If this is online backup then server may not have started with --log-bin.
```

```
You may specify its location with --log-bin-index option.
```

```
mysqlbackup: INFO: Copying meta file /home/admin/backup-tmp/backup-my.cnf.
```

```
mysqlbackup: INFO: Copying meta file /home/admin/backup-tmp/meta/backup_create.xml.
```

```
150930 19:19:28 mysqlbackup: INFO: Found checkpoint at lsn 2478740.
```

```
150930 19:19:28 mysqlbackup: INFO: Starting log scan from lsn 2478592.
```

```
150930 19:19:28 mysqlbackup: INFO: Copying log...
```

```
150930 19:19:28 mysqlbackup: INFO: Log copied, lsn 2478749.
```

```
150930 19:19:28 mysqlbackup: INFO: Copying /var/lib/mysql/ibdata1 (Barracuda file format).
```

```
mysqlbackup: INFO: Processing General table spaces in /var/lib/mysql/
```

```
150930 19:19:28 mysqlbackup: INFO: Copying /var/lib/mysql/mydb/mytable.ibd (Barracuda file format).
```

```
150930 19:19:28 mysqlbackup: INFO: Copying /var/lib/mysql/mysql/engine_cost.ibd (Barracuda file format).
```

```
150930 19:19:28 mysqlbackup: INFO: Copying /var/lib/mysql/mysql/gtid_executed.ibd (Barracuda file format).
```

```
150930 19:19:28 mysqlbackup: INFO: Copying /var/lib/mysql/mysql/help_category.ibd (Barracuda file format).
```

```
150930 19:19:28 mysqlbackup: INFO: Copying /var/lib/mysql/mysql/help_keyword.ibd (Barracuda file format).
```

```
150930 19:19:28 mysqlbackup: INFO: Copying /var/lib/mysql/mysql/help_relation.ibd (Barracuda file format).
```

```
150930 19:19:28 mysqlbackup: INFO: Copying /var/lib/mysql/mysql/help_topic.ibd (Barracuda file format).
```

```
150930 19:19:28 mysqlbackup: INFO: Copying /var/lib/mysql/mysql/innodb_index_stats.ibd (Barracuda file format).
```

```
150930 19:19:28 mysqlbackup: INFO: Copying /var/lib/mysql/mysql/innodb_table_stats.ibd (Barracuda file format).
```

```
150930 19:19:28 mysqlbackup: INFO: Copying /var/lib/mysql/mysql/plugin.ibd (Barracuda file format).
```

```
150930 19:19:28 mysqlbackup: INFO: Copying /var/lib/mysql/mysql/server_cost.ibd (Barracuda file format).
```

```
150930 19:19:28 mysqlbackup: INFO: Copying /var/lib/mysql/mysql/servers.ibd (Barracuda file format).
```

```
150930 19:19:28 mysqlbackup: INFO: Copying /var/lib/mysql/mysql/slave_master_info.ibd (Barracuda file format).
```

```
150930 19:19:28 mysqlbackup: INFO: Copying /var/lib/mysql/mysql/slave_relay_log_info.ibd (Barracuda file format).
```

```

150930 19:19:28 mysqlbackup: INFO: Copying /var/lib/mysql/mysql/slave_worker_info.ibd (Barracuda file format).
150930 19:19:28 mysqlbackup: INFO: Copying /var/lib/mysql/mysql/time_zone.ibd (Barracuda file format).
150930 19:19:28 mysqlbackup: INFO: Copying /var/lib/mysql/mysql/time_zone_leap_second.ibd (Barracuda file format).
150930 19:19:28 mysqlbackup: INFO: Copying /var/lib/mysql/mysql/time_zone_name.ibd (Barracuda file format).
150930 19:19:28 mysqlbackup: INFO: Copying /var/lib/mysql/mysql/time_zone_transition.ibd (Barracuda file format).
150930 19:19:28 mysqlbackup: INFO: Copying /var/lib/mysql/mysql/time_zone_transition_type.ibd (Barracuda file format).
150930 19:19:29 mysqlbackup: INFO: Copying /var/lib/mysql/sys/sys_config.ibd (Barracuda file format).
150930 19:19:29 mysqlbackup: INFO: Completing the copy of innodb files.
mysqlbackup: INFO: Processing General table spaces in /var/lib/mysql/
mysqlbackup: INFO: Now complete gen ts
150930 19:19:30 mysqlbackup: INFO: Preparing to lock tables: Connected to mysqld server.
150930 19:19:30 mysqlbackup: INFO: Starting to lock all the tables...
150930 19:19:30 mysqlbackup: INFO: All tables are locked and flushed to disk
150930 19:19:30 mysqlbackup: INFO: Opening backup source directory '/var/lib/mysql/'
150930 19:19:30 mysqlbackup: INFO: Starting to backup all non-innodb files in
subdirectories of '/var/lib/mysql/'
150930 19:19:30 mysqlbackup: INFO: Adding database directory: datadir/mydb
150930 19:19:30 mysqlbackup: INFO: Adding database directory: datadir/mysql
150930 19:19:30 mysqlbackup: INFO: Adding database directory: datadir/performance_schema
150930 19:19:30 mysqlbackup: INFO: Adding database directory: datadir/sys
150930 19:19:30 mysqlbackup: INFO: Completing the copy of all non-innodb files.
mysqlbackup: INFO: Processing General table spaces in /var/lib/mysql/
mysqlbackup: INFO: Now complete gen ts
mysqlbackup: INFO: Processing General table spaces in /home/admin/backup-tmp/datadir
mysqlbackup: INFO: Processing General table spaces in /home/admin/backup-tmp/datadir
mysqlbackup: INFO: Now complete gen ts
150930 19:19:32 mysqlbackup: INFO: A copied database page was modified at 2478707.
(This is the highest lsn found on page)
Scanned log up to lsn 2478749.
Was able to parse the log up to lsn 2478749.
150930 19:19:32 mysqlbackup: INFO: All tables unlocked
150930 19:19:32 mysqlbackup: INFO: All MySQL tables were locked for 2.089 seconds.
150930 19:19:32 mysqlbackup: INFO: Reading all global variables from the server.
150930 19:19:32 mysqlbackup: INFO: Completed reading of all global variables from the server.
150930 19:19:32 mysqlbackup: INFO: Creating server config files server-my.cnf and server-all.cnf in /home/admin/
mysqlbackup: INFO: Copying meta file /home/admin/backup-tmp/meta/backup_variables.txt.
mysqlbackup: INFO: Copying meta file /home/admin/backup-tmp/datadir/ibbackup_logfile.
mysqlbackup: INFO: Copying meta file /home/admin/backup-tmp/server-all.cnf.
mysqlbackup: INFO: Copying meta file /home/admin/backup-tmp/server-my.cnf.
mysqlbackup: INFO: Copying meta file /home/admin/backup-tmp/meta/backup_content.xml.
mysqlbackup: INFO: Copying meta file /home/admin/backup-tmp/meta/image_files.xml.
150930 19:19:35 mysqlbackup: INFO: Full Image Backup operation completed successfully.
150930 19:19:35 mysqlbackup: INFO: Backup image created successfully.
mysqlbackup: INFO: Image Path = /home/admin/backups/my.mbi

-----
Parameters Summary
-----
Start LSN          : 2478592
End LSN            : 2478749
-----

mysqlbackup completed OK!

```

### 4.2.2 Verifying a Backup

You can check the integrity of your backup using the `validate` command. The following is a sample command for validating an encrypted backup image and the output for the successful validation:

```

$ mysqlbackup -backup-image=/meb/backups/image.mbi --decrypt --key-file=/meb/enckeyfile validate

140219 11:22:44 mysqlbackup: INFO: Validating image ... /logs/img.bi
140219 11:22:44 mysqlbackup: INFO: Validate: [Dir]: meta
140219 11:22:45 mysqlbackup: INFO: Total files as specified in image: 44
mysqlbackup: INFO: datadir/tpch/tabnorm7.ibd Validated...

```

```
mysqlbackup: INFO: datadir/tpch/tabnorm8.ibd Validated...
mysqlbackup: INFO: datadir/tpch/tabnorm9.ibd Validated...
.....
140219 11:22:45 mysqlbackup: INFO: Validate operation completed successfully.
140219 11:22:45 mysqlbackup: INFO: Backup Image validation successful.
mysqlbackup: INFO: Source Image Path = /logs/img.bi
mysqlbackup completed OK!
```

This is a sample output for a checksum mismatch in the header:

```
mysqlbackup: ERROR: Checksum mismatch.
Computed checksum: ###          Checksum in image: ### mysqlbackup: ERROR: Problem verifying checksum o
Image Path = /meb/backups/image.mbi
mysqlbackup: ERROR: Backup image validation failed.
```

This is a sample output for an image containing corrupted `.ibd` files:

```
mysqlbackup: ERROR: datadir/db2/bigtab1.ibd has corrupt page number : 64   page number from page header : 0
mysqlbackup: ERROR: datadir/db2/bigtab1.ibd is corrupt and has : 10 corrupt pages
mysqlbackup: ERROR: datadir/db2/t1.ibd has corrupt page number : 4   page number from   page header : 0
.....
mysqlbackup: ERROR: datadir/db2/t1.ibd is corrupt and has : 5 corrupt pages
mysqlbackup: ERROR: datadir/ibdata1 has corrupt page number : 63   page number from page header : 63
mysqlbackup: ERROR: datadir/ibdata1 has corrupt page number : 7   page number from page header : 7
.....
mysqlbackup: ERROR: datadir/ibdata1 is corrupt and has : 10 corrupt pages
mysqlbackup failed with errors!
```

This is a sample output for a successful validation for a compressed backup directory

```
mysqlbackup: INFO: /backups/backup-dir/datadir/tpch/tabnorm5.ibz Validated...
mysqlbackup: INFO: /backups/backup-dir/datadir/tpch/tabnorm6.ibz Validated...
mysqlbackup: INFO: /backups/backup-dir/datadir/tpch/tabnorm7.ibz Validated...
mysqlbackup: INFO: /backups/backup-dir/datadir/tpch/tabnorm8.ibz Validated...
mysqlbackup: INFO: /backups/backup-dir/datadir/tpch/tabnorm9.ibz Validated...
mysqlbackup: INFO: /backups/backup-dir/datadir/tpch/tabrowformat.ibz Validated...
140219 11:22:45 mysqlbackup: INFO: Validate backup directory operation completed successfully.
```

Furthermore, you can also verify that your backup has been successful by restoring the backup data on a different server and run the MySQL daemon (`mysqld`) on the new data directory. You can then execute `SHOW` statements to verify the database and table structures, and execute queries to verify further details of the database. See [Section 4.2.3, “Restoring a Database”](#) for basic steps for restoring a backup, and see [Chapter 5, \*Recovering or Restoring a Database\*](#) for more detailed instructions. Running the `mysqld` daemon on the restored data requires a valid configuration file, which you specify with the `--defaults-file` option of the `mysqld` command. You can reuse most of the settings from the original `my.cnf` file of the backed up MySQL instance, combined with the settings from the `backup-my.cnf` file, which was created in the temporary directory you specified with `--backup-dir` when you created a single-image backup (see [Section 4.2.1, “Backing Up an Entire MySQL Instance”](#)) and contains a small subset of parameters required by `mysqlbackup`. Create a new configuration file by concatenating the two files mentioned above into a new one, and use that file on the server on which you perform the verification. Edit the file to make sure the `datadir` parameter points to the right location on the verification server. Edit the values for port, socket, and so on if you need to use different connection settings on the verification server.



### Warning

Do not try to verify a backup by using the `backup directory` a data directory to start a MySQL server. This will crash the server, and might also corrupt your backup. See [Appendix A, \*Frequently Asked Questions for MySQL Enterprise Backup\*](#).

## 4.2.3 Restoring a Database

To restore a MySQL instance from a backup:

- Shut down the database server.
- Use, for example, `mysqlbackup` with the `copy-back-and-apply-log` command, which makes the raw backup into a prepared backup by updating it to a consistent state and then copies the tables, indexes, metadata, and any other required files back to their original locations, as defined by the original MySQL configuration file. For the different combinations of options that you can specify as part of this operation, see [Section 12.3, “Restore Operations”](#).

In the example below, the single-file backup created in the example given in [Section 4.2.1, “Backing Up an Entire MySQL Instance”](#) is restored using the `copy-back-and-apply-log` command. Besides the usual connection parameters, the following options are used in the command:

- `--defaults-file` supplies the configuration file `backup-my.cnf`, which was created in the temporary directory you specified with `--backup-dir` when you created the single-image backup (see [Section 4.2.2, “Verifying a Backup”](#) for more explanations) and is now at the location specified with the option. Without specifying the `--defaults-file` option, the restore can still be performed, as long as the related InnoDB parameters for the backed-up and the target server match; you are going to receive some warnings though from `mysqlbackup`.
- `--datadir` supplies the location of the data directory for restoring the data.
- `--backup-image` provides the path of the single-file backup.
- `--backup-dir` provides the location of an empty folder to store some temporary files created during the restore procedure.

```
$ ./mysqlbackup --defaults-file=/home/admin/backup-tmp/backup-my.cnf --datadir=/var/lib/mysql \
--backup-image=/home/admin/backups/my.mbi --backup-dir=/home/admin/restore-tmp copy-back-and-apply-log
MySQL Enterprise Backup version 4.0.0 Linux-3.2.0-25-generic-pae-i686 [2015/09/29]
Copyright (c) 2003, 2015, Oracle and/or its affiliates. All Rights Reserved.

mysqlbackup: INFO: Starting with following command line ...
./bin/mysqlbackup --defaults-file=/home/admin/backup-tmp/backup-my.cnf
--datadir=/var/lib/mysql --backup-image=/home/admin/backups/my.mbi
--backup-dir=/home/admin/restore-tmp copy-back-and-apply-log

mysqlbackup: INFO:
IMPORTANT: Please check that mysqlbackup run completes successfully.
          At the end of a successful 'copy-back-and-apply-log' run mysqlbackup
          prints "mysqlbackup completed OK!".

mysqlbackup: INFO: Backup Image MEB version string: 4.0.0 [2015/09/29]
151008 12:35:23 mysqlbackup: INFO: MEB logfile created at /home/admin/restore-tmp/meta/MEB_2015-10-08.12-35-23

-----
Server Repository Options:
-----
datadir = /var/lib/mysql
innodb_data_home_dir = /var/lib/mysql
innodb_data_file_path = ibdata1:12M:autoextend
innodb_log_group_home_dir = /var/lib/mysql
innodb_log_files_in_group = 2
innodb_log_file_size = 50331648
innodb_page_size = 16384
innodb_checksum_algorithm = crc32
-----

Backup Config Options:
-----
datadir = /home/admin/restore-tmp/datadir
```

```

innodb_data_home_dir = /home/admin/restore-tmp/datadir
innodb_data_file_path = ibdata1:12M:autoextend
innodb_log_group_home_dir = /home/admin/restore-tmp/datadir
innodb_log_files_in_group = 2
innodb_log_file_size = 50331648
innodb_page_size = 16384
innodb_checksum_algorithm = crc32

mysqlbackup: INFO: Creating 14 buffers each of size 16777216.
151008 12:35:23 mysqlbackup: INFO: Copy-back-and-apply-log operation starts with following threads
  1 read-threads    6 process-threads    1 write-threads
151008 12:35:23 mysqlbackup: INFO: Copying database directory: meta
151008 12:35:23 mysqlbackup: INFO: Copying datadir/ibdata1.
151008 12:35:23 mysqlbackup: INFO: Copying datadir/mydb/mytable.ibd.
151008 12:35:23 mysqlbackup: INFO: Copying database directory: datadir/mydb
151008 12:35:23 mysqlbackup: INFO: Copying datadir/mysql/engine_cost.ibd.
151008 12:35:23 mysqlbackup: INFO: Copying database directory: datadir/mysql
151008 12:35:23 mysqlbackup: INFO: Copying datadir/mysql/gtid_executed.ibd.
151008 12:35:23 mysqlbackup: INFO: Copying datadir/mysql/help_category.ibd.
151008 12:35:23 mysqlbackup: INFO: Copying datadir/mysql/help_keyword.ibd.
151008 12:35:23 mysqlbackup: INFO: Copying datadir/mysql/help_relation.ibd.
151008 12:35:23 mysqlbackup: INFO: Copying datadir/mysql/help_topic.ibd.
151008 12:35:23 mysqlbackup: INFO: Copying datadir/mysql/innodb_index_stats.ibd.
151008 12:35:23 mysqlbackup: INFO: Copying datadir/mysql/innodb_table_stats.ibd.
151008 12:35:23 mysqlbackup: INFO: Copying datadir/mysql/plugin.ibd.
151008 12:35:23 mysqlbackup: INFO: Copying datadir/mysql/server_cost.ibd.
151008 12:35:23 mysqlbackup: INFO: Copying datadir/mysql/servers.ibd.
151008 12:35:24 mysqlbackup: INFO: Copying datadir/mysql/slave_master_info.ibd.
151008 12:35:24 mysqlbackup: INFO: Copying datadir/mysql/slave_relay_log_info.ibd.
151008 12:35:24 mysqlbackup: INFO: Copying datadir/mysql/slave_worker_info.ibd.
151008 12:35:24 mysqlbackup: INFO: Copying datadir/mysql/time_zone.ibd.
151008 12:35:24 mysqlbackup: INFO: Copying datadir/mysql/time_zone_leap_second.ibd.
151008 12:35:24 mysqlbackup: INFO: Copying datadir/mysql/time_zone_name.ibd.
151008 12:35:24 mysqlbackup: INFO: Copying datadir/mysql/time_zone_transition.ibd.
151008 12:35:24 mysqlbackup: INFO: Copying datadir/mysql/time_zone_transition_type.ibd.
151008 12:35:24 mysqlbackup: INFO: Copying datadir/sys/sys_config.ibd.
151008 12:35:24 mysqlbackup: INFO: Copying database directory: datadir/performance_schema
151008 12:35:24 mysqlbackup: INFO: Copying database directory: datadir/sys
151008 12:35:24 mysqlbackup: INFO: Copying database directory: datadir/mydb
151008 12:35:24 mysqlbackup: INFO: Copying database directory: datadir/mysql
151008 12:35:24 mysqlbackup: INFO: Copying database directory: datadir/performance_schema
151008 12:35:25 mysqlbackup: INFO: Copying database directory: datadir/sys
151008 12:35:26 mysqlbackup: INFO: Total files as specified in image: 297
151008 12:35:26 mysqlbackup: INFO: Creating server config files server-my.cnf and server-all.cnf in /var/lib/
151008 12:35:26 mysqlbackup: INFO: Copy-back operation completed successfully.
mysqlbackup: INFO: Source Image Path = /home/admin/backups/my.mbi

mysqlbackup: INFO: Restoring ...5.7.9 version
mysqlbackup: INFO: Creating 14 buffers each of size 65536.
151008 12:35:26 mysqlbackup: INFO: Apply-log operation starts with following threads
  1 read-threads    1 process-threads
mysqlbackup: INFO: Using up to 100 MB of memory.
151008 12:35:26 mysqlbackup: INFO: ibbackup_logfile's creation parameters:
    start lsn 2478592, end lsn 2478749,
    start checkpoint 2478740.
mysqlbackup: INFO: InnoDB: Doing recovery: scanned up to log sequence number 2478749.
mysqlbackup: INFO: InnoDB: Starting an apply batch of log records to the database...
.
InnoDB: Progress in percent: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28
mysqlbackup: INFO: InnoDB: Setting log file size to 50331648.
mysqlbackup: INFO: InnoDB: Setting log file size to 50331648.
151008 12:35:29 mysqlbackup: INFO: We were able to parse ibbackup_logfile up to
    lsn 2478749.
151008 12:35:29 mysqlbackup: INFO: The first data file is '/var/lib/mysql/ibdata1'
    and the new created log files are at '/var/lib/mysql'
mysqlbackup: INFO: Restoring ...5.7.9 version

```



```
151008 12:35:29 mysqlbackup: INFO: Apply-log operation completed successfully.
151008 12:35:29 mysqlbackup: INFO: Full Backup has been restored successfully.

mysqlbackup completed OK!
```

Now the original database directory is restored from the backup, and you can restart the database server.

## 4.3 Backup Scenarios and Examples

### 4.3.1 Making a Single-File Backup

To avoid having a large number of backup files to keep track, store, and transport, `mysqlbackup` conveniently creates backups in a single-file format. It can also pack an existing backup directory into a single file, unpack the single file back to a backup directory, list the contents of a single-file backup, verify the contents of a single-file backup against embedded checksums, or extract a single file into a directory tree. For the syntax of the relevant `mysqlbackup` options, see [Section 13.9, “Single-File Backup Options”](#).

*Advanced:* While `mysqlbackup` can also create a directory backup (see description for the `backup` command for details) instead of a single-file backup, the single-file format is preferable in most cases: a single-file backup is easier to handle and store, and certain functions of `mysqlbackup` are not supported for directory backups—for example, backup to cloud and backup to tape using the System Backup to Tape (SBT) API. Throughout the manual, directory backup is mostly treated as an advanced topic, and information and examples for directory backups are marked with the *Advanced* tag, like this paragraph.

Because the single-file backup can be streamed or piped to another process such as a tape backup or a command like `scp`, you can use the technique to put the backup onto another storage device or server and avoid significant storage overhead on the original database server.

To create a single-file backup, use the `backup-to-image` command. The following examples illustrate how to perform a single-file backup and other related operations.

#### Example 4.1 Single-File Backup to Absolute Path

This command creates a single backup image on the given absolute path. It still requires `--backup-dir`, which is used to hold temporary output, status, and metadata files.

```
mysqlbackup --backup-image=/backups/sales.mbi --backup-dir=/backup-tmp backup-to-image
```

#### Example 4.2 Single-File Backup to Relative Path

When a relative path instead of an absolute path was supplied with the `--backup-image` option, the path is taken to be relative to the backup directory. Therefore, in this example, the resulting single-file backup is created as `/backups/sales.mbi`.

```
mysqlbackup --backup-image=sales.mbi --backup-dir=/backups backup-to-image
```

#### Example 4.3 Single-File Backup to Standard Output

The following command dumps the backup output to standard output. Again, the folder specified with the `--backup-dir` option is used as a temporary directory.



```
mysqlbackup --backup-dir=/backups --backup-image=- backup-to-image > /backup/mybackup.mbi
```

#### Example 4.4 Convert Existing Backup Directory to Single Image

The `backup-dir` directory is bundled into the `/backup/my.mbi` file.

```
mysqlbackup --backup-image=/backup/my.mbi --backup-dir=/var/mysql/backup backup-dir-to-image
```

#### Example 4.5 Extract Existing Image to Backup Directory

The image contents are unpacked into `backup-dir`.

```
mysqlbackup --backup-dir=/var/backup --backup-image=/backup/my.mbi image-to-backup-dir
```

#### Example 4.6 List Single-File Backup Contents

The image contents are listed, with each line indicating a file or directory entry.

```
mysqlbackup --backup-image=/backup/my.mbi list-image
```

#### Example 4.7 Validate a Single-File Backup

The following command verifies that the single-file backup is not corrupted, truncated, or damaged by validating the checksum value for each data page in the backup.

```
mysqlbackup --backup-image=/logs/fullimage.mi validate
```

#### Example 4.8 Extract Single-File Backup into Current Directory

The following command extracts all contents from a single-file backup into the current working directory.

```
mysqlbackup --backup-image=/var/my.mbi extract
```

#### Example 4.9 Extract Single-File Backup into a Backup Directory

This command behaves like the `image-to-backup-dir` option by extracting all contents of a single-file backup into the directory specified with the `--backup-dir` option.

```
mysqlbackup --backup-image=/var/my.mbi --backup-dir=/var/backup extract
```

#### Example 4.10 Selective Extract of Single File

The following command extracts the single file `meta/comments.txt` from the backup image `my.mbi` into the local path `./meta/comments.txt`.

```
mysqlbackup --backup-image=/var/my.mbi \  
--src-entry=meta/comments.txt extract
```

The following command extracts the `meta/comments.txt` file from the backup image `my.mbi` into a specified path `/tmp/mycomments.txt` by using the `--dst-entry` option.

```
mysqlbackup --backup-image=/var/my.mbi \
--src-entry=meta/comments.txt \
--dst-entry=/tmp/mycomments.txt extract
```

The following command dumps the contents of `meta/comments.txt` (which is inside the single-file backup `my.mbi`) to standard output.

```
mysqlbackup --backup-image=/var/my.mbi --src-entry=meta/comments.txt --dst-entry=- extract
```

#### Example 4.11 Selective Extract of Single Directory

The following command extracts a single directory `meta` from the backup image `my.mbi` into a local file system path `./meta`. All contents in the `meta` directory are extracted, including any subdirectories.

```
mysqlbackup --backup-image=/backup/my.mbi --src-entry=meta extract
```

The following command extracts all contents of the `meta` directory, including all its files and subdirectories, into the directory `/tmp/my-meta`.

```
mysqlbackup --backup-image=/backup/my.mbi --src-entry=meta \
--dst-entry=/tmp/my-meta extract
```

#### Example 4.12 Dealing with Absolute Path Names

Since absolute pathnames are extracted to the same paths in local system, it could be a problem if you do not have write permission for that path. You can remap absolute paths as follows:

```
mysqlbackup --backup-image=/backup/my.mbi --src-entry=/ --dst-entry=/myroot extract
mysqlbackup --backup-image=/backup/my.mbi --src-entry=. extract
```

The first command extracts all absolute paths to `/myroot` directory in the local system. The second command extracts all relative paths to the current directory.

### 4.3.1.1 Streaming the Backup Data to Another Device or Server

To limit the storage overhead on the database server, you can transfer the backup data to a different server without ever storing it locally. You can achieve that with a single-file backup. To send the single-file backup to standard output, use the `mysqlbackup` command `backup-to-image` without specifying the `--backup-image` option. (You can also specify `--backup-image=-` to make it obvious that the data is sent to stdout.) To stream the data, you use the single-file backup in combination with operating system features such as pipes, `ssh`, `scp`, and so on, which take the input from standard output and create an equivalent file on a remote system. You can either store the single-file backup directly on the remote system, or invoke `mysqlbackup` with the `copy-back-and-apply-log` command on the other end to restore the backup to a remote MySQL server.

#### Example 4.13 Single-File Backup to a Remote Host

The following command streams the backup as a single-file output to a remote host, where it may be saved directly to a tape device. `--backup-dir=/tmp` designates the directory for storing temporary files rather than the final output file.

```
mysqlbackup --backup-image=- --backup-dir=/tmp backup-to-image | \
ssh user@host command arg1 arg2...
```

For simplicity, all the connection and other necessary options are assumed to be specified in the default configuration file. To have the desired operations run on the remote system, substitute the combination of command, device, and so on that you use as part of your normal archiving procedure, such as [dd](#) or [tar](#).

#### Example 4.14 Single-file Backup to a Remote MySQL Server

The following command streams the backup as a single backup file to be restored on a remote MySQL server:

```
mysqlbackup --backup-dir=backup --backup-image=- --compress backup-to-image | \
ssh <user name>@<remote host name> 'mysqlbackup --backup-dir=backup_tmp --datadir=/data \
--innodb_log_group_home_dir=. \
--innodb_log_files_in_group=<innodb_log_files_in_group_of_backedup_server> \
--innodb_log_file_size=<innodb_log_file_size_of_backedup_server> \
--innodb_data_file_path=<innodb_data_file_path_of_backedup_server> \
--uncompress --backup-image=- copy-back-and-apply-log'
```

#### Example 4.15 Stream a Backup Directory to a Remote MySQL Server

The following command streams a backup directory as a single backup file to be restored on a remote MySQL server:

```
mysqlbackup --backup-image=- --backup-dir=/path/to/my/backup backup-dir-to-image | \
ssh <user name>@<remote host name> \
'mysqlbackup --backup-dir=backup_tmp --datadir=/data --backup-image=- copy-back-and-apply-log'
```

### 4.3.1.2 Backing Up to Tape

Tape drives are affordable, high-capacity storage devices for backup data. MySQL Enterprise Backup can interface with media management software (MMS) such as Oracle Secure Backup (OSB) to drive MySQL backup and restore jobs. The media management software must support Version 2 or higher of the System Backup to Tape (SBT) interface.

For information about doing tape backups in combination with MMS products such as Oracle Secure Backup, see [Chapter 9, Using MySQL Enterprise Backup with Media Management Software \(MMS\) Products](#).

### 4.3.1.3 Backing Up to Cloud Storage

MySQL Enterprise Backup supports cloud backups. Only single-file backups can be created on and restored from cloud storage. All [mysqlbackup](#) options compatible with single-file operations (including, for example, the [--incremental](#), [--compression](#), [--partial](#), and [--encryption](#) options) can be used with cloud backups or restores.

Currently, MySQL Enterprise Backup supports two types of cloud storage services: Amazon S3 and OpenStack Object Storage (Swift).

MySQL Enterprise Backup 4.0 supports the Swift v1.0 API, and also the OpenStack Identity (Keystone) API v2.0 for authentication. It also supports authentication using Swift's TempAuth system. Backups are stored as dynamic large objects in Swift, with each backup larger than 5G being split into multiple parts with names in the form of [<object\\_name>\\_part\\_<number>](#). See [the OpenStack documentation](#) for details.

A cloud backup is created using the cloud options for [mysqlbackup](#), which are described in details in [Section 13.14, “Cloud Storage Options”](#). Here are some sample commands for creating a cloud backup:

### Example 4.16 Creating a Cloud Backup on Amazon S3

```
mysqlbackup\
--cloud-service=s3 --cloud-aws-region=<aws region> \
--cloud-access-key-id=<aws access key id> --cloud-secret-access-key=< aws secret access key> \
--cloud-bucket=<s3 bucket name> --cloud-object-key=<aws object key> \
--backup-dir=/home/user/dba/s3backuptmpdir \
--backup-image=- \
backup-to-image
```

### Example 4.17 Creating a Cloud Backup on an OpenStack Object Storage

This example uses the Keystone identity service for authenticating user credentials.

```
mysqlbackup \
--include-tables=testdb.t1 --use-tts=with-full-locking \
--cloud-service=openstack --cloud-container=<swift container> \
--cloud-user-id=<keystone user> --cloud-password=<keystone password> \
--cloud-region=<keystone region> --cloud-tenant=<keystone tenant> \
--cloud-identity-url=<keystone url> \
--cloud-trace=1 --cloud-object=image_800.mbi \
--backup-dir=/home/user/dba/s3backuptmpdir \
--backup-image=- \
backup-to-image
```

Besides `backup-to-image`, all other `mysqlbackup` operations for single-file backups (`backup-dir-to-image`, `list-image`, `validate`, `image-to-backup-dir`, `extract`, `copy-back`, and `copy-back-and-apply-log`) can also be performed with cloud storage. For example:

### Example 4.18 Extract an Existing Image from Amazon S3 Cloud Storage to Backup Directory

Extract a backup image from Amazon S3, using the `--backup-dir` option to specify the directory into which the image will be extracted:

```
mysqlbackup\
--cloud-service=s3 --cloud-aws-region=<aws region> \
--cloud-access-key-id=<aws access key id> --cloud-secret-access-key=< aws secret access key> \
--cloud-bucket=<s3 bucket name> --cloud-object-key=<aws object key> \
--backup-dir=/home/user/dba/s3backuptmpdir \
--backup-image=- \
image-to-backup-dir
```

See [Section 5.1.5, “Restoring a Backup from Cloud Storage to a MySQL Server”](#) on how to restore a backup image from a cloud storage.

## 4.3.2 Making a Full Backup

Most backup strategies start with a complete backup of the MySQL server, from which you can restore all databases and tables. After you have created a [full backup](#), you might perform [incremental backups](#) (which are smaller and faster) for the next several backup tasks. You then make a full backup periodically to begin the cycle again.

For sample commands for making a full backup, see [Section 4.2.1, “Backing Up an Entire MySQL Instance”](#).

This section outlines some of the things to consider when deciding on a strategy for creating full backups. As we shall see, factors like speed, capacity, and convenience are all relevant for your decisions.

## Options on Command Line or in Configuration File?

For clarity, the examples in this manual often show some of the command-line options that are used with the `mysqlbackup` commands. For convenience and consistency, you can include those options that remain unchanged for most backup jobs into the `[mysqlbackup]` section of the MySQL configuration file that you supply to `mysqlbackup`. `mysqlbackup` also picks up the options from the `[mysqld]` section if they are present there. Putting the options into a configuration file can simplify backup administration for you: for example, putting port information into a configuration file, you can avoid the need to edit your backup scripts each time the database instance switches to a different port. See [Chapter 14, Configuration Files and Parameters](#) for details about the use of configuration files.

## Using a Single Backup Directory or Timestamped Subdirectories?

For convenience, the `--with-timestamp` option creates uniquely named subdirectories under the [backup directory](#) to hold the backup data (permanent or temporary) and metadata. The timestamped subdirectories make it simpler to establish retention periods, allowing easy removal and archiving of backup data that has passed a certain age.

If you do use a single backup directory (that is, if you omit the `--with-timestamp` option), specify a new, unique directory name for each backup job.

For incremental backups that uses the `--incremental-base` option to specify the directory containing the previous backup, in order to make the directory names predictable, you might prefer to not use the `--with-timestamp` option and generate a sequence of directory names with your backup script instead .

## Always Full Backup, or Full Backup plus Incremental Backups?

If your InnoDB data volume is small, or if your database is so busy that a high percentage of data changes between backups, you might want to run a full backup each time. However, you can usually save time and storage space by running periodic full backups and then several incremental backups in between them, as described in [Section 4.3.3, “Making a Differential or Incremental Backup”](#).

## Use Compression or Not?

Creating a compressed backup can save you considerable storage space and reduce I/O usage significantly. And with the LZ4 compression method, the overhead for processing compression is quite low. In cases where database backups are moving from a faster disk system where the active database files sit to a possibly slower storage, compression will often significantly lower the overall backup time. It can result in reduced restoration time as well. In general, we recommend LZ4 compression over no compression for most users, as LZ4-based backups often finish in a shorter time period. However, test out MySQL Enterprise Backup within your environment to determine what is the most efficient approach. For more discussions on compressed backups, see [Section 4.3.4, “Making a Compressed Backup”](#).

### 4.3.3 Making a Differential or Incremental Backup

Assuming a good portion of the data on your MySQL server remains unchanged over time, you can increase the speed and reduce the required storage space for your regular backups by backing up not all the data on the server each time, but only the changes to the data which have taken place over time. In order to that, after making first a full backup that contains all data, you can do one of the following:

- **Performing a series of differential backups.** Each [differential backup](#) includes all the changes made to the data since the last full backup was performed. To restore data up to, for example, time `t`, you simply restore first the full backup, and then, on top of it, the differential backup taken for time `t`.
- **Perform a series of incremental backup.** Each [incremental backup](#) only includes the changes since the previous backup, which can itself be a full or incremental backup. The first backup in an incremental

series is always then a differential backup; but after that, each incremental backup only contains the changes made since that last incremental backup. Each subsequent incremental backup is thus usually smaller in size than a differential backup, and is faster to make; that allows you to make very frequent incremental backups, and then enables you to restore the database to a more precise point in time when necessary. However, restoring data with incremental backups might take longer and more work: in general, to restore data up to, for example, time *t*, you start with restoring the full backup, and then restore the incremental backups one by one, until you are finished with the incremental backup taken for time *t*.

MySQL Enterprise Backup supports both incremental and differential backups. You should decide on which backup strategy to adopt by looking at such factors like how much storage space you have, how quickly you have to be able to restore data, and so on.

MySQL Enterprise Backup treats differential backup as a special case of incremental backup that has a full backup as its base. To create a differential backup, simply follow the instructions below for performing incremental backups, and make sure you specify a full backup as the base of your incremental backup using the methods we describe below; you should also ignore any instructions that only apply to the handling of multiple incremental backups.

See [Section 13.7, “Incremental Backup Options”](#), for descriptions of the `mysqlbackup` options used for incremental backups. An Incremental backup is enabled with one of the two options: `--incremental` and `--incremental-with-redo-log-only` option. See [Creating Incremental Backups Using Only the Redo Log](#) for their differences.

When creating an incremental backup, you have to indicate to `mysqlbackup` the point in time of the previous full or incremental backup. For convenience, you can use the `--incremental-base` option to automatically derive the necessary [log sequence number \(LSN\)](#) from the metadata stored in a previous backup directory or on the server. Or, you can specify an explicit LSN value using the `--start-lsn` option, providing to `mysqlbackup` the ending LSN from a previous full or incremental backup.

To prepare the backup data to be restored, you combine all incremental backups with an original full backup. Typically, you perform a new full backup after a designated period of time, after which you can discard the older incremental backup data.

## Creating Incremental Backups Using Only the Redo Log

The `--incremental-with-redo-log-only` might offer some benefits over the `--incremental` option for creating an incremental backup:

- The changes to InnoDB tables are determined based on the contents of the [InnoDB redo log](#). Since the redo log files have a fixed size that you know in advance, it can require less I/O to read the changes from them than to scan the InnoDB tablespace files to locate the changed pages, depending on the size of your database, amount of DML activity, and size of the redo log files.
- Since the redo log files act as a circular buffer, with records of older changes being overwritten as new DML operations take place, you must take new incremental backups on a predictable schedule dictated by the size of the log files and the amount of redo data generated for your workload. Otherwise, the redo log might not reach back far enough to record all the changes since the previous incremental backup, in which case `mysqlbackup` will quickly determine that it cannot proceed and will return an error. Your backup script should be able to catch that error and then perform an incremental backup with the `--incremental` option instead.

For example:

- To calculate the size of the redo log, issue the command `SHOW VARIABLES LIKE 'innodb_log_file%'` and, based on the output, multiply the `innodb_log_file_size` setting



by the value of `innodb_log_files_in_group`. To compute the redo log size at the physical level, look into the `datadir` directory of the MySQL instance and sum up the sizes of the files matching the pattern `ib_logfile*`.

- The InnoDB `LSN` value corresponds to the number of bytes written to the redo log. To check the LSN at some point in time, issue the command `SHOW ENGINE INNODB STATUS` and look under the `LOG` heading. While planning your backup strategy, record the LSN values periodically and subtract the earlier value from the current one to calculate how much redo data is generated each hour, day, and so on.

Prior to MySQL 5.5, it was common practice to keep the redo logs fairly small to avoid a long startup time when the MySQL server was killed rather than shut down normally. With MySQL 5.5 and higher, the performance of [crash recovery](#) is significantly improved, as described in [Optimizing InnoDB Configuration Variables](#), so that you can make your redo log files bigger if that helps your backup strategy and your database workload.

- This type of incremental backup is not so forgiving of too-low `--start-lsn` values as the standard `--incremental` option is. For example, you cannot make a full backup and then make a series of `--incremental-with-redo-log-only` backups all using the same `--start-lsn` value. Make sure to specify the precise end LSN of the previous backup as the start LSN of the next incremental backup; do not use arbitrary values.



#### Note

To ensure the LSN values match up exactly between successive incremental backups, it is recommended that you always use the `--incremental-base` option when you use the `--incremental-with-redo-log-only` option.

- To judge whether this type of incremental backup is practical and efficient for a particular MySQL instance:
  - Measure how fast the data changes within the InnoDB redo log files. Check the `LSN` periodically to decide how much redo data accumulates over the course of some number of hours or days.
  - Compare the rate of redo log accumulation with the size of the redo log files. Use this ratio to see how often to take an incremental backup, in order to avoid the likelihood of the backup failing because the historical data are not available in the redo log. For example, if you are producing 1GB of redo log data per day, and the combined size of your redo log files is 7GB, you would schedule incremental backups more frequently than once a week. You might perform incremental backups every day or two, to avoid a potential issue when a sudden flurry of updates produced more redo log data than usual.
  - Benchmark incremental backup times using both the `--incremental` and `--incremental-with-redo-log-only` options, to confirm if the redo log backup technique performs faster and with less overhead than the traditional incremental backup method. The result could depend on the size of your data, the amount of DML activity, and the size of your redo log files. Do your testing on a server with a realistic data volume and a realistic workload. For example, if you have huge redo log files, reading them in the course of an incremental backup could take as long as reading the InnoDB data files using the traditional incremental technique. Conversely, if your data volume is large, reading all the data files to find the few changed pages could be less efficient than processing the much smaller redo log files.

## Other Considerations for Incremental Backups

The incremental backup feature is primarily intended for InnoDB tables, or non-InnoDB tables that are read-only or rarely updated. Incremental backups detect changes at the level of [pages](#) in the InnoDB [data files](#), as opposed to table rows; each page that has changed is backed up. Thus, the space and time savings are not exactly proportional to the percentage of changed InnoDB rows or columns.

For non-InnoDB files, the entire file is included in an incremental backup if that file has changed since the previous backup, which means the savings for backup resources are less significant when comparing with the case with InnoDB tables.

You cannot perform incremental backups with the `--compress` option.

## Examples of Incremental Backups

This example uses `mysqlbackup` to make an incremental backup of a MySQL server, including all databases and tables. We show two alternatives, one using the `--incremental-base` option and the other using the `--start-lsn` option.

With the `--incremental-base` option, you do not have to keep track of LSN values between one backup and the next. Instead, you can just specify the backup directory of the previous backup (either full or incremental), and `mysqlbackup` figures out the starting point for this backup based on the metadata of the earlier one. Because you need a known set of directory names, you might want to use hardcoded names or generate a sequence of names in your own backup script, rather than using the `--with-timestamp` option.

Note that even if your last backup was a single-file, you can still use `--incremental-base` by specifying for `dir:directory_path` the location of the temporary directory you supplied with the `--backup-dir` option during the full backup:

As an alternative to specifying `--incremental-base=dir:directory_path`, you can tell `mysqlbackup` to query the `end_lsn` value from the last successful **non-TTS backup** as recorded in the `backup_history` table on the server using `--incremental-base=history:last_backup` (this required that the last backup was made with `mysqlbackup` connected to the server).

You can also use the `--start-lsn` option to specify where the incremental backup should start. You have to record the LSN of the previous backup reported by `mysqlbackup` at the end of the backup:

```
mysqlbackup: Was able to parse the log up to lsn 2654255716
```

The number is also recorded in the `meta/backup_variables.txt` file in the folder specified by `--backup-dir` during the backup. Supply then that number to `mysqlbackup` using the `--start-lsn` option. The incremental backup then includes all changes that came *after* the specified LSN.

To create an incremental backup image with the `--start-lsn` option, use the following command, specifying with `--incremental-backup-dir` the backup directory, which, in this case, is a directory for storing the metadata for the backup and some temporary files:

```
$ mysqlbackup --defaults-file=/home/pekka/.my.cnf --incremental \
--start-lsn=2654255716 \
--with-timestamp \
--incremental-backup-dir=/incr-tmp \
--backup-image=/incr-backup/incremental_image.bi
backup-to-image
```

In the following example though, because `--backup-image` does not provide a full path to the image file to be created, the incremental backup image is created under the folder specified by `--incremental-backup-dir`:

```
$ mysqlbackup --defaults-file=/home/pekka/.my.cnf --incremental \
```



```
--start-lsn=2654255716 \
--with-timestamp \
--incremental-backup-dir=/incr-images \
--backup-image=incremental_image1.bi
backup-to-image
```

In the following example, the `--incremental-base=history:last_backup` option is used, given which `mysqlbackup` fetches the LSN of the last successful (non-TTS) full or partial backup from the `mysql.backup_history` table and performs an incremental backup basing on that.

```
$ mysqlbackup --defaults-file=/home/pekka/.my.cnf \
--incremental --incremental-base=history:last_backup \
--backup-dir=/home/pekka/temp_dir \
--backup-image=incremental_image1.bi
backup-to-image
```

*Advanced:* Use the following commands to create incremental directory backups, using the `--incremental-base` and the `--start-lsn` options:

```
$ mysqlbackup --defaults-file=/home/pekka/.my.cnf --incremental \
--incremental-base=dir:/incr-backup/wednesday \
--incremental-backup-dir=/incr-backup/thursday \
backup
```

## Maintaining a backup schedule:

- On a regular schedule determined by date or amount of database activity, take more incremental or differential backups.
- Optionally, periodically start the cycle over again by taking a [full, uncompressed](#) or [compressed](#) backup. Typically, this milestone happens when you can archive and clear out your oldest backup data.

On how to restore your database using the incremental backups, see [Section 5.1.3, “Restoring an Incremental Backup”](#)

## 4.3.4 Making a Compressed Backup

To save disk space, you can compress InnoDB backup data files by using the `--compress` option of `mysqlbackup`. Compression lets you keep more sets of backup data on hand and save transmission time when sending the backup data to another server. The downside includes the extra CPU overhead during the backup itself and the extra time needed for the restoration process for uncompressing the data.

The backup compression feature works only for InnoDB tables. After the InnoDB tablespace files are compressed during backup, they receive the `.ibz` extension. To avoid wasting CPU cycles without saving additional disk space, `--compress` does not attempt to compress tables that were already-compressed on the server (see [Creating Compressed Tables](#)); nevertheless, such tablespace files are also saved with the `.ibz` extension inside the backup.



### Note

When there is unused space within an InnoDB tablespace file, the entire file is copied during an uncompressed backup. Perform a compressed backup to avoid the storage overhead for the unused space.

The binary log and relay log files are compressed and saved with the `.bz` extension when being included in a compressed backup.

You can only use the `--compress` option for [full backups](#), not for [incremental backups](#).

You can also select the compression algorithm to use by the `--compress-method` option and, when using the ZLIB or LZMA compression algorithm, the level of compression by the `--compress-level` option. See [Section 13.6, “Compression Options”](#) for details.

This is a sample command for making a compressed single-file backup:

```
mysqlbackup --defaults-file=/etc/my.cnf --compress --compress-level=5 \
--backup-image=backup.img backup-to-image
```

*Advanced:* This is a sample command for making a compressed backup:

```
mysqlbackup --defaults-file=/etc/my.cnf --compress --compress-level=5 backup
```

### 4.3.5 Making a Partial Backup



#### Note

To facilitate the creation of partial backups, MySQL Enterprise Backup has introduced two new options for partial backup since version 3.10: `--include-tables` and `--exclude-tables`. The new options are intended for replacing the older options of `--include`, `--databases`, `--databases-list-file`, and `--only-innodb-with-frm`, which are incompatible with the new options and will be deprecated in the upcoming releases. In the discussions below we assume the new options are used for partial backups. For reference purpose, we have included information on the older options at the end of this section in [Making a Partial Backup with Legacy Options](#).

By default, all the files under the database subdirectories in the data directory are included in the backup, so that the backup includes data from all MySQL storage engines, any third-party storage engines, and even any non-database files in that directory. This section explains options you can use to selectively back up or exclude data.

There are various ways to create different kinds of partial backup with MySQL Enterprise Backup:

- Including or excluding specific tables by their names. This uses the `--include-tables` or `--exclude-tables` option.

Each table is checked against the regular expression specified with the `--include-tables` or `--exclude-tables` option. If the regular expression matches the fully qualified name of the table (in the form of `db_name.table_name`), the table is included or excluded for the backup. The regular expression syntax used is the extended form specified in the POSIX 1003.2 standard. The options have been implemented with the RE2 regular expression library.

- Including some or all InnoDB tables, but not other table types. This uses the `--only-innodb` option.
- Leaving out files that are present in the MySQL data directory but not actually part of the MySQL instance. This uses the `--only-known-file-types` option.
- Achieving a multiple of selection effects by using a combination of the above mentioned options.
- Backing up a selection of InnoDB tables using [transportable tablespaces \(TTS\)](#). This uses the `--use-tts` and the `--include-tables` or `--exclude-tables` (or both) options.

For syntax details on all the options involved, see [Section 13.8, “Partial Backup and Restore Options”](#).



### Important

Typically, a partial backup is more difficult to restore than a full backup, because the backup data might not include the necessary interrelated pieces to constitute a complete MySQL instance. In particular, InnoDB tables have internal IDs and other data values that can only be restored to the same instance, not a different MySQL server. Always fully test the recovery procedure for any partial backups to understand the relevant procedures and restrictions.



### Important

Because the InnoDB system tablespace holds metadata about InnoDB tables from all databases in an instance, restoring a partial backup on a server that includes other databases could cause the system to lose track of those InnoDB tables in the other databases. Always restore partial backups on a fresh MySQL server instance without any other InnoDB tables that you want to preserve.

The following are some command samples for partial backups.

Including all tables with names starting with “emp” into the backup:

```
$ mysqlbackup \
--host=localhost --user=mysqluser --protocol=TCP --port=3306 \
--backup-dir=$MEB_TEMP_BACKUP_DIR --backup-image=$MEB_BACKUPS_DIR/my.mbi \
--include-tables='\.emp' \
backup-to-image
```

Taking a backup of all tables except tables from the “mysql” and “performance\_schema” databases:

```
$ mysqlbackup \
--host=localhost --user=mysqluser --protocol=TCP --port=3306 \
--backup-dir=$MEB_TEMP_BACKUP_DIR --backup-image=$MEB_BACKUPS_DIR/my.mbi \
--exclude-tables='^(mysql|performance_schema)\.' \
backup-to-image
```

Taking a backup of all tables in the “sales” database, but excludes the table with the name “hardware”

```
$ mysqlbackup \
--host=localhost --user=mysqluser --protocol=TCP --port=3306 \
--backup-dir=$MEB_TEMP_BACKUP_DIR --backup-image=$MEB_BACKUPS_DIR/my.mbi \
--include-tables='^sales\.' --exclude-tables='^sales\.hardware$' \
backup-to-image
```

Taking a backup of all tables in the “sales reps” database, but excludes the table with the name “euro-asia” (special characters like spaces or dashes are supported by the partial backup options since release 3.12.1):

```
$ mysqlbackup \
--host=localhost --user=mysqluser --protocol=TCP --port=3306 \
--backup-dir=$MEB_TEMP_BACKUP_DIR --backup-image=$MEB_BACKUPS_DIR/my.mbi \
--include-tables='^sales reps\.' --exclude-tables='^sales reps\.euro-asia' \
backup-to-image
```

Backing up all InnoDB tables, but not `.frm` files:

```
$ mysqlbackup --defaults-file=/home/pekka/.my.cnf --only-innodb backup-to-image
```

You can also make [compressed](#) and other kinds of selective backups by using the appropriate command options.

## Making a Partial Backup with the Legacy Options



### Important

Information in this subsection is only for using the legacy options of `--include`, `--databases`, `--databases-list-file`, and `--only-innodb-with-frm`, which will be deprecated in the upcoming issues. For creating partial backups, it is strongly recommended that the new options of `--include-tables` and `--exclude-tables` be used instead. Note that you cannot combine the legacy and the new partial-backup options in a single command.

MySQL Enterprise Backup can make different kinds of partial backup using the legacy partial-backup options:

- Including certain InnoDB tables but not others. This operation involves the `--include`, `--only-innodb`, and `--only-innodb-with-frm` options.
- Including certain non-InnoDB tables from selected databases but not others. This operation involves the `--databases` and `--databases-list-file` options.

For syntax details on all these options, see [Legacy Partial Backup Options](#).



### Note

Typically, a partial backup is more difficult to restore than a full backup, because the backup data might not include the necessary interrelated pieces to constitute a complete MySQL instance. In particular, InnoDB tables have internal IDs and other data values that can only be restored to the same instance, not a different MySQL server. Always fully test the recovery procedure for any partial backups to understand the relevant procedures and restrictions.

With its `--include` option, `mysqlbackup` can make a backup that includes some InnoDB tables but not others:

- A partial backup with the `--include` option always contains the InnoDB system tablespace and all the tables inside it.
- For the InnoDB tables stored outside the system tablespace, the partial backup includes only those tables whose names match the regular expression specified with the `--include` option.

This operation requires the tables being left out to be stored in separate `table_name.ibd` files. To put an InnoDB table outside the system tablespace, create it while the `innodb_file_per_table` MySQL configuration option is enabled. Each `.ibd` file holds the data and indexes of one table only.

Those InnoDB tables created with `innodb_file_per_table` turned off are stored as usual in the InnoDB [system tablespace](#), and cannot be left out of the backup.

For each table with a per-table data file a string of the form `db_name.table_name` is checked against the regular expression specified with the `--include` option. If the regular expression matches the complete string `db_name.table_name`, the table is included in the backup. The regular expression syntax used is the extended form specified in the POSIX 1003.2 standard. On Unix-like systems, quote the regular expression appropriately to prevent interpretation of shell meta-characters. This feature has been implemented with the RE2 regular expression library.

The backup directory produced contains a backup log file and copies of InnoDB data files.

**IMPORTANT:** Although `mysqlbackup` supports taking partial backups, be careful when restoring a database from a partial backup. `mysqlbackup` copies also the `.frm` files of those tables that are not included in the backup, except when you do partial backups using, for example, the `--databases` option. If you use `mysqlbackup` with the `--include` option, before restoring the database, delete from the backup data the `.frm` files for any tables that are not included in the backup.

**IMPORTANT:** Because the InnoDB system tablespace holds metadata about InnoDB tables from all databases in an instance, restoring a partial backup on a server that includes other databases could cause the system to lose track of those InnoDB tables in other databases. Always restore partial backups on a fresh MySQL server instance without any other InnoDB tables that you want to preserve.

The `--only-innodb` and `--only-innodb-with-frm` options back up InnoDB tables only, skipping those of other storage engines. You might also use them together with the `--include` option to make selective backup of InnoDB tables while excluding all other files created by other storage engines.

#### Example 4.19 Making an Uncompressed Partial Backup of InnoDB Tables

In this example, we have configured MySQL so that some InnoDB tables have their own tablespaces. We make a partial backup including only those InnoDB tables in `test` database whose name starts with `ib`. The contents of the database directory for `test` database are shown below. The directory contains a MySQL description file (`.frm` file) for each of the tables (`alex1`, `alex2`, `alex3`, `blobt3`, `ibstest0`, `ibstest09`, `ibtest11a`, `ibtest11b`, `ibtest11c`, and `ibtest11d`) in the database. Of these 10 tables six (`alex1`, `alex2`, `alex3`, `blobt3`, `ibstest0`, `ibstest09`) are stored in per-table data files (`.ibd` files).

```
$ ls /sqldata/mts/test
alex1.frm  alex2.ibd  blobt3.frm  ibstest0.ibd  ibtest11a.frm  ibtest11d.frm
alex1.ibd  alex3.frm  blobt3.ibd  ibtest09.frm  ibtest11b.frm
alex2.frm  alex3.ibd  ibstest0.frm  ibtest09.ibd  ibtest11c.frm
```

We run the `mysqlbackup` with the `--include` option:

```
# Back up some InnoDB tables but not any .frm files.
$ mysqlbackup --defaults-file=/home/pekka/.my.cnf --include='test\ib.*' --only-innodb backup
...many lines of output...
mysqlbackup: Scanned log up to lsn 2666737471.
mysqlbackup: Was able to parse the log up to lsn 2666737471.
mysqlbackup: Maximum page number for a log record 0
101208 17:17:45 mysqlbackup: Full backup completed!

# Back up some InnoDB tables and the .frm files for the backed-up tables only.
$ mysqlbackup --defaults-file=/home/pekka/.my.cnf --include='test\ib.*' \
  --only-innodb-with-frm=related backup
...many lines of output...
mysqlbackup: Scanned log up to lsn 2666737471.
mysqlbackup: Was able to parse the log up to lsn 2666737471.
mysqlbackup: Maximum page number for a log record 0
101208 17:17:45 mysqlbackup: Full backup completed!
```

The backup directory contains only backups of `ibstest` and `ibtest09` tables. Other InnoDB tables did not match the include pattern `test\ib.*`. Notice, however, that the tables `ibtest11a`, `ibtest11b`, `ibtest11c`, `ibtest11d` are in the backup even though they are not visible in the directory shown below, because they are stored in the system tablespace (`ibdata1` file) which is always included in the backup.

```
# With the --only-innodb option:
```

```
$ ls /sqldata-backup/test
ibstest0.ibd  ibtest09.ibd

# With the --only-innodb-with-frm=related option:
$ ls /sqldata-backup/test
ibstest0.frm  ibtest09.frm
ibstest0.ibd  ibtest09.ibd
```

### Example 4.20 Making a Compressed Partial Backup

We have configured MySQL so that every InnoDB table has its own tablespace. We make a partial backup including only those InnoDB tables whose name starts with `alex` or `blob`. The contents of the database directory for `test` database is shown below.

```
$ ls /sqldata/mts/test
alex1.frm  alex2.ibd  blobt3.frm  ibstest0.ibd  ibtest11a.frm  ibtest11d.frm
alex1.ibd  alex3.frm  blobt3.ibd  ibtest09.frm  ibtest11b.frm
alex2.frm  alex3.ibd  ibstest0.frm  ibtest09.ibd  ibtest11c.frm
```

We run `mysqlbackup` with the `--compress` and `--include` options:

```
$ mysqlbackup --defaults-file=/home/pekka/.my.cnf --compress \
  --include='.*\.(alex|blob).*' --only-innodb backup
...many lines of output...
mysqlbackup: Scanned log up to lsn 2666737471.
mysqlbackup: Was able to parse the log up to lsn 2666737471.
mysqlbackup: Maximum page number for a log record 0

mysqlbackup: Compressed 147 MB of data files to 15 MB (compression 89%).

101208 17:18:04  mysqlbackup: Full backup completed!
```

The backup directory for the database `test` is shown below. The `.ibz` files are compressed per-table data files.

```
$ ls /sqldata-backup/test
alex1.ibz  alex2.ibz  alex3.ibz  blobt3.ibz
```

The `--databases` and `--databases-list-file` options of `mysqlbackup` let you back up non-InnoDB tables only from selected databases, rather than across the entire MySQL instance. (To filter InnoDB tables, use the `--include` option instead.) With `--databases`, you specify a space-separated list of database names, with the entire list enclosed in double quotation marks. With `--databases-list-file`, you specify the path of a file containing the list of database names, one per line.

Some or all of the database names can be qualified with table names, to only back up selected non-InnoDB tables from those databases.

If you specify this option, make sure you include the same set of databases for every backup (especially incremental backups), so that you do not restore out-of-date versions of any databases.

## 4.3.6 Making an Optimistic Backup

Optimistic backup is a feature introduced in MySQL Enterprise Backup 3.11 for improving performance for backing up and restoring huge databases in which only a small number of tables are modified frequently.

During a hot backup of a huge database (say, in the order of terabytes), huge redo log files could be generated on the server when the backup is in progress. As the redo log files grow faster than they can be

processed by `mysqlbackup`, the backup operation can actually fail when `mysqlbackup` cannot catch up with the redo log cycles and LSNs get overwritten by the server before they are read by `mysqlbackup`. Moreover, the `apply-log` step for [preparing a backup for restoration](#) can take a very long time as `mysqlbackup` has huge `ibbackup_logfile` files (created from the big redo log files) to apply to the backup. The problems are intensified when the I/O resources available for reading and writing the redo logs are scarce during the backup and restoration processes.

Optimistic backup relieves the problems by dividing the backup process into two internal phases, which are transparent to the users:

1. Optimistic phase: In this first phase, tables that are unlikely to be modified during the backup process (referred to as the “inactive tables” below, identified by the user with the `optimistic-time` option or, by exclusion, with the `optimistic-busy-tables` option) are backed up without locking the MySQL instance. And because those tables are not expected to be changed before the backup is finished, redo logs, undo logs, and system table spaces are not backed up by `mysqlbackup` in this phase.
2. Normal phase: In this second phase, tables that are not backed up in the first phase (referred to as the “busy tables” below) are being backed up in a manner similar to how they are processed in an ordinary backup: the InnoDB files are copied first, and then other relevant files and copied or processed with the MySQL instance locked. The redo logs, undo logs, and the system tablespace are also backed up in this phase.

An optimistic backup occurs whenever the `optimistic-time` or `optimistic-busy-tables` option is used. For how to use the options, see detailed descriptions for them in [Section 13.10, “Performance / Scalability / Capacity Options”](#). If, as expected, the inactive tables identified by the options are not changed a lot during the backup, the overall backup time (time for the backup operation plus time for the `apply-log` operation) can be reduced significantly comparing with an ordinary backup because the `apply-log` operation will be much faster. However, if it turns out that the inactive tables identified are modified a lot during the backup process, benefits of performing an optimistic back up will become limited and, in the worst case, an optimistic backup might actually take longer to perform and, for a single-file backup, the size of the backup will be larger when comparing with an ordinary backup. Therefore, users should be careful in identifying which tables are “inactive” and which are “busy” when trying to perform an optimistic backup.



#### Note

An optimistic backup cannot be performed for an incremental backup or a backup using [transportable tablespaces \(TTS\)](#).

The following examples illustrate how to make an optimistic backup.

#### Example 4.21 Optimistic Backup Using the Option `optimistic-time=YYMMDDHHMMSS`

In this example, tables that have been modified since the noon of May 16, 2011 are treated as busy tables and backed up in the normal phase of an optimistic backup, and all other tables are backed up in the optimistic phase:

```
mysqlbackup --defaults-file=/etc/my.cnf --optimistic-time=110516120000 backup-to-image
```

#### Example 4.22 Optimistic Backup Using the Option `optimistic-time=now`

In this example, all tables are treated as inactive tables and backed up in the optimistic phase of an optimistic backup:

```
mysqlbackup --defaults-file=/etc/my.cnf --optimistic-time=now backup-to-image
```



### Example 4.23 Optimistic Backup Using the `optimistic-busy-tables` Option

In this example, tables in `mydatabase` that are suffixed by `mytables-` in their names are treated as busy tables and backed up in the normal phase of an optimistic backup, and all other tables are backed up in the optimistic phase:

```
mysqlbackup --defaults-file=/etc/my.cnf --optimistic-busy-tables='mydatabase\mytables-.*' backup
```

When you use both the `optimistic-time` and `optimistic-busy-tables` options and they come into conflict on determining which tables are to be busy tables, `optimistic-busy-tables` takes precedence over `optimistic-time`. For example:

### Example 4.24 Optimistic and Partial Backup Using both the `optimistic-busy-tables` and `optimistic-time` Options

In this example, tables in `mydatabase` that are suffixed by `mytables-` in their names are treated as busy tables and backed up in the normal phase, even if they have not been modified since May 16, 2010, the time specified by `optimistic-time`:

```
mysqlbackup --defaults-file=/etc/my.cnf --optimistic-busy-tables='mydatabase\mytables-.*' \
--optimistic-time=100516 backup
```

## 4.3.7 Making a Back Up of In-Memory Database Data

The `--exec-when-locked` option of `mysqlbackup` lets you specify a command (together with the desired command arguments) to run near the end of the backup while the database is still locked. This command can copy or create additional files in the backup directory. For example, you can use this option to back up `MEMORY` tables with the `mysqldump` command, storing the output in the backup directory. To delay any redirection or variable substitution until the command is executed, enclose the entire option value within single quotes.

## 4.3.8 Making Scheduled Backups

Maintaining a regular backup schedule is an important measure for preventing data loss for your MySQL server. This section discusses some simple means for setting up a schedule for running MySQL Enterprise Backup.

*For Linux and other Unix-like platforms:* you can set up a cron job on your system for scheduled backups. There are two types of cron jobs: to set up a user cron job, which is owned and run by a particular user, do the following:

- Log on as the user who runs MySQL Enterprise Backup and use the following command to invoke an editor for creating (or modifying) a crontab:

```
shell> crontab -e
```

- In the editor, add an entry similar to the following one to the crontab, and then save your changes (make sure contents in both lines below appear in a single line in the crontab):

```
@daily /path-to-mysqlbackup/mysqlbackup -uroot --backup-dir=/path-to-backup-folder/cronbackups
--with-timestamp --backup-image=my.mib backup-to-image &>/dev/null
```

This crontab entry invokes `mysqlbackup` to create a backup under the `cronbackups` directory at `00:00:00` everyday. Outputs from the `stderr` and `stdout` streams are redirected to `/dev/null`, so they will not invoke other actions on the part of the Cron server (for example, email notifications to the user).



To set up a system cron job, which is owned and run by `root`, create a file under the `/etc/cron.d` folder and put into it a similar crontab entry as the one above, adding the user (`root` in the following example) before the `mysqlbackup` command:

```
@daily root /path-to-mysqlbackup/mysqlbackup -uroot --backup-=/path-to-backup-folder/cronbackups \
--with-timestamp --backup-image=my.mib backup-to-image &>/dev/null
```

Check your platform's documentation on further details on the different ways to set up cron jobs for various types of schedules.

*For Windows platforms:* Use the Task Scheduler for the purpose. Check the documentation for you Windows platform for instructions.

## 4.4 Making Backups with a Distributed File System (DFS) or Storage Access Network (SAN)

When system administrators attempt to set up MySQL and MySQL Enterprise Backup in an environment that uses a distributed file system (DFS) or a storage access network (SAN), the MySQL server, the server's data directory, MySQL Enterprise Backup, and the backup directory may end up existing on different physical servers. When that happens, the operations of `mysqlbackup` might be impacted. The operation most likely to be adversely affected is hot backup, the success of which depends on:

1. Each page of a data file is copied consistently, that is, all the bytes in the page correspond to the same LSN.
2. No copied page is older than the time that marks the beginning of the temporal duration the backup is supposed to cover.
3. The redo log is copied consistently, meaning a continuous segment of redo log is copied, and it includes all the changes from the beginning of the temporal period that the backup is to cover until the end of the backup operation. Each block of the copied redo log has to be consistent.

Condition 1 is easily achievable with most DFSs or SANs of reasonable performance. Condition 2 though can remain unfulfilled even when condition 1 has been satisfied: for example, `mysqlbackup` could copy all the pages of a tablespace correctly except for one page for which `mysqlbackup` has included an old version into the copy. If the LSN of that old version of the page is smaller than the LSN first seen by `mysqlbackup` at the beginning of the backup process, the resulting backup will be defective. This example shows that `mysqlbackup` may have problem performing a hot backup unless it can see the writes to the file system being executed in the correct order, that is, the order in which the server executed them.

Regarding condition 3, unlike data file pages, redo log blocks are written sequentially, which means condition 3 is easier to fulfill than conditions 1 and 2. However, if `mysqlbackup` reaches the highest LSN in the copied data file pages before encountering the end of the redo log, the backup fails. A failure occurs also if `mysqlbackup` reads a corrupted log block at any time during the copying of the redo log. Both these failures can occur if `mysqlbackup` does not see the same history of the file system states as the MySQL server does.

Therefore, to use `mysqlbackup` with a DFS or SAN, it is important to make sure that `mysqlbackup` sees all the writes to the file system in the same order as the MySQL server does. The condition is most likely to be satisfied when `mysqlbackup` and the MySQL server are running on the same server node, and it is unlikely to be always fulfilled when it is otherwise.



---

# Chapter 5 Recovering or Restoring a Database

## Table of Contents

5.1 Performing a Restore Operation .....	55
5.1.1 Restoring a Compressed Backup .....	56
5.1.2 Restoring an Encrypted Backup Image .....	56
5.1.3 Restoring an Incremental Backup .....	57
5.1.4 Restoring Backups Created with the <code>--use-tts</code> Option .....	58
5.1.5 Restoring a Backup from Cloud Storage to a MySQL Server .....	58
5.1.6 Advanced: Restoring a Directory Backup .....	59
5.2 Point-in-Time Recovery from a Hot Backup .....	60
5.3 Restoring a Backup with a Database Upgrade or Downgrade .....	61

The ultimate purpose of backup data is to help recover from a database issue or to create a clone of the original database in another location (typically, to run report queries or to create a new replication slave). This section describes the procedures to handle those scenarios.

After a serious database issue, you might need to perform a recovery under severe time pressure. It is critical to confirm in advance:

- How long the recovery will take, including any steps to transfer, unpack, and otherwise process the data.
- That you have practiced and documented all steps of the recovery process, so that you can do it correctly in one try. If a hardware issue requires restoring the data to a different server, verify all privileges, storage capacity, and so on, on that server ahead of time.
- That you have periodically verified the accuracy and completeness of the backup data, so that the system will be up and running soon after being recovered.

## 5.1 Performing a Restore Operation

The `mysqlbackup` commands to perform a restore operation are `copy-back-and-apply-log` and `copy-back` (for directory backup only; see [Section 5.1.6, “Advanced: Restoring a Directory Backup”](#)). Normally, the restoration process requires the database server to be already shut down (or, at least not operating on the directory you are restoring the data to), except for restorations of backups created with the `--use-tts` option; see [explanations](#) below. The process copies the data files, logs, and other backed-up files from the backup directory back to their original locations, and performs any required post-processing on them. For any restore operation, the options `datadir`, `innodb_log_files_in_group`, `innodb_log_file_size`, and `innodb_data_file_path` must be specified either in the target server's configuration file, in the file specified by the `--defaults-file` option, or as command-line options.

### Example 5.1 Restoring a Database

```
mysqlbackup --defaults-file=<backupDir>/backup-my.cnf -uroot --backup-image=<image_name> \  
--backup-dir=<backupTmpDir> --datadir=<restoreDir> copy-back-and-apply-log
```

The `copy-back-and-apply-log` command achieves two things:

- Extracts the backup from the image file (the `--backup-dir` option specifies the temporary directory into which the backup image will be extracted) and copies it to the data directory on the server to be restored.

- Performs an [apply log operation](#) to the restored data to bring them up-to-date.

The restored data includes the [backup\\_history](#) table, where MySQL Enterprise Backup records details of each backup. The table allows you to perform future incremental backups using the `--incremental-base=history:last_backup` option.



#### Important

- When performing a full restore (for example, when the backup data is used to set up a new MySQL server or used to replace all data of an existing MySQL server), make sure the target data directories are all clean, containing no old or unwanted data files (this might require manual removal of files at the locations specified by both the `--datadir` and `--innodb_data_file_path` options); otherwise, add the `--force [105]` option to the restore command to overwrite the old data. The same cleanup is not required for restoring backups created with the `--use-tts` option (in which case other requirements described in [Section 5.1.4, “Restoring Backups Created with the --use-tts Option”](#) apply though), and usually not necessary for restoring a partial backup.
- Due to a known issue (Bug# 20485910), restoring a partial image backup created with MySQL Enterprise Backup 3.11 or earlier requires using the `--force` option.

The following subsections describes a number of different scenarios for restoring a backup.

### 5.1.1 Restoring a Compressed Backup

Restore a compressed backup image named `<image_name>`, using the `--backup-dir` option to specify the temporary directory into which the image will be extracted:

```
mysqlbackup --defaults-file=<backupDir>/backup-my.cnf -uroot --backup-image=<image_name> \
--backup-dir=<backupTmpDir> --datadir=<restoreDir> --uncompress copy-back-and-apply-log
```

*Advanced:* Do the same for a compressed directory backup at `<backupDir>` to `<restoreDir>` on the server using `copy-back-and-apply-log`:

#### Example 5.2 Restoring a Compressed Backup

```
mysqlbackup --defaults-file=<my.cnf> -uroot --backup-dir=<backupDir> --datadir=<restoreDir> \
--uncompress copy-back-and-apply-log
```

See [Section 4.3.4, “Making a Compressed Backup”](#) and [Section 13.6, “Compression Options”](#) for more details on compressed backups.

### 5.1.2 Restoring an Encrypted Backup Image

Restore an encrypted backup image named `<image_name>` to `<restoreDir>` on the server with `copy-back-and-apply-log`, using the encryption key contained in a file named `<keyFile>`:

#### Example 5.3 Restoring an Encrypted Backup Image

```
mysqlbackup --defaults-file=<backupDir>/backup-my.cnf --backup-image=<image_name> \
--backup-dir=<backupTmpDir> --datadir=<restoreDir> --decrypt --key-file=<keyFile> copy-back-and-apply-log
```

See [Section 13.13, “Encryption Options”](#) for more details on backup encryption and decryption.

### 5.1.3 Restoring an Incremental Backup

There are different ways to use incremental backups to restore a database under different scenarios. The preferred method is to first restore the full backup and make it up-to-date to the time at which the full backup was performed using the `copy-back-and-apply-log` command (see [Example 5.1, “Restoring a Database”](#) on how to do it), then use `copy-back-and-apply-log` again to restore the incremental backup image on top of the full backup that was just restored:

#### Example 5.4 Restoring an Incremental Backup Image

```
mysqlbackup --defaults-file=<backupDir>/backup-my.cnf -uroot --backup-image=<inc_image_name> \
--incremental-backup-dir=<incBackupTmpDir> --datadir=<restoreDir> --incremental \
copy-back-and-apply-log
```

In this example, the incremental backup image named `<inc_image_name>` was restored to `<restoreDir>` on the server (where the full backup that the incremental backup image was based on has already been restored). The `--incremental-backup-dir` option is used to specify the temporary directory into which the image will be extracted (you can use `--backup-dir` for the same purpose). Repeat the step with other incremental backup images that you have, until the data has been restored to a desired point in time.

### Advanced: Restoring an Incremental Backup Directory

Incremental directory backups can be restored in a series of `copy-back-and-apply-log` command, as illustrated above for single-file backups. Alternatively, at anytime after an incremental backup is taken and before the data is restored, you can bring your full backup up-to-date with your incremental backup. First, apply to the full backup any changes that occurred while the backup was running:

```
$ mysqlbackup --backup-dir=/full-backup/2010-12-08_17-14-11 apply-log
..many lines of output...
101208 17:15:10 mysqlbackup: Full backup prepared for recovery successfully!

101208 17:15:10 mysqlbackup: mysqlbackup completed OK!
```

Then, we apply the changes from the incremental backup using the `apply-incremental-backup` command:

```
$ mysqlbackup --incremental-backup-dir=/incr-backup/2010-12-08_17-14-48
--backup-dir=/full-backup/2010-12-08_17-14-11 apply-incremental-backup
..many lines of output...
101208 17:15:12 mysqlbackup: mysqlbackup completed OK!
```

Now, the data files in the full backup directory are fully up-to-date, as of the time of the last incremental backup. You can keep updating it with more incremental backups, so it is ready to be restored anytime.

### Binary Log and Relay Log Restore

When an incremental backup is being restored using either the `copy-back-and-apply-log` or `apply-incremental-backup` command, the binary log (and also the relay log, in the case of a slave server), if included in the incremental backup, is also restored to the target server by default. This default behavior is overridden when either (1) the `--skip-binlog` option (or the `--skip-relaylog` option for the relay log) is used with the restore command, or (2) if the full backup the incremental backup was based on or any prior incremental backup that came in between the full backup and this incremental backup has the binary log (or relay log) missing.

See [Section 4.3.3, “Making a Differential or Incremental Backup”](#), and [Section 13.7, “Incremental Backup Options”](#), for more details on incremental backups.

### 5.1.4 Restoring Backups Created with the `--use-tts` Option

There are some special requirements for restoring backups created with [transportable tablespaces \(TTS\)](#) (that is, created with the `--use-tts` option):

- The destination server must be running.
- Make sure that the required parameters for connecting to the server (port number, socket name, etc.) are provided as command-line options for `mysqlbackup`, or are specified in the `[client]` section of a default file.
- The destination server must be using the same page size that was used on the server on which the backup was made.
- The `innodb_file_per_table` option must be enabled on the destination server.
- The tables being restored must not exist on the destination server.
- A restore fails if the InnoDB file format of a per-table data file (`.ibd` file) to be restored does not match the value of the `innodb_file_format` system variable on the destination server. In that case, use the `--force [105]` option with the restore commands to change temporarily the value of `innodb_file_format` on the server, in order to allow restores of per-table data files regardless of their format.

Selected tables can be restored from a backup created with [transportable tablespaces \(TTS\)](#) using the `--include-tables` and `--exclude-tables` options. The following command restores all tables in the “sales” database from the backup, but excludes the table with the name “hardware”:

#### Example 5.5 Restoring Selected Tables from a TTS Image Backup

```
mysqlbackup --socket=/tmp/restoreserver.sock --datadir=/logs/restoreserverdata --backup-dir=/logs/backup \
--include-tables='^sales\.' --exclude-tables='^sales\.hardware$' copy-back-and-apply-log
```

This following commands rename a table when restoring it from a TTS Backup by using the `--rename` option:

#### Example 5.6 Restoring and Renaming a Table from a TTS Backup

```
# Using fully qualified table names:
mysqlbackup --socket=/tmp/restoreserver.sock --datadir=/logs/restoreserverdata \
--backup-dir=/BackupDirTemp --backup-image=/home/dbadmin/backups/tts-backup.mbi \
--include-tables='^sales\.cars' --rename="sales.cars to sales.autos" copy-back-and-apply-log

# It works the same if database names are omitted in the argument for --rename:
mysqlbackup --socket=/tmp/restoreserver.sock --datadir=/logs/restoreserverdata \
--backup-dir=/BackupDirTemp --backup-image=/home/dbadmin/backups/tts-backup.mbi \
--include-tables='^sales\.cars' --rename="cars to autos" copy-back-and-apply-log
```

### 5.1.5 Restoring a Backup from Cloud Storage to a MySQL Server

To restore a backup image from cloud storage to `datadir` on the server, use the [cloud storage options](#), and the `--backup-dir` option to specify the temporary directory into which the image will be extracted:

**Example 5.7 Restoring a Single-file Backup from Amazon S3 to a MySQL Server**

```
mysqlbackup \
--defaults-file=/bkups/backupdir/backupmy.cnf \
--cloud-service=s3 --cloud-aws-region=<aws region> \
--cloud-access-key-id=<aws access key id> --cloud-secret-access-key=<aws secret access key> \
--cloud-bucket=<s3 bucket name> --cloud-object-key=<aws object key> \
--backup-dir=/home/user/dba/s3backuptmpdir \
--datadir=/home/user/dba/datadir \
--backup-image=- \
copy-back-and-apply-log
```

**Example 5.8 Restoring a Single-file Backup from an OpenStack Object Storage to a MySQL Server**

```
mysqlbackup \
--defaults-file=/bkups/backupdir/backupmy.cnf \
--cloud-service=openstack --cloud-container=<swift container> \
--cloud-user-id=<keystone user> --cloud-password=<keystone password> \
--cloud-region=<keystone region> --cloud-tenant=<keystone tenant> \
--cloud-identity-url=<keystone url> --cloud-object=image_800.mbi \
--backup-dir=/home/user/dba/swiftbackuptmpdir \
--datadir=/home/user/dba/datadir \
--backup-image=- \
copy-back-and-apply-log
```

## 5.1.6 Advanced: Restoring a Directory Backup

A directory backup, just like a single-file backup, can be restored using the `copy-back-and-apply-log` command as explained at the beginning of [Section 5.1, “Performing a Restore Operation”](#).

**Example 5.9 Restoring a Backup Directory using `copy-back-and-apply-log`**

```
mysqlbackup --defaults-file=/usr/local/mysql/my.cnf \
--backup-dir=/export/backups/full \
copy-back-and-apply-log
```

However, two alternatives exist for directory backups:

- Perform the [apply log operation](#) on the [raw backup](#) right after the backup, or anytime before restore, using the `apply-log` command. You can run this step on the same database server where you did the backup, or transfer the raw backup files to a different system first, to limit the CPU and storage overhead on the database server. Here are some examples of doing that, on different kinds of directory backups:

**Example 5.10 Applying the Log to a Backup**

This example runs `mysqlbackup` to roll forward the data files so that the data is ready to be restored:

```
mysqlbackup --backup-dir=/export/backups/2011-06-21__8-36-58 apply-log
```

That command creates InnoDB log files (`ib_logfile*`) within the backup directory and applies log records to the InnoDB data files (`ibdata*` and `*.ibd`).

**Example 5.11 Applying the Log to a Compressed Backup**

If the backup is compressed, as in [Section 4.3.4, “Making a Compressed Backup”](#), specify the `--uncompress` option to `mysqlbackup` when applying the log to the backup:

```
mysqlbackup --backup-dir=/export/backups/compressed --uncompress apply-log
```



### Note

Since the [apply log operation](#) does not modify any of the original files in the backup, nothing is lost if the operation fails for some reason (for example, insufficient disk space). After fixing the problem, you can safely retry [apply-log](#) and by specifying the [--force \[105\]](#) option, which allows the data and log files created by the failed [apply log operation](#) to be overwritten.

- For simple backups (which are not compressed and non-incremental), you can combine the initial backup and the [apply-log](#) steps using the [backup-and-apply-log](#) command.

After the backup has been prepared, you can now restore it using the [copy-back](#) command:

```
mysqlbackup --defaults-file=/usr/local/mysql/my.cnf \
--backup-dir=/export/backups/full \
copy-back
```

## 5.2 Point-in-Time Recovery from a Hot Backup

Using MySQL Enterprise Backup and the [binary log \[11\]](#) files included by default in the incremental backups it creates (except when they are created with the [--use-tts](#) option), you can restore your database to its state at an arbitrary time  $t_R$  in between a full backup and an incremental backup or in between two incremental backups. To do so:

1. Binary logging must be enabled in MySQL before taking the full backup that serves as the base for this restore operation.
2. Using the full and incremental backups that were created before  $t_R$ , restore the database to a time as close to  $t_R$  as possible and *note that time* (let us call it  $t_{LB}$ ). Make sure the restored database is in a consistent state. This can be achieved by, for example, using the [copy-back-and-apply-log](#) command for backup restore.
3. Roll forward the database to its state at the desired time  $t_R$  using the [binary log \[11\]](#) file[s] included in the incremental backup that covers  $t_R$ . Use the [mysqlbinlog](#) utility to extract from the [binary log \[11\]](#) file[s] all the SQL activities that happened in between the  $t_{LB}$  (which you noted in step 2 above) and  $t_R$ , specifying those times with the [--start-datetime](#) and [--stop-datetime](#) options, and pipe the output to the [mysql](#) client, to be replayed on the server:

```
mysqlbinlog --start-datetime=" $t_{LB}$ " \
--stop-datetime=" $t_R$ " \
binlog.000005 binlog.000006 binlog.000007 | mysql -u root -p
```

An alternative is to identify the beginning and endpoint of the extraction not by the start and stop times, but by their corresponding binary log positions:

```
mysqlbinlog --start-position="binary-log-position-corresponding-to- \$t\_{LB}\$ " \
--stop-position="binary-log-position-corresponding-to- \$t\_R\$ " \
binlog.000005 binlog.000006 binlog.000007 | mysql -u root -p
```

Note that you need to pipe all the binary log files in a single connection to the server. You can also dump all the SQL activities to a single file first, and then pipe or play the file to the [mysql](#) client.

For more tips on using the binary log for point-in-time recovery, see [Point-in-Time \(Incremental\) Recovery Using the Binary Log](#).



## 5.3 Restoring a Backup with a Database Upgrade or Downgrade



### Note

Currently, MySQL Enterprise Backup 4.0 only supports MySQL 5.7.9. Upgrade and downgrade information in this section is not applicable to MySQL Enterprise Backup 4.0 until another GA versions of MySQL 5.7 is released later, or until MySQL Enterprise Backup 4.0 starts supporting MySQL 5.6 and earlier.

You can facilitate a server upgrade or downgrade by using MySQL Enterprise Backup to make a backup of your server data, restore it to a data directory, and start a different version of MySQL server on the old data. However, you may encounter technical challenges during a server upgrade or downgrade, and it is beyond the function of MySQL Enterprise Backup, as a backup tool, to ensure a successful server upgrade or downgrade. Users interested in the topic are advised to consult the MySQL server manual, especially the [Upgrading MySQL](#) and [Downgrading MySQL](#) sections.

Here are a number of things that users should pay attention to when restoring a backup with a database upgrade or downgrade:

- Restoring a database to an older MySQL version (i.e., server downgrading) should only be performed when the original and the final server versions are in the same release series (for example, going from 5.5.30 to 5.5.29). Downgrading to a lower series (for example, from 5.6.33 to 5.5.33) might cause server crashes or data corruption.
- Restoring a database to a newer MySQL version (i.e. server upgrading) requires the following steps, the skipping of which might crash the target server:
  - Back up the data on the older server.
  - Restore the data to the directory you plan to use as the newer server's data directory by running a [copy-back-and-apply-log](#) operation on the backup.
  - Restart the older server, using the intended data directory for the newer server as its own data directory. With this step, the older server prepares the data for an upgrade by performing some clean-up steps on the data, similar to what the server would do during a [crash recovery](#).
  - Perform a slow shutdown for the older server by, issuing the [SET GLOBAL innodb\\_fast\\_shutdown=0](#) statement and then shutting the server down. This will ensure that all dirty pages are flushed, and hence there will be no redo log processing for the newer server.
  - Install the newer server.
  - Start the newer server.
  - Run [upgrade steps](#) as documented in the MySQL reference manual on the restored data. Make sure the [mysql\\_upgrade](#) that comes with your newer server is applied.

After performing these steps, check your data to make sure that your restore is successful.



---

# Chapter 6 Using MySQL Enterprise Backup with Replication

## Table of Contents

6.1 Setting Up a New Replication Slave .....	63
6.2 Backing up and Restoring a Slave Database .....	64
6.3 Restoring a Master Database .....	65

Backup and restore operations are especially important in systems that use MySQL replication to synchronize data across a master server and a set of slave servers. In a replication configuration, MySQL Enterprise Backup helps you manage images for the entire system, set up new slave servers, or restore a master server in an efficient way that avoids unnecessary work for the slave servers. On the other hand, having multiple slave servers to choose from gives you more flexibility about where to perform backups. When the binary log is enabled, you have more flexibility about restoring the database to a specific point in time, even a time that is later than that of the last backup.

## 6.1 Setting Up a New Replication Slave

MySQL Enterprise Backup allows you to set up a replication slave by backing up the master and restoring the backup on a new slave server, without having to stop the master.

*For servers NOT using GTID:*

1. Take a backup of the master and then use, for example, the `copy-back-and-apply-log` command, to restore the backup and the log files to the right directories on the new slave and prepare the data.
2. Edit the `my.cnf` file of the new slave and put `skip-slave-start` and `event_scheduler=off` under the `[mysqld]` section.
3. Start the new slave `mysqld` (version `>= 5.1`). It prints the latest MySQL binary log position the backup knows of.

```
...
InnoDB: Last MySQL binlog file position 0 128760128, file name ./hundin-bin.006
...
```

Note that InnoDB only stores the binary log position information to its tablespace at a transaction commit. **To make InnoDB aware of the current binary log position, you must run at least one transaction while binary logging is enabled.**

4. Use the `CHANGE MASTER TO` SQL command on the slave to initialize it properly. For example:

```
CHANGE MASTER TO
MASTER_LOG_FILE='hundin-bin.006',
MASTER_LOG_POS=128760128;
```

5. Set the statuses of any events that were copied from the master to `SLAVESIDE_DISABLED`. For example:

```
mysql> UPDATE mysql.event SET status = 'SLAVESIDE_DISABLED';
```

6. Remove the line `skip-slave-start` and `event_scheduler=off` entries you added to the `my.cnf` file of the slave in step 2.

7. Restart the slave server. Replication starts.

*For servers using GTIDs (supported by MySQL Server 5.6 and after):*

1. Take a backup of the master and then use, for example, the `copy-back-and-apply-log` command, to restore the backup and the log files to the right directories on a new GTID-enabled slave and prepare the data.
2. Edit the `my.cnf` file of the new slave and put `skip-slave-start`, `event_scheduler=off`, and `sql_log_bin=0` under the `[mysqld]` section.
3. Start the new slave server.
4. Connect to the slave server with the `mysql` client.
5. When a server using the GTID feature is backed up, `mysqlbackup` produces a file named `backup_gtid_executed.sql` inside the `meta` folder under the backup directory. The file contains a SQL statement that sets the `GTID_PURGED` configuration option on the slave:

```
# On a new slave, issue the following command if GTIDs are enabled:
SET @@GLOBAL.GTID_PURGED='f65db8e2-0e1a-11e5-a980-080027755380:1-3';
```

It also contains a commented-out `CHANGE MASTER` command for initializing the slave. Uncomment the command and add any needed connection and authentication parameters to it; for example:

```
# Use the following command if you want to use the GTID handshake protocol:
CHANGE MASTER TO MASTER_HOST="127.0.0.1", MASTER_USER="muser", MASTER_PASSWORD="mpass", MASTER_PORT=18675,
```

Execute the file with the `mysql` client

```
mysql> source /path-to-backup_gtid_executed.sql/backup_gtid_executed.sql
```

6. Set the statuses of any events that were copied from the master to `SLAVESIDE_DISABLED`. For example:
- ```
mysql> UPDATE mysql.event SET status = 'SLAVESIDE_DISABLED';
```
7. Remove the `skip-slave-start`, `event_scheduler=off`, and `sql_log_bin=0` entries you added to the `my.cnf` file of the slave in step 2.
  8. Restart the slave server. Replication starts.

For more information on the GTIDs, see [GTID feature](#).

## 6.2 Backing up and Restoring a Slave Database

To backup a slave database, add the `--slave-info` option to your backup command.

To restore the backup on a slave server, follow the same steps outlined in [Section 6.1, “Setting Up a New Replication Slave”](#).



### Warning

In a statement-based replication (SBR) setting (see [Replication Formats](#) for details), whenever the SQL thread for the slave is stopped (for example, during a slave shutdown that has to occur when you restore a backup to a slave server), all temporary tables that are open then will not get replicated on the slave (see

[Replication and Temporary Tables](#) for details). That means it is not always safe to backup a slave server in an SBR setting with `mysqlbackup` if temporary tables might be created on the master, as the temporary tables might be missing from the backup, making it inconsistent. Therefore, only use `mysqlbackup` with an SBR slave if you know no temporary tables are created on the master.

## 6.3 Restoring a Master Database

To fix a corruption problem in a replication master database, you can restore the backup, taking care not to propagate unnecessary SQL operations to the slave servers:

1. Shut down the master database and then use, for example, the `copy-back-and-apply-log` command, to restore a backup of it and prepare the data.
2. Edit the master `my.cnf` file and comment out `log-bin`, so that the slaves do not receive twice the binary log needed to recover the master.
3. Replication in the slaves must be stopped temporarily while you pipe the binary log to the master. In the slaves, do:

```
mysql> STOP SLAVE;
```

4. Start the master `mysqld` on the restored backup:

```
$ mysqld
...
InnoDB: Doing recovery: scanned up to log sequence number 0 64300044
InnoDB: Last MySQL binlog file position 0 5585832, file name
./omnibook-bin.002
...
```

InnoDB prints the binary log file (`./omnibook-bin.002` in this case) and the position (`5585832` in this case) it was able to recover to.

5. Pipe the remaining of the binary log files to the restored backup. For example, if there are two more binary log files, `omnibook-bin.003` and `omnibook-bin.004` that come after `omnibook-bin.002`, pipe them all with a single connection to the server (see [Point-in-Time \(Incremental\) Recovery Using the Binary Log](#) for more instructions on using `mysqlbinlog`):

```
$ mysqlbinlog --start-position=5585833 /mysqldata/omnibook-bin.002 \
/mysqldata/omnibook-bin.003 /mysqldata/omnibook-bin.004 | mysql
```

The number of remaining binary log files varies depending on the length of the timespan between the last backup and the time to which you want to bring the database up to date. The longer the timespan, the more remaining binary log files there may be. All the binary log files, containing all the continuous binary log positions in that timespan, are required for a successful restore.

6. The master database is now recovered. Shut down the master and edit `my.cnf` to uncomment `log-bin`.
7. Start the master again.
8. Start replication in the slaves again:

```
mysql> START SLAVE;
```



---

# Chapter 7 Performance Considerations for MySQL Enterprise Backup

## Table of Contents

|                                          |    |
|------------------------------------------|----|
| 7.1 Optimizing Backup Performance .....  | 67 |
| 7.2 Optimizing Restore Performance ..... | 70 |

This chapter describes the performance considerations for backing up and restoring databases using MySQL Enterprise Backup.

## 7.1 Optimizing Backup Performance

This section describes the performance considerations for backing up a database with MySQL Enterprise Backup. When optimizing and tuning the backup procedure, measure both the raw performance (how long it takes the backup to complete) and the amount of overhead on the database server. When measuring backup performance, consider:

- The limits imposed by your backup procedures. For example, if you take a backup every 8 hours, the backup must take less than 8 hours to finish.
- The limits imposed by your network and storage infrastructure. For example, if you need to fit many backups on a particular storage device, you might use compressed backups, even if that made the backup process slower.
- The tradeoff between backup time and restore time. You might choose a set of options resulting in a slightly slower backup, if those options enable the restore to be much faster. See [Section 7.2, “Optimizing Restore Performance”](#) for performance information for the restore process.

## Full or Incremental Backup

After taking a full backup, subsequent backups can be performed more quickly by doing incremental backups, where only the changed data is backed up. For an incremental backup, specify the `--incremental` or `--incremental-with-redo-log-only` option to `mysqlbackup`. See [Section 13.7, “Incremental Backup Options”](#) for information about these options. For usage instructions for the backup and apply stages of incremental backups, see [Section 4.3.3, “Making a Differential or Incremental Backup”](#).

## Compressed Backup

Compressing the backup data before transmitting it to another server involves additional CPU overhead on the database server where the backup takes place, but less network traffic and less disk I/O on the server that is the final destination for the backup data. Consider the load on your database server, the bandwidth of your network, and the relative capacities of the database and destination servers when deciding whether or not to use compression. See [Section 4.3.4, “Making a Compressed Backup”](#) and [Section 13.6, “Compression Options”](#) for information about creating compressed backups.

Compression involves a tradeoff between backup performance and restore performance. In an emergency, the time needed to uncompress the backup data before restoring it might be unacceptable. There might also be storage issues if there is not enough free space on the database server to hold both the compressed backup and the uncompressed data. Thus, the more critical the data is, the more likely that

you might choose not to use compression: accepting a slower, larger backup to ensure that the restore process is as fast and reliable as possible.

## Single-File Backups

The single-file backup by itself is not necessarily faster than the traditional type of backup that produces a directory tree of output files. Its performance advantage comes from combining different steps that you might otherwise have to perform in sequence, such as combining the backup data into a single output file and transferring it to another server. See [Section 12.5, “Other Single-File Backup Operations”](#) for the options related to single-file backups, and [Section 4.3.1, “Making a Single-File Backup”](#) for usage instructions.

## InnoDB Configuration Options Settings

Prior to MySQL 5.5, it was common practice to keep the redo logs fairly small to avoid long startup times when the MySQL server was killed rather than shut down normally. In MySQL 5.5 and higher, the performance of [crash recovery](#) is significantly improved, as explained in [Optimizing InnoDB Configuration Variables](#). With those releases, you can make your redo log files bigger if that helps your backup strategy and your database workload.

As discussed later, there are a number of reasons why you might prefer to run with the setting `innodb_file_per_table=1`.

## Parallel Backup

`mysqlbackup` can take advantage of modern multicore CPUs and operating system threads to perform backup operations in parallel. See [Section 13.10, “Performance / Scalability / Capacity Options”](#) for the options to control how many threads are used for different aspects of the backup process. If you see that there is unused system capacity during backups, consider increasing the values for these options and testing whether doing so increases backup performance:

- When tuning and testing backup performance using a RAID storage configuration, consider the combination of option settings `--read-threads=3 --process-threads=6 --write-threads=3`. Compare against the combination `--read-threads=1 --process-threads=6 --write-threads=1`.
- When tuning and testing backup performance using a non-RAID storage configuration, consider the combination of option settings `--read-threads=1 --process-threads=6 --write-threads=1`.
- When you increase the values for any of the 3 “threads” options, also increase the value of the `--limit-memory` option, to give the extra threads enough memory to do their work.
- If the CPU is not too busy (less than 80% CPU utilization), increase the value of the `--process-threads` option.
- If the storage device that you are backing up from (the source drive) can handle more I/O requests, increase the value of the `--read-threads` option.
- If the storage device that you are backing up to (the destination drive) can handle more I/O requests, increase the value of the `--write-threads` option.

Depending on your operating system, you can measure resource utilization using commands such as `top`, `iostat`, `sar`, `dtrace`, or a graphical performance monitor. Do not increase the number of read or write threads `iowait` once the system `iowait` value reaches approximately 20%.



## MyISAM Considerations



### Important

- Although `mysqlbackup` backs up InnoDB tables without interrupting database use, the final stage that copies non-InnoDB files (such as MyISAM tables and `.frm` files) temporarily puts the database into a read-only state, using the statement `FLUSH TABLES WITH READ LOCK`. For best backup performance and minimal impact on database processing:
  1. Do not run long `SELECT` queries or other SQL statements at the time of the backup run.
  2. Keep your MyISAM tables relatively small and primarily for read-only or read-mostly work.

Then the locked phase at the end of a `mysqlbackup` run is short (maybe a few seconds), and does not disturb the normal processing of `mysqld` much. If the preceding conditions are not met in your database application, use the `--only-innodb` or `--only-innodb-with-frm` option to back up only InnoDB tables, or use the `--no-locking` option to back up non-InnoDB files. Note that MyISAM, `.frm`, and other files copied under the `--no-locking` setting cannot be guaranteed to be consistent, if they are updated during this final phase of the backup.
- For a large database, a backup run might take a long time. Always check that `mysqlbackup` has completed successfully, either by verifying that `mysqlbackup` returned exit code 0, or by observing that `mysqlbackup` has printed the text “mysqlbackup completed OK!”.
- Schedule backups during periods when no DDL operations involving tables are running. See [Section B.1, “Limitations of MySQL Enterprise Backup”](#) for restrictions on backups at the same time as DDL operations.

## Network Performance

For data processing operations, you might know the conventional advice that Unix sockets are faster than TCP/IP for communicating with the database. Although the `mysqlbackup` command supports the options `--protocol=tcp`, `--protocol=socket`, and `--protocol=pipe`, these options do not have a significant effect on backup or restore performance. These processes involve file-copy operations rather than client/server network traffic. The database communication controlled by the `--protocol` option is low-volume. For example, `mysqlbackup` retrieves information about database parameters through the database connection, but not table or index data.

## Data Size

If certain tables or databases contain non-critical information, or are rarely updated, you can leave them out of your most frequent backups and back them up on a less frequent schedule. See [Section 13.8, “Partial Backup and Restore Options”](#) for information about the relevant options, and [Section 4.3.5, “Making a Partial Backup”](#) for instructions about leaving out data from specific tables, databases, or storage engines. Partial backups are faster because they copy, compress, and transmit a smaller volume of data.

To minimize the overall size of InnoDB data files, consider enabling the MySQL configuration option `innodb_file_per_table`. This option can minimize data size for InnoDB tables in several ways:

- It prevents the [InnoDB](#) system tablespace from ballooning in size, allocating disk space that can afterwards only be used by MySQL. For example, sometimes huge amounts of data are only needed temporarily, or are loaded by mistake or during experimentation. Without the [innodb\\_file\\_per\\_table](#) option, the system tablespace expands to hold all this data, and never shrinks afterward.
- It immediately frees the disk space taken up by an [InnoDB](#) table and its indexes when the table is dropped or truncated. Each table and its associated indexes are represented by a [.ibd file](#) that is deleted or emptied by these DDL operations.
- It allows unused space within a [.ibd file](#) to be reclaimed by the [OPTIMIZE TABLE](#) statement, when substantial amounts of data are removed or indexes are dropped.
- It enables partial backups where you back up some [InnoDB](#) tables and not others, as discussed in [Section 4.3.5, “Making a Partial Backup”](#).

Avoid creating indexes that are not used by queries. Because indexes take up space in the backup data, unnecessary indexes slow down the backup process. (The copying and scanning mechanisms used by [mysqlbackup](#) do not rely on indexes to do their work.) For example, it is typically not helpful to create an index on each column of a table, because only one index is used by any query. Because the primary key columns are included in each [InnoDB](#) secondary index, it wastes space to define primary keys composed of numerous or lengthy columns, or multiple secondary indexes with different permutations of the same columns.

## Advanced: Apply-Log Phase (for Directory Backups only)

If you store the backup data on a separate machine, and that machine is not as busy the machine hosting the database server, you can offload some postprocessing work (the [apply-log](#) phase) to that separate machine. [Apply-log Operation](#)

There is always a performance tradeoff between doing the apply-log phase immediately after the initial backup (makes restore faster), or postponing it until right before the restore (makes backup faster). In an emergency, restore performance is the most important consideration. Thus, the more crucial the data is, the more important it is to run the apply-log phase immediately after the backup. Either combine the backup and apply-log phases on the same server by specifying the [backup-and-apply-log](#) option, or perform the fast initial backup, transfer the backup data to another server, and then perform the apply-log phase using one of the options from [Apply-log Operation](#).

## 7.2 Optimizing Restore Performance

This section describes the performance considerations for restoring a database with MySQL Enterprise Backup. This subject is important because:

- The restore operation is the phase of the backup-restore cycle that tends to vary substantially between different backup methods. For example, backup performance might be acceptable using [mysqldump](#), but [mysqldump](#) typically takes much longer than MySQL Enterprise Backup for a restore operation.
- The restore operation is often performed during an emergency, where it is critical to minimize the downtime of the application or web site.
- The restore operation is always performed with the database server shut down.
- The restore operation is mainly dependent on low-level considerations, such as I/O and network speed for transferring files, and CPU speed, processor cores, and so on for uncompressing data.

For the combination of options you can specify for a restore job, see [Section 12.3, “Restore Operations”](#).

## Restoring Different Classes of Backup Data

Restoring a partial backup takes less time than restoring a full backup, because there is less data to physically copy. See [Section 13.8, “Partial Backup and Restore Options”](#) for information about making partial backups.

Restoring a compressed backup takes more time than restoring an uncompressed backup, because the time needed to uncompress the data is typically greater than any time saved by transferring less data across the network. If you need to rearrange your storage to free up enough space to uncompress the backup before restoring it, include that administration work in your estimate of the total time required. In an emergency, the time needed to uncompress the backup data before restoring it might be unacceptable. on the database server to hold both the compressed backup and the uncompressed data. Thus, the more critical the data is, the more likely that you might choose not to use compression: accepting a slower, larger backup to ensure that the restore process is as fast and reliable as possible. See [Section 13.6, “Compression Options”](#) for information about making compressed backups.

The unpacking process to restore a single-file backup is typically not expensive either in terms of raw speed or extra storage. Each file is unpacked directly to its final destination, the same as if it was copied individually. Thus, if you can speed up the backup substantially or decrease its storage requirements by using single-file backups, that typically does not involve a tradeoff with restore time. See [Section 12.5, “Other Single-File Backup Operations”](#) for information about making single-file backups.

## The Apply-Log Phase (for Directory Backups only)

See [Advanced: Apply-Log Phase \(for Directory Backups only\)](#) for performance considerations regarding the apply-log phase.

## Network Performance

For data processing operations, you might know the conventional advice that Unix sockets are faster than TCP/IP for communicating with the database. Although the `mysqlbackup` command supports the options `--protocol=tcp`, `--protocol=socket`, and `--protocol=pipe`, these options do not have a significant effect on backup or restore performance. These processes involve file-copy operations rather than client/server network traffic. The database communication controlled by the `--protocol` option is low-volume. For example, `mysqlbackup` retrieves information about database parameters through the database connection, but not table or index data.

## Parallel Restore

`mysqlbackup` can take advantage of modern multicore CPUs and operating system threads to perform backup operations in parallel. See [Section 13.10, “Performance / Scalability / Capacity Options”](#) for the options to control how many threads are used for different aspects of the restore process. If you see that there is unused system capacity during a restore, consider increasing the values for these options and testing whether doing so increases restore performance:

- When tuning and testing backup performance using a RAID storage configuration, consider the combination of option settings `--read-threads=3` `--process-threads=6` `--write-threads=3`. Compare against the combination `--read-threads=1` `--process-threads=6` `--write-threads=1`.
- When tuning and testing backup performance using a non-RAID storage configuration, consider the combination of option settings `--read-threads=1` `--process-threads=6` `--write-threads=1`.
- When you increase the values for any of the 3 “threads” options, also increase the value of the `--limit-memory` option, to give the extra threads enough memory to do their work.

- If the CPU is not too busy (less than 80% CPU utilization), increase the value of the `--process-threads` option.
- If the storage device that you are restoring from (the source drive) can handle more I/O requests, increase the value of the `--read-threads` option.
- If the storage device that you are restoring to (the destination drive) can handle more I/O requests, increase the value of the `--write-threads` option.

Depending on your operating system, you can measure resource utilization using commands such as `top`, `iostat`, `sar`, `dtrace`, or a graphical performance monitor. Do not increase the number of read or write threads `iowait` once the system `iowait` value reaches approximately 20%.

---

## Chapter 8 Encryption for Backups

In order to enhance security for backed up data, MySQL Enterprise Backup provides encryption for single-file backups. The encryption can also be applied when creating a partial, compressed, or incremental single-file backups, and for [streaming backup data to another device or server](#).

The encryption is performed with Advanced Encryption Standard (AES) block cipher in CBC mode, with a key string of 64 hexadecimal digits supplied by the user. Decryption is performed using the same key. The key can be created manually just by putting together 64 random hexadecimal bytes, or it can be generated by [shasum](#) (or similar programs for hash calculations that work on your platform) by supplying it with a keyphrase:

```
$ echo -n "my secret passphrase" | shasum -a 256
a7e845b0854294da9aa743b807cb67b19647c1195ea8120369f3d12c70468f29 -
```

Note that the “-” at the end is not part of the key and should be ignored. Supply the key to [mysqlbackup](#) with the `--key` option, or paste the key into a key file and supply the file's pathname to [mysqlbackup](#) with the `--key-file` option.

To generate a key randomly, you can use tools like OpenSSL:

```
$ openssl rand 32 -hex
8f3ca9b850ec6366f4a54feba99f2dc42fa79577158911fe8cd641ffff1e63d6
```

To put an OpenSSL-generated key into a key file, you can do the following:

```
$ openssl rand 32 -hex >keyfile
$ cat keyfile
6a1d325e6ef0577f3400b7cd624ae574f5186d0da2eeb946895de418297ed75b
```

The encryption function uses MySQL Enterprise Backup's own encryption format, which means decryption is possible only by using MySQL Enterprise Backup. For Unix-like operating systems, different magic numbers are used to identify encrypted and unencrypted backup files. For examples, you can add these lines to the `/etc/magic` file of your operating system:

```
0 string MBackupP\n MySQL Enterprise Backup backup image
0 string MebEncR\n MySQL Enterprise Backup encrypted backup
```

The `file` command can then be used to identify the file types:

```
$ file /backups/image1 /backups/image2
/backups/image1: MySQL Enterprise Backup backup image
/backups/image2: MySQL Enterprise Backup encrypted backup
```

The command options used for encryption and decryption are `--encrypt`, `--decrypt`, `--key`, and `--key-file`. These options can be used with various operations on backup images. See [Section 13.13, “Encryption Options”](#) for details.

The following is a sample command for creating an encrypted backup:

```
mysqlbackup --backup-image=/backups/image.enc --encrypt
--key=23D987F3A047B475C900127148F9E0394857983645192874A2B3049570C12A34
--backup-dir=/var/tmp/backup backup-to-image
```

To use a key file for the same task:

```
mysqlbackup --backup-image=/backups/image.enc --encrypt
--key-file=/meb/key --backup-dir=/var/tmp/backup backup-to-image
```

---

To decrypt a backup when extracting it:

```
mysqlbackup --backup-image=/backups/image.enc --decrypt  
--key-file=/meb/key --backup-dir=/backups/extract-dir extract
```

To validate an encrypted backup image:

```
mysqlbackup --backup-image=/logs/encimage.bi --decrypt --key-file=/meb/enckey validate
```

---

# Chapter 9 Using MySQL Enterprise Backup with Media Management Software (MMS) Products

## Table of Contents

|                                                        |    |
|--------------------------------------------------------|----|
| 9.1 Backing Up to Tape with Oracle Secure Backup ..... | 75 |
|--------------------------------------------------------|----|

This section describes how you can use MySQL Enterprise Backup in combination with media management software (MMS) products for creating backups for your database. Such products are typically used for managing large volumes of backup data, often with high-capacity backup devices such as tape drives. Also see [this whitepaper](#) for more discussions on using an MMS product with `mysqlbackup` through the System Backup to Tape (SBT) interface.

## 9.1 Backing Up to Tape with Oracle Secure Backup

Tape drives are affordable, high-capacity storage devices for backup data. MySQL Enterprise Backup can interface with media management software (MMS) such as Oracle Secure Backup (OSB) to drive MySQL backup and restore jobs. The media management software must support Version 2 or higher of the System Backup to Tape (SBT) API.

On the MySQL Enterprise Backup side, you run the backup job as a single-file backup using the `--backup-image` parameter, with the prefix `sbt:` in front of the filename, and optionally pass other `--sbt-*` parameters to `mysqlbackup` to control various aspects of the SBT processing. The `--sbt-*` options are listed in [Section 13.9, “Single-File Backup Options”](#).

On the OSB side, you can schedule MySQL Enterprise Backup jobs by specifying a configurable command that calls `mysqlbackup`. You control OSB features such as encryption by defining a “storage selector” that applies those features to a particular backup, and passing the name of the storage selector to OSB using the MySQL Enterprise Backup parameter `--sbt-database-name=storage_selector`.

### To back up MySQL data to tape:

- Specify the `--backup-image=sbt:name` parameter of `mysqlbackup` to uniquely identify the backup data. The `sbt:` prefix sends the backup data to the MMS rather than a local file, and the remainder of the argument value is used as the unique backup name within the MMS.
- Specify the `--sbt-database-name` parameter of `mysqlbackup` to enable the OSB operator to configure a storage selector for backups from this MySQL source. (This parameter refers to a “storage selector” defined by the OSB operator, not to any MySQL database name.) By default, `mysqlbackup` supplies a value of `MySQL` for this MMS parameter. The argument to this option is limited to 8 bytes.
- If you have multiple media management programs installed, to select the specific SBT library to use, specify the `--sbt-lib-path` parameter of the `mysqlbackup` command. If you do not specify the `--sbt-lib-path` parameter, `mysqlbackup` uses the normal operating system paths and environment variables to locate the SBT library, which is named `libobk.so` on Linux and Unix systems and `ORASBT.DLL` on Windows systems. When you specify `--sbt-lib-path`, you can use a different filename for the library in addition to specifying the path.
- Specify any other product-specific settings that are normally controlled by environment variables using the `--sbt-environment` option.

Each time an online backup is made to a tape using the SBT API, besides recording the backup in the `mysql.backup_history` and the `mysql.backup_progress` tables, an entry is also made

to the `mysql.backup_sbt_history` table on the backed up MySQL instance. That facilitates the management of tape backups by allowing easy look-ups for information on them. The definition of the `backup_sbt_history` table is shown below:

```
mysql> DESCRIBE TABLE `backup_sbt_history`;
```

| Field              | Type          | Null | Key | Default             | Extra |
|--------------------|---------------|------|-----|---------------------|-------|
| backup_id          | bigint(20)    | NO   |     | NULL                |       |
| backup_file_name   | varchar(4096) | NO   |     | NULL                |       |
| file_creation_time | timestamp     | NO   |     | 0000-00-00 00:00:00 |       |
| file_expiry_time   | timestamp     | NO   |     | 0000-00-00 00:00:00 |       |
| volume_label       | varchar(64)   | NO   |     | NULL                |       |
| sbt_error_msg      | varchar(4096) | NO   |     | NULL                |       |
| sbt_error_code     | int(11)       | NO   |     | NULL                |       |

Here are the descriptions for the fields in the table:

- `backup_id`: The backup's ID, which is also recorded in the backup's entries in the `mysql.backup_history` and the `mysql.backup_progress` tables.
- `backup_file_name`: The file name provided by the user through the `--backup-image=sbt:name` option.
- `file_creation_time`: Creation date and time for the tape backup.
- `file_expiry_time`: Expiration date and time for the tape backup.
- `volume_label`: Volume label for the physical medium which contains the tape backup.
- `sbt_error_msg`: Error message, when an error occurs while retrieving information for the tape backup.
- `sbt_error_code`: Error code, when an error occurs while retrieving information for the tape backup.

Multiple entries, one for each volume label, are created in the `mysql.backup_sbt_history` table, if the backup is split across multiple volumes.

Here are some sample entries in the `mysql.backup_sbt_history` table:

```
mysql> select * from mysql.backup_sbt_history;
```

| backup_id         | backup_file_name | file_creation_time  | file_expiry_time    | volume_label | sbt_error_ |
|-------------------|------------------|---------------------|---------------------|--------------|------------|
| 14377247250385175 | test1.mbi        | 2015-07-24 09:59:00 | 0000-00-00 00:00:00 | MySQL-000018 |            |
| 14377259679196369 | test2.mbi        | 2015-07-24 10:19:43 | 0000-00-00 00:00:00 | MySQL-000018 |            |
| 14377265350728997 | test3.mbi        | 2015-07-24 10:29:10 | 0000-00-00 00:00:00 | MySQL-000018 |            |
| 14377265350728997 | test3.mbi        | 2015-07-24 10:29:10 | 0000-00-00 00:00:00 | MySQL-000019 |            |

4 rows in set (0.00 sec)

### To restore MySQL data from tape:

- Specify the `--backup-image=sbt:name` parameter of `mysqlbackup` as part of the restore operation. Use the same `name` value which was used during the backup. This single parameter retrieves the appropriate data from the appropriate tape device.
- Optionally use the `--sbt-lib-path` option, using the same value as for the backup operation.
- Specify any other product-specific settings that are normally controlled by environment variables using the `--sbt-environment` option.



For product-specific information about Oracle Secure Backup, see [the Oracle Secure Backup documentation](#).

**Example 9.1 Sample `mysqlbackup` Commands Using MySQL Enterprise Backup with Oracle Secure Backup**

```
# Uses libobk.so or ORASBT.DLL, at standard locations:
mysqlbackup --port=3306 --protocol=tcp --user=root --password \
  --backup-image=sbt:backup-shoeprod-2011-05-30 \
  --backup-dir=/backup backup-to-image

# Associates this backup with storage selector 'shoeprod':
mysqlbackup --port=3306 --protocol=tcp --user=root --password \
  --backup-image=sbt:backup-shoeprod-2011-05-30 \
  --sbt-database-name=shoeprod \
  --backup-dir=/backup backup-to-image

# Uses an alternative SBT library, /opt/Other-MMS.so:
mysqlbackup --port=3306 --protocol=tcp --user=root --password \
  --backup-image=sbt:backup-shoeprod-2011-05-30 \
  --sbt-lib-path=/opt/Other-MMS.so \
  --backup-dir=/backup backup-to-image
```



---

# Chapter 10 Troubleshooting for MySQL Enterprise Backup

## Table of Contents

|                                                             |    |
|-------------------------------------------------------------|----|
| 10.1 Monitoring Backups with MySQL Enterprise Monitor ..... | 79 |
| 10.2 Error codes of MySQL Enterprise Backup .....           | 79 |
| 10.3 Working Around Corruption Problems .....               | 80 |
| 10.4 Using the MySQL Enterprise Backup Logs .....           | 80 |
| 10.5 Using the MySQL Enterprise Backup Manifest .....       | 82 |

To troubleshoot issues regarding backup and restore with the MySQL Enterprise Backup product, consider the following aspects:

- Before troubleshooting any problem, familiarize yourself with the known limits and restrictions on the product, in [Appendix B, MySQL Enterprise Backup Limitations](#).
- If `mysqlbackup` encounters problems during operating system calls, it returns the corresponding OS error codes. You might need to consult your operating system's documentation for the meaning of those error codes and how to handle them.
- The output from `mysqlbackup` is sent to `stderr` rather than `stdout`. By default, the same output is also saved to a log file in the `backup_dir` for use in error diagnosis. See [Section 13.11, "Message Logging Options"](#) for details on how to configure this logging feature.
- Incremental backups require care to specify a sequence of time periods. You must record the final LSN value at the end of each backup, and specify that value in the next incremental backup. You must also make sure that the full backup you restore is prepared correctly first, so that it contains all the changes from the sequence of incremental backups.
- As `mysqlbackup` proceeds, it writes progress information into the `mysql.backup_progress` table. When the command finishes the backup operation, it records status information in the `mysql.backup_history` table. You can query those tables to monitor ongoing backup jobs, see how much time has been used for various stages, and check if any errors have occurred.

## 10.1 Monitoring Backups with MySQL Enterprise Monitor

With the combination of the MySQL Enterprise Backup and MySQL Enterprise Monitor products, you can monitor the progress and history of backup jobs without writing your own queries or scripts:

- The MySQL Enterprise Monitor graphs `Backup Run Time` and `Backup Locked Time` chart how long the phases of backup jobs take.
- The MySQL Enterprise Monitor rules `MySQL Enterprise Backup Failed`, `MySQL Enterprise Backup Succeeded`, `MySQL Enterprise Backup Lock Time Excessive`, `Incremental MySQL Enterprise Backups Not Enabled`, and `Last Full MySQL Enterprise Backup Too Old` alert you to issues related to backup jobs.

The monitoring capability requires MySQL Enterprise Monitor 2.3.4 and higher. For information about these MySQL Enterprise Monitor features, see the [MySQL Enterprise Monitor User's Guide](#).

## 10.2 Error codes of MySQL Enterprise Backup

The return code of the MySQL Enterprise Backup (`mysqlbackup`) process is 0 if the backup or restore run succeeds. If the run fails for any reason, the return code is 1.

## 10.3 Working Around Corruption Problems

Sometimes the operating system or the hardware can corrupt a data file page at a location that does not cause a database error, but prevents `mysqlbackup` from completing:

```
mysqlbackup: Re-reading page at offset 0 3185082368 in /sqldata/mts/ibdata15
mysqlbackup: Re-reading page at offset 0 3185082368 in /sqldata/mts/ibdata15
mysqlbackup: Error: page at offset 0 3185082368 in /sqldata/mts/ibdata15 seems corrupt!
```

A corruption problem can have different causes. Here are some suggestions for dealing with it:

- The problem can occur if the MySQL server is too busy. Before trying other solutions, you might want to perform the backup again using some non-default settings for the following `mysqlbackup` options:
  - `--page-reread-time=MS`. Try set the value to, for example, "0.05", for faster rereads during checksum failures.
  - `--page-reread-count=retry_limit`. Try set the value to, for example, "1000", to allow more rereads during checksum failures before MySQL Enterprise Backup gives up and throws an error.
- Scrambled data in memory can cause the problem even though the data on disk is actually uncorrupted. Reboot the database server and the storage device to see if the problem persists.
- If the problem persists after the database server and the storage device have been restarted, you might really have a corruption on your disk. You might consider restoring data from an earlier backup and "roll forward" the recent changes to bring the database back to its current state.
- If you want to make MySQL Enterprise Backup finish a backup anyway before you go and investigate the root cause of the issue, you can rewrite the checksum values on the disk by running the `innochecksum` utility on the server:

```
innochecksum --no-checksum --write=crc32
```

The option `--no-checksum` disable the verification function of the tool, and the option `--write=crc32` makes `innochecksum` rewrite the checksum values on the disk.

**IMPORTANT:** Do not treat corruption problems as a minor annoyance. Find out what is wrong with the system that causes the corruption—however, such troubleshooting is beyond the scope of this manual.

## 10.4 Using the MySQL Enterprise Backup Logs

The `mysql.backup_progress` table lets you monitor backup jobs as they run, and the `mysql.backup_history` table lets you see the results of completed jobs. Because these tables are created with the CSV storage engine, you can query them from SQL, or parse the text files from an application or script.

To skip updating these tables for a backup operation, use the `--no-history-logging` option.

### backup\_progress Table

Each row in the `backup_progress` table records a state change or message from a running backup job. The `backup_progress` table has the following columns:

- backup\_id
- tool\_name
- error\_code
- error\_message
- current\_time
- current\_state

Because the CSV storage engine cannot represent `NULL` values directly, the logs use a -1 value instead, for example in the `binlog_pos` column if binary logging is not enabled.

Use the `backup_id` value to group together the information for a single backup operation, and to correlate with the corresponding row in the `backup_history` table after the job is finished.

Use the `error_code` and `error_message` values to track the progress of the job, and detect if a serious error occurs that requires stopping the backup operation.

Use the `current_time` and `current_state` values to measure how long each part of the backup operation takes, to help with planning the time intervals for future backups.

## backup\_history Table

Each row in the `backup_history` table records the details of one completed backup job, produced by `mysqlbackup` command. The `backup_history` table has the following columns:

- backup\_id
- tool\_name
- start\_time
- end\_time
- binlog\_pos
- binlog\_file
- compression\_level
- engines
- innodb\_data\_file\_path
- innodb\_file\_format
- start\_lsn
- end\_lsn
- backup\_type
- backup\_format
- mysql\_data\_dir
- innodb\_data\_home\_dir

- `innodb_log_group_home_dir`
- `innodb_log_files_in_group`
- `innodb_log_file_size`
- `backup_destination`
- `lock_time`
- `exit_state`
- `last_error`
- `last_error_code`

Use the `end_lsn` value to automate operations related to incremental backups. When you take a full or incremental backup, you can specify the end LSN from that backup as the starting LSN for the next incremental backup with the `start-lsn` option (an alternative to specifying the start LSN value is to use the `incremental-base` option).

Use the values that correspond to backup-related configuration settings, such as `mysql_data_dir`, `innodb_data_home_dir`, and `backup_destination`, to confirm that the backups are using the right source and destination directories.

Use the values `exit_state`, `last_error`, and `last_error_code` to evaluate the success or failure of each backup.

If `last_error` is 'NO\_ERROR', the backup operation was successful. In case of any errors, you can retrieve the full list of errors for that backup operation from the `backup_progress` table.

## 10.5 Using the MySQL Enterprise Backup Manifest

Each backup directory includes some files in the `meta` subdirectory that detail how the backup was produced, and what files it contains. The files containing this information are known collectively as the `manifest`.

`mysqlbackup` produces these files for use by database management tools; it does not consult or modify the manifest files after creating them. Management tools can use the manifest during diagnosis and troubleshooting procedures, for example where the original MySQL instance has been lost entirely and the recovery process is more involved than copying files back to a working MySQL server.

The files in the manifest include:

- `backup_create.xml`: information about the backup operation.
- `backup_content.xml`: information about the files in the backup. This information is only complete and consistent when the backup operation succeeds. A management tool might use this information to confirm which tables are part of a backup. A management tool might compare the checksum recorded in the manifest for a single-file backup against the checksum for the file after the single-file backup is unpacked. The file also contains details of all the plugins defined on the backed-up server, by which users should make sure the same plugins are defined in the same manner on the target server for restoration.
- `image_files.xml`: information about the files in a single-file backup. (Only produced for backups taken with the `backup-to-image` and `backup-dir-to-image` commands.) A management tool might use the paths recorded in this file to plan or automate the unpacking of a single-file backup using

the `image-to-backup-dir` or `extract` commands, or to remap the paths of extracted files with the `--src-entry` and `--dst-entry` options.





---

## Part III `mysqlbackup` Command Reference

---

---

---

---

## Table of Contents

|    |                                                          |     |
|----|----------------------------------------------------------|-----|
| 11 | <code>mysqlbackup</code> .....                           | 89  |
| 12 | <code>mysqlbackup</code> commands .....                  | 91  |
|    | 12.1 Backup Operations .....                             | 91  |
|    | 12.2 Update Operations .....                             | 92  |
|    | 12.3 Restore Operations .....                            | 93  |
|    | 12.4 Validation Operations .....                         | 95  |
|    | 12.5 Other Single-File Backup Operations .....           | 96  |
| 13 | <code>mysqlbackup</code> Command-Line Options .....      | 99  |
|    | 13.1 Standard Options .....                              | 104 |
|    | 13.2 Connection Options .....                            | 106 |
|    | 13.3 Server Repository Options .....                     | 107 |
|    | 13.4 Backup Repository Options .....                     | 109 |
|    | 13.5 Metadata Options .....                              | 113 |
|    | 13.6 Compression Options .....                           | 114 |
|    | 13.7 Incremental Backup Options .....                    | 115 |
|    | 13.8 Partial Backup and Restore Options .....            | 117 |
|    | 13.9 Single-File Backup Options .....                    | 123 |
|    | 13.10 Performance / Scalability / Capacity Options ..... | 125 |
|    | 13.11 Message Logging Options .....                      | 132 |
|    | 13.12 Progress Report Options .....                      | 133 |
|    | 13.13 Encryption Options .....                           | 136 |
|    | 13.14 Cloud Storage Options .....                        | 137 |
|    | 13.15 Options for Special Backup Types .....             | 139 |
| 14 | Configuration Files and Parameters .....                 | 141 |



---

## Chapter 11 `mysqlbackup`

The `mysqlbackup` client is an easy-to-use tool for all backup and restore operations. During backup operations, `mysqlbackup` backs up:

- All InnoDB tables and indexes, including:
  - The InnoDB `system tablespace`, which, by default contains all the InnoDB tables.
  - Any separate data files produced with the InnoDB `file-per-table` setting. Each one contains one table and its associated indexes. Each data file can use either the original `Antelope` or the new `Barracuda` file format.
- All MyISAM tables and indexes.
- Tables managed by other storage engines.
- Other files underneath the MySQL data directory, such as the `.frm` files that record the structure of each table.
- Any other files in the database subdirectories under the server's data directory.

In addition to creating backups, `mysqlbackup` can pack and unpack backup data, apply to the backup data any changes to InnoDB tables that occurred during the backup operation, and restore data, index, and log files back to their original locations, or to other places.

Here are some sample commands to start a backup operation with `mysqlbackup` are:

```
# Information about data files can be retrieved through the database connection.
# Specify connection options on the command line.
mysqlbackup --user=dba --password --port=3306 \
  --with-timestamp --backup-dir=/export/backups \
  backup

# Or we can include the above options in the configuration file
# under the [mysqlbackup] section, and just specify the configuration file
# and the 'backup' operation.
mysqlbackup --defaults-file=/usr/local/mysql/my.cnf backup

# Or we can specify the configuration file as above, but
# override some of those options on the command line.
mysqlbackup --defaults-file=/usr/local/mysql/my.cnf \
  --compress --user=backupadmin --password --port=18080 \
  backup
```

The `--user` and the `--password` you specify are used to connect to the MySQL server. This MySQL user must have certain privileges in the MySQL server, as described in [Section 4.1.2, “Grant MySQL Privileges to Backup Administrator”](#).

The `--with-timestamp` option places the backup in a subdirectory created under the directory you specified above. The name of the backup subdirectory is formed from the date and the clock time of the backup run.

For the meanings of other command-line options, see [Chapter 13, `mysqlbackup` Command-Line Options](#). For information about configuration files, see [Chapter 14, `Configuration Files and Parameters`](#).

Make sure that the user or the cron job running `mysqlbackup` has the rights to copy files from the MySQL database directories to the backup directory.

---

Make sure that your connection timeouts are long enough so that the `mysqlbackup` command can keep the connection to the server open for the duration of the backup run. `mysqlbackup` pings the server after copying each database to keep the connection alive.



### Important

- Although `mysqlbackup` backs up InnoDB tables without interrupting database use, the final stage that copies non-InnoDB files (such as MyISAM tables and `.frm` files) temporarily puts the database into a read-only state, using the statement `FLUSH TABLES WITH READ LOCK`. For best backup performance and minimal impact on database processing:

1. Do not run long `SELECT` queries or other SQL statements at the time of the backup run.
2. Keep your MyISAM tables relatively small and primarily for read-only or read-mostly work.

Then the locked phase at the end of a `mysqlbackup` run is short (maybe a few seconds), and does not disturb the normal processing of `mysqld` much. If the preceding conditions are not met in your database application, use the `--only-innodb` option to back up only InnoDB tables, or use the `--no-locking` option to back up non-InnoDB files. Note that MyISAM, `.frm`, and other files copied under the `--no-locking` setting cannot be guaranteed to be consistent, if they are updated during this final phase of the backup.

- For a large database, a backup run might take a long time. Always check that the `mysqlbackup` command has been completed successfully by verifying that `mysqlbackup` has returned the exit code 0, or by observing that `mysqlbackup` has printed the text “mysqlbackup completed OK!”.
- `mysqlbackup` is not the same as the former “MySQL Backup” open source project from the MySQL 6.0 source tree. The MySQL Enterprise Backup product supersedes the MySQL Backup initiative.
- Schedule backups during periods when no DDL operations involving tables are running. See [Section B.1, “Limitations of MySQL Enterprise Backup”](#) for restrictions on creating backups in parallel with the DDL operations.

---

## Chapter 12 `mysqlbackup` commands

### Table of Contents

|                                                |    |
|------------------------------------------------|----|
| 12.1 Backup Operations .....                   | 91 |
| 12.2 Update Operations .....                   | 92 |
| 12.3 Restore Operations .....                  | 93 |
| 12.4 Validation Operations .....               | 95 |
| 12.5 Other Single-File Backup Operations ..... | 96 |

These are commands for the major operations for `mysqlbackup`. Only one of them can be specified for each `mysqlbackup` invocation, and, unlike the command options, the name of a command is not preceded by any dashes.

Each of these commands has its own set of required or allowed command options. For example, the `backup` command typically requires connection information to the database server. The `apply-log` and other commands that operate on the backup data after it is produced require the options that specify where the backup data is located.

The major groups of commands are:

- Backup operations: `backup`, `backup-and-apply-log`, `backup-to-image`
- Update operations: `apply-log`, `apply-incremental-backup`
- Restore operations: `copy-back`, `copy-back-and-apply-log`
- Validation operation: `validate`
- Single-file backup operations: `image-to-backup-dir`, `backup-dir-to-image`, `list-image`, `extract`

### 12.1 Backup Operations

The backup operations are the most frequently performed tasks by MySQL Enterprise Backup. Various kinds of backups can be performed by adding different options, like using `--compress` or `--incremental` for compressed or incremental backups. Here is the syntax for the `mysqlbackup` commands for performing a backup operation:

```
mysqlbackup [STD-OPTIONS]
            [CONNECTION-OPTIONS]
            [SERVER-REPOSITORY-OPTIONS]
            [BACKUP-REPOSITORY-OPTIONS]
            [METADATA-OPTIONS]
            [COMPRESSION-OPTIONS]
            [SPECIAL-BACKUP-TYPES-OPTIONS]
            [INCREMENTAL-BACKUP-OPTIONS]
            [PARTIAL-BACKUP-RESTORE-OPTIONS]
            [SINGLE-FILE-BACKUP-OPTIONS]
            [PERFORMANCE-SCALABILITY-CAPACITY-OPTIONS]
            [MESSAGE-LOGGING-OPTIONS]
            [PROGRESS-REPORT-OPTIONS]
            [ENCRYPTION-OPTIONS]
            [CLOUD-STORAGE-OPTIONS]
            backup-to-image

mysqlbackup [STD-OPTIONS]
            [CONNECTION-OPTIONS]
```

```
[SERVER-REPOSITORY-OPTIONS]
[BACKUP-REPOSITORY-OPTIONS]
[METADATA-OPTIONS]
[COMPRESSION-OPTIONS]
[SPECIAL-BACKUP-TYPES-OPTIONS]
[INCREMENTAL-BACKUP-OPTIONS]
[PARTIAL-BACKUP-RESTORE-OPTIONS]
[SINGLE-FILE-BACKUP-OPTIONS]
[PERFORMANCE-SCALABILITY-CAPACITY-OPTIONS]
[MESSAGE-LOGGING-OPTIONS]
[PROGRESS-REPORT-OPTIONS]
backup | backup-and-apply-log
```

- `backup-to-image`

Produces a single-file backup holding the backup data. In most cases, single-file backups are preferred over directory backups, which are created using the `backup` command.

The command requires the `--backup-image` option to specify the destination file. Can be used to stream the backup to a storage device or another system without ever storing the data on the database server. You can specify `--backup-image=-`, representing standard output, allowing the output to be piped to another command. To avoid mixing normal informational messages with backup output, the `--help` message, errors, alerts, and normal informational messages are always printed to standard error stream.

- `backup`

Backs up data to a directory. In most cases, single-file backups, which are created using the `backup-to-image` command, are preferred over directory backups.

The command only performs the initial phase of a complete backup process. The second phase is performed later by running `mysqlbackup` again with the `apply-log` command, which makes the backup consistent.

- `backup-and-apply-log`

A combination of `backup` and `apply-log`. It cannot be used for a compressed or an incremental backup.

## 12.2 Update Operations

There are two types of operations to bring your backup data up-to-date:

### Apply-log Operation

After a backup job was first completed, the backup data might not be in a consistent state, because data could have been inserted, updated, or deleted while the backup was running. This initial backup file is known as the [raw backup](#). During a backup, `mysqlbackup` also copies the accumulated InnoDB log to a file called `ibbackup_logfile`. In an apply-log operation, the `ibbackup_logfile` file is used to “roll forward” the raw data files, so that every page in the data files corresponds to the same log sequence number of the InnoDB log. This is similar to the operation that takes place during a [crash recovery](#).

For single-file backups, the apply-log operation is usually performed as part of the `copy-back-and-apply-log` command. For directory backups, the `copy-back-and-apply-log` command can also be used, but you also have the two alternatives of

- Performing the apply-log operation together with the back up using the `backup-and-apply-log` command (not applicable for incremental or compressed directory backups)



- Performing the `apply-log` operation separately with the `apply-log` command on the raw backup, before running the `copy-back` command.

```
mysqlbackup [STD-OPTIONS]
             [--limit-memory=MB] [--uncompress] [--backup-dir=PATH]
             [MESSAGE-LOGGING-OPTIONS]
             [PROGRESS-REPORT-OPTIONS]
             apply-log
```

- `apply-log`

*Advanced:* Brings the InnoDB tables in the directory backup up-to-date, including any changes made to the data while the backup was running.

### Example 12.1 Apply Log to Full Backup

```
mysqlbackup --backup-dir=/path/to/backup apply-log
```

It reads the `backup-my.cnf` file inside `backup-dir` to understand the backup. The `my.cnf` default files have no effect other than supplying the `limit-memory=MB` value, which limits usage of memory while doing the `apply-log` operation.

## Apply-incremental-backup Operation

*Advanced:* Use the `apply-incremental-backup` to update a backup directory with data in an incremental backup directory:

```
mysqlbackup [STD-OPTIONS]
             [--incremental-backup-dir=PATH] [--backup-dir=PATH]
             [--limit-memory=MB] [--uncompress]
             [MESSAGE-LOGGING-OPTIONS]
             [PROGRESS-REPORT-OPTIONS]
             apply-incremental-backup
```

- `apply-incremental-backup`

*Advanced:* Brings up-to-date a directory backup, specified by the `--backup-dir` option, using the data from an incremental backup directory, specified with the `--incremental-backup-dir` option. See [Section 5.1.3, “Restoring an Incremental Backup”](#) for instructions on restoring incremental backups.

For a single-file incremental backup, you typically use the `copy-back-and-apply-log` command with the `--incremental` option to apply the data in the incremental image backup to the full backup that has already been restored to the data directory of the target server.

## 12.3 Restore Operations

The restore operations restores the data files from a backup to their original locations on the database server, or to other desired locations. Normally, the restoration process requires the database server to be already shut down (or, at least not operating on the directory you are restoring the data to), except for restorations of backups created with the `--use-tts` option. The options `datadir`, `innodb_log_files_in_group`, and `innodb_log_file_size` must be specified either in the target server's configuration file, in the file specified by the `--defaults-file` option, or as command-line options. For usage examples, see [Chapter 5, \*Recovering or Restoring a Database\*](#).

```
mysqlbackup [STD-OPTIONS]
            [SERVER-REPOSITORY-OPTIONS]
            [--backup-image=IMAGE]
            [--backup-dir=PATH]
            [MESSAGE-LOGGING-OPTIONS]
            [PARTIAL-BACKUP-RESTORE-OPTIONS]
            [PROGRESS-REPORT-OPTIONS]
            [ENCRYPTION-OPTIONS]
            [CLOUD-STORAGE-OPTIONS]
            copy-back-and-apply-log
mysqlbackup [STD-OPTIONS]
            [SERVER-REPOSITORY-OPTIONS]
            [--backup-dir=PATH]
            [MESSAGE-LOGGING-OPTIONS]
            [PARTIAL-BACKUP-RESTORE-OPTIONS]
            [PROGRESS-REPORT-OPTIONS]
            [CLOUD-STORAGE-OPTIONS]
            copy-back
```

- `copy-back-and-apply-log`

In a single step, restores a [single-file backup](#) specified by the `--backup-image` option or a backup from the directory specified by the `--backup-dir` option to a server's data directory and performs an `apply-log` operation to the restored data to bring them up-to-date. Comparing with a multi-step approach for restoring a [single-file backup](#) (which typically consists of performing the successive steps of `extract`, `uncompress`, `apply-log`, and `copy-back` for restoring compressed image, or `extract`, `apply-log`, and `copy-back` for uncompressed image), the command makes the restoration process simpler and faster, and also saves the disk space required.

The following are some special requirements for different kinds of backup restoration using `copy-back-and-apply-log`:

- To restore a compressed directory or image, include the `--uncompress` option in the command line.
- To restore a single-file backup, besides specifying the location of the backup image with the `--backup-image` option, also supply with the `--backup-dir` option the location of a folder that will be used for storing temporary files produced during the restoration process.
- To restore a single-file incremental backup, assuming the full backup (on which the incremental backup was based) has already been restored:
  - Include the `--incremental` option in the command line.
  - Specifies the location of the incremental backup image with the `--backup-image` option
  - Supplies with the `--backup-dir` or `--incremental-backup-dir` option the location of a folder that will be used for storing temporary files produced during the restoration process.
- *Advanced:* To restore an incremental backup directory, assuming the full backup (on which the incremental backup was based) has already been restored:
  - Include the `--incremental` option in the command line.
  - Use either the `--backup-dir` or `--incremental-backup-dir` option to specify the incremental backup directory.
- To restore a backup created with the `--use-tts` option:
  - See the general requirements described in [Section 5.1.4, “Restoring Backups Created with the `--use-tts` Option”](#).

- When restoring an image backup created with the option setting `--use-tts=with-minimum-locking`, also supply with the `--backup-dir` option the location of a folder that will be used for extracting temporarily all tables in the backup and for performing an `apply-log` operation to make the data up-to-date before restoring them to the server's data directory.
- When restoring a backup directory created with the option `--use-tts`, an `apply-log` operation will be performed on the backup directory. That means the backup taken will be altered during the process, and users might want to make an extra copy of the backup directory before proceeding with the restoration, in order to prevent the loss of backup data in case something goes wrong.

Also note that:

- Backups created with the `--skip-unused-pages` option cannot be restored using `copy-back-and-apply-log`.
- For image backups taken with MySQL Enterprise Backup 3.8.2 or earlier, per-table `.ibd` files pointed to by `.isl` files in a backup are restored by `copy-back-and-apply-log` to the server's data directory rather than the locations pointed to by the `.isl` files.
- Due to a known issue, when restoring a compressed backup created with MySQL Enterprise Backup 3.9 or earlier and containing any InnoDB tables that were created on the server as compressed tables (by using the `ROW_FORMAT=COMPRESSED` option, the `KEY_BLOCK_SIZE=` option, or both), do not use `copy-back-and-apply-log`; instead, perform an `apply-log` first, and then a `copy-back`. See entry for Bug# 17992297 in the [MySQL Enterprise Backup 3.10.0 changelog](#) for details.

At the end of the `copy-back-and-apply-log` operation, the file `backup_variables.txt` is being created or updated in the data directory. This file contains metadata about the restored contents and is being used by successive single-step restores of incremental backups; it should not be deleted or modified by users.

For some sample commands for restoring different kinds of backups with the `copy-back-and-apply-log` command, see [Section 5.1, “Performing a Restore Operation”](#).

- `copy-back`

Restores files from a directory backup to their original locations within the MySQL server.

Some clean-up efforts on the target directory for restoration might be needed before performing a full restore (for example, when the backup data is used to set up a new MySQL server or to replace all data of an existing MySQL server). See the discussions [here \[56\]](#) for details.

There are some special requirements when restoring backups created with the `--use-tts` option; see [Section 5.1.4, “Restoring Backups Created with the `--use-tts` Option”](#) for details.



### Warning

When restoring a server for [replication](#) purpose, if the backed-up server has used the `innodb_undo_directory` option to put the undo logs outside of the data directory, when using the file `server-my.cnf` or `server-all.cnf` for the `--defaults-file` option with `copy-back` or `copy-back-and-apply-log`, care should be taken to configure correctly the `innodb_undo_directory` option in the file. Otherwise, the data or log files on the original server might be overwritten by accident.

## 12.4 Validation Operations

To ensure the integrity of the backup data, MySQL Enterprise Backup provides a `validate` command for validating a backup by the checksum values of its data pages after the backup is created or transferred to another system.

```
mysqlbackup [STD-OPTIONS]
            [--backup-dir=PATH] [--backup-image=IMAGE]
            [MESSAGE-LOGGING-OPTIONS]
            [PROGRESS-REPORT-OPTIONS]
            [CLOUD-STORAGE-OPTIONS]
            validate
```

- `validate`

Verifies that a backup is not corrupted, truncated, or damaged. This operation validates the checksum value for each data page in a backup.

To avoid spending excessive time and resources on files that are too heavily corrupted, `mysqlbackup` stops validating a `.ibd` file after more than twenty corrupted pages are found in it, and proceeds to the next file instead. In that case, the operation's summary will not give a full count of corrupted pages, but only says "at least 20 pages are corrupted."

The operation also has the following limitations:

- If any `.ibd` or `.frm` files are missing from the data directory during a backup or have been deleted from a backup after the backup was made, the `validate` operation will not be able to detect the problem.
- If a backup has been corrupted by removing or truncating pages from any of the `.ibd` files inside, the `validate` operation will not be able to detect the problem.
- For any backup directory, the operation can only validate the InnoDB data files (`ibdata*` and `*.ibd` files) in it. Problems with other file types within a backup directory (for example, `.frm` file corruptions) are not detected.

Here is a sample command for validating a backup image:

```
mysqlbackup -uroot --backup-image=/logs/fullimage.mi validate
```

*Advanced:* Here is a sample command for validating a backup directory:

```
mysqlbackup -uroot --backup-dir=/logs/backupext validate
```

For more usage examples for the `validate` command, see [Section 4.2.2, "Verifying a Backup"](#)

## 12.5 Other Single-File Backup Operations

Besides the commands for creating and restoring single-file backups (namely, `backup-to-image` and `copy-back-and-apply-log`), `mysqlbackup` provides a number of other commands for you to work with single-file backups. They are explained below.

```
mysqlbackup [STD-OPTIONS]
            [--backup-image=IMAGE] [--backup-dir=PATH]
            [MESSAGE-LOGGING-OPTIONS]
            [PROGRESS-REPORT-OPTIONS]
            [ENCRYPTION-OPTIONS]
            [CLOUD-STORAGE-OPTIONS]
            image-to-backup-dir
```

```

mysqlbackup [STD-OPTIONS]
             [--backup-dir=PATH] [--backup-image=IMAGE]
             [MESSAGE-LOGGING-OPTIONS]
             [PROGRESS-REPORT-OPTIONS]
             [ENCRYPTION-OPTIONS]
             [CLOUD-STORAGE-OPTIONS]
             backup-dir-to-image

mysqlbackup [STD-OPTIONS]
             [--backup-image=IMAGE] [--src-entry=PATH]
             [MESSAGE-LOGGING-OPTIONS]
             [ENCRYPTION-OPTIONS]
             [CLOUD-STORAGE-OPTIONS]
             list-image

mysqlbackup [STD-OPTIONS]
             [--backup-image=IMAGE]
             [--backup-dir=PATH]
             [--src-entry=PATH] [--dst-entry=PATH]
             [MESSAGE-LOGGING-OPTIONS]
             [PROGRESS-REPORT-OPTIONS]
             [ENCRYPTION-OPTIONS]
             [CLOUD-STORAGE-OPTIONS]
             extract

```

- `image-to-backup-dir`

Unpacks a single-file backup to a full backup directory structure. You specify the paths to both the image file and the destination directory for the unpacking. For usage examples, see [Section 4.3.1, “Making a Single-File Backup”](#).

- `backup-dir-to-image`

Packs an existing backup directory into a single file. The value for the `--backup-image` parameter should either be “-”(stands for standard output) or an absolute path outside of the `backup-dir` directory. Specify a `--backup-image` value of - (standard output) to stream an existing backup directory structure to a tape device or a command that transfers the backup to another server. For usage examples, see [Section 4.3.1, “Making a Single-File Backup”](#).

- `list-image`

Display the contents of a single-file backup. Lists all files and directories in the image. The `--src-entry=name` can be used to list a specific file or directory; if the name is a directory, all its files and subdirectories inside the image are recursively listed. For usage examples, see [Section 4.3.1, “Making a Single-File Backup”](#).



#### Note

The `list-image` operation can be performed on a cloud backup only if the cloud proxy supports HTTP range headers.

- `extract`

Unpacks an individual file or directory from a single-file backup. It is useful for troubleshooting, or for restorations that do not require the full set of backup data. The resulting file or directory goes into the current directory, or into `backup-dir` if specified. All files and directory contents in the image with absolute path names are extracted into the same absolute paths on the local system. For usage examples, see [Section 4.3.1, “Making a Single-File Backup”](#).

The `--src-entry=path` option can be used for selective extraction of a single file or single directory in image. Specify the path as it appears in the image.

The `--dst-entry=path` option, along with `--src-entry=path` option can be used to extract a single file or single directory into a user-specified file or directory respectively. If the `--src-entry` option is used, but `--dst-entry` option is omitted, then the selected file or directory is extracted to the same path in the local file system.

The default destination for the extract is the current working directory. It can be overridden by the `--backup-dir` option. All the files with relative pathnames in the image are extracted to pathnames relative to the destination directory.

If the image contains some entries with absolute pathnames, those entries are extracted to the same absolute pathnames even if `--backup-dir` option is specified. The `--dst-entry` option must be used to relocate an absolute pathname.

---

## Chapter 13 `mysqlbackup` Command-Line Options

### Table of Contents

|                                                          |     |
|----------------------------------------------------------|-----|
| 13.1 Standard Options .....                              | 104 |
| 13.2 Connection Options .....                            | 106 |
| 13.3 Server Repository Options .....                     | 107 |
| 13.4 Backup Repository Options .....                     | 109 |
| 13.5 Metadata Options .....                              | 113 |
| 13.6 Compression Options .....                           | 114 |
| 13.7 Incremental Backup Options .....                    | 115 |
| 13.8 Partial Backup and Restore Options .....            | 117 |
| 13.9 Single-File Backup Options .....                    | 123 |
| 13.10 Performance / Scalability / Capacity Options ..... | 125 |
| 13.11 Message Logging Options .....                      | 132 |
| 13.12 Progress Report Options .....                      | 133 |
| 13.13 Encryption Options .....                           | 136 |
| 13.14 Cloud Storage Options .....                        | 137 |
| 13.15 Options for Special Backup Types .....             | 139 |

The following sections describe the command-line options for the different modes of operation of `mysqlbackup`.

The table below list all the command options for `mysqlbackup`. Use the hyperlinks at the option names to jump to the detailed descriptions for the options.



#### Note

The command options can also be specified in configuration files; see explanations in [Chapter 14, Configuration Files and Parameters](#). `mysqlbackup` follows the MySQL standard practice for handling duplicate options, whether specified in a configuration file, on the command line, or both. Options are processed first from configuration files, then from the command line. If an option is specified more than once, the last instance takes precedence.

**Table 13.1 List of All Options**

| Format                                             | Description                                                                | Introduced |
|----------------------------------------------------|----------------------------------------------------------------------------|------------|
| <a href="#">--backup-dir</a>                       | The directory to store the backup data.                                    |            |
| <a href="#">--backup-image</a>                     | Specifies the path name of the backup image.                               |            |
| <a href="#">--backup_innodb_checksum_algorithm</a> | The name of the checksum algorithm used for validating InnoDB tablespaces. |            |
| <a href="#">--backup_innodb_data_file_path</a>     | Specifies InnoDB system tablespace files' path and size in backup.         |            |
| <a href="#">--backup_innodb_data_home_dir</a>      | Backup base directory for all InnoDB data files in the system tablespace.  |            |
| <a href="#">--backup_innodb_log_file_size</a>      | The size in bytes of each InnoDB backup log file.                          |            |
| <a href="#">--backup_innodb_log_files_in_group</a> | Number of InnoDB log files in backup.                                      |            |

| Format                                 | Description                                                                                                             | Introduced |
|----------------------------------------|-------------------------------------------------------------------------------------------------------------------------|------------|
| --<br>backup_innodb_log_group_home_dir | Backup directory for InnoDB log files.                                                                                  |            |
| --backup_innodb_page_size              | The page size for all InnoDB tablespaces in a MySQL instance.                                                           |            |
| --<br>backup_innodb_undo_directory     | The relative or absolute directory path where InnoDB creates separate tablespaces for the undo logs.                    |            |
| --backup_innodb_undo_logs              | Number of rollback segments in the system tablespace that InnoDB uses within a transaction.                             |            |
| --<br>backup_innodb_undo_tablespaces   | The number of tablespace files that the undo logs are divided between when a non-zero innodb_undo_logs setting is used. |            |
| --character-sets-dir                   | Directory for character set files.                                                                                      |            |
| --cloud-access-key-id                  | AWS access key ID for logging onto Amazon S3.                                                                           |            |
| --cloud-aws-region                     | Region for Amazon Web Services that mysqlbackup access for S3.                                                          |            |
| --cloud-bucket                         | The storage bucket on Amazon S3 for the backup image.                                                                   |            |
| --cloud-container [138]                | The Swift container for the backup image.                                                                               | 3.12.0     |
| --cloud-identity-url [139]             | The URL of the Keystone identity service.                                                                               | 3.12.0     |
| --cloud-object [138]                   | The Swift object for the backup image.                                                                                  | 3.12.0     |
| --cloud-object-key                     | The Amazon S3 object key for the backup image.                                                                          |            |
| --cloud-password [138]                 | Password for user specified by --cloud-user-id.                                                                         | 3.12.0     |
| --cloud-proxy                          | Proxy address and port number for overriding the environment's default proxy settings for accessing cloud service.      |            |
| --cloud-region [139]                   | The Keystone region for the user specified by --cloud-user-id.                                                          | 3.12.0     |
| --cloud-secret-access-key              | AWS secret access key.                                                                                                  |            |
| --cloud-service                        | Cloud service for data backup or restoration.                                                                           |            |
| --cloud-tempauth-url [139]             | The URL of the identity service for authenticating user credentials with Swift's TempAuth authentication system.        | 3.12.0     |
| --cloud-tenant [139]                   | The Keystone tenant for the user specified by --cloud-user-id.                                                          | 3.12.0     |
| --cloud-trace                          | Print trace information for cloud operations.                                                                           |            |
| --cloud-user-id [138]                  | User ID for accessing Swift.                                                                                            | 3.12.0     |
| --comments                             | Specifies comments string.                                                                                              |            |
| --comments-file                        | Specifies path to comments file.                                                                                        |            |
| --compress                             | Create backup in compressed format.                                                                                     |            |
| --compress-level                       | Specifies the level of compression.                                                                                     |            |
| --compress-method                      | Specifies the compression algorithm.                                                                                    |            |
| --connect-if-online                    | Use connection only if available.                                                                                       |            |
| --connect_timeout                      | Connection timeout in seconds.                                                                                          |            |



| Format                                        | Description                                                                                                                                            | Introduced |
|-----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|------------|
| <code>--databases</code>                      | [Legacy] Specifies the list of non-InnoDB tables to back up.                                                                                           |            |
| <code>--databases-list-file</code>            | [Legacy] Specifies the pathname of a file that lists the non-InnoDB tables to be backed up.                                                            |            |
| <code>--datadir</code>                        | Path to mysql server data directory.                                                                                                                   |            |
| <code>--debug</code>                          | Print debug information.                                                                                                                               |            |
| <code>--decrypt</code>                        | Decrypt backup image written in an MEB Secure File.                                                                                                    |            |
| <code>--default-character-set</code>          | Set the default character set.                                                                                                                         |            |
| <code>--defaults-extra-file</code>            | Read this file after the global files are read.                                                                                                        |            |
| <code>--defaults-file</code>                  | Only read default options from the given file.                                                                                                         |            |
| <code>--defaults-group-suffix</code>          | Also read option groups with the usual names and a suffix of str.                                                                                      |            |
| <code>--disable-manifest</code>               | Disable generation of manifest files for a backup operation.                                                                                           |            |
| <code>--dst-entry</code>                      | Used with single-file backups to extract a single file or directory to a user-specified path.                                                          |            |
| <code>--encrypt</code>                        | Encrypt backup image and write it in an MEB Secure File.                                                                                               |            |
| <code>--exclude-tables</code>                 | Exclude in backup tables whose names match the regular expression REGEXP.                                                                              |            |
| <code>--exec-when-locked</code>               | Execute the specified utility when all tables are locked near the end of the execution.                                                                |            |
| <code>--force [105]</code>                    | Force overwriting of data, log, or image files, depending on the operation.                                                                            |            |
| <code>--help</code>                           | Display help.                                                                                                                                          |            |
| <code>--host</code>                           | Host name to connect.                                                                                                                                  |            |
| <code>--include</code>                        | [Legacy] Backup only those per-table innodb data files which match the regular expression REGEXP.                                                      |            |
| <code>--include-tables</code>                 | Include in backup tables which match the regular expression REGEXP.                                                                                    |            |
| <code>--incremental</code>                    | Specifies that the associated backup or backup-to-image operation is incremental.                                                                      |            |
| <code>--incremental-backup-dir</code>         | Specifies the location under which to store data from an incremental backup.                                                                           |            |
| <code>--incremental-base</code>               | The specification of base backup for --incremental option.                                                                                             |            |
| <code>--incremental-with-redo-log-only</code> | Specifies the incremental backup of InnoDB tables to be based on copying redo log to the backup, without including any InnoDB datafiles in the backup. |            |
| <code>--innodb_checksum_algorithm</code>      | The name of the checksum algorithm used for validating InnoDB tablespaces.                                                                             |            |
| <code>--innodb_data_file_path</code>          | Specifies InnoDB system tablespace files' path and size.                                                                                               |            |
| <code>--innodb_data_home_dir</code>           | Specifies base directory for all InnoDB data files in the shared system tablespace.                                                                    |            |
| <code>--innodb_log_file_size</code>           | The size in bytes of each InnoDB log file in the log group.                                                                                            |            |
| <code>--innodb_log_files_in_group</code>      | The number of InnoDB log files.                                                                                                                        |            |

| Format                                   | Description                                                                                                                                                                            | Introduced |
|------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------|
| <code>--innodb_log_group_home_dir</code> | The directory path to InnoDB log files.                                                                                                                                                |            |
| <code>--innodb_page_size</code>          | The page size for all InnoDB tablespaces in a MySQL instance.                                                                                                                          |            |
| <code>--key</code>                       | The symmetric key used for encryption and decryption.                                                                                                                                  |            |
| <code>--key-file</code>                  | The pathname of a file that contains the symmetric key used for encryption and decryption.                                                                                             |            |
| <code>--limit-memory</code>              | The memory in MB available for the MEB operation.                                                                                                                                      |            |
| <code>--log-bin-index</code>             | Specifies the absolute path of the index file on the MySQL server that lists all the used binary log files (for MySQL 5.5, and for all offline backups).                               |            |
| <code>--login-path</code>                | Read options from the named login path in the <code>.mylogin.cnf</code> login file.                                                                                                    |            |
| <code>--master-info-file</code>          | Specifies the absolute path of the information file in which a slave records information about its master (for offline backups of slave servers only). information file.               |            |
| <code>--messages-logdir</code>           | Specifies the path name of an existing directory for storing the message log.                                                                                                          |            |
| <code>--no-connection</code>             | Do not connect to server.                                                                                                                                                              |            |
| <code>--no-defaults</code>               | Do not read default options from any given file.                                                                                                                                       |            |
| <code>--no-history-logging</code>        | Disable history logging even if connection is available.                                                                                                                               |            |
| <code>--no-locking</code>                | Disable locking of tables even if connection is available.                                                                                                                             |            |
| <code>--number-of-buffers</code>         | Specifies the exact number of memory buffers to be used for the backup operation.                                                                                                      |            |
| <code>--on-disk-full</code>              | Specifies the behavior when a backup process encounters a disk-full condition.                                                                                                         |            |
| <code>--only-innodb</code>               | Back up only InnoDB data and log files.                                                                                                                                                |            |
| <code>--only-innodb-with-frm</code>      | [Legacy] Back up only InnoDB data, log files, and the <code>.frm</code> files associated with the InnoDB tables.                                                                       |            |
| <code>--only-known-file-types</code>     | Includes only files of a list of known types in the backup.                                                                                                                            |            |
| <code>--optimistic-busy-tables</code>    | Perform an optimistic backup, using the regular expression specified with the option to select tables that will be skipped in the first phase of an optimistic backup.                 |            |
| <code>--optimistic-time</code>           | Perform an optimistic backup with the value specified with the option as the optimistic time—a time after which tables that have not been modified are believed to be inactive tables. |            |
| <code>--page-reread-count</code>         | Maximum number of page re-reads.                                                                                                                                                       |            |
| <code>--page-reread-time</code>          | Wait time before a page re-read.                                                                                                                                                       |            |
| <code>--password</code>                  | Connection password.                                                                                                                                                                   |            |
| <code>--pipe</code>                      | alias for <code>--protocol=pipe</code> .                                                                                                                                               |            |
| <code>--port</code>                      | TCP portnumber to connect to.                                                                                                                                                          |            |

| Format                                 | Description                                                                                                                                                | Introduced |
|----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|------------|
| <code>--print-defaults</code>          | Print a list of option values supplied by defaults files and exit.                                                                                         |            |
| <code>--process-threads</code>         | Specifies the number of process-threads for the backup operation.                                                                                          |            |
| <code>--progress-interval</code>       | Interval between progress reports in seconds.                                                                                                              |            |
| <code>--protocol</code>                | Connection protocol.                                                                                                                                       |            |
| <code>--read-threads</code>            | Specifies the number of read-threads for the backup operation.                                                                                             |            |
| <code>--relay-log-index</code>         | Specifies the absolute path of the index file on the MySQL server that lists all the used relay log files (for offline backups of slave servers only).     |            |
| <code>--relaylog-info-file</code>      | Specifies the absolute path of the information file in which a slave records information about the relay logs (for offline backups of slave servers only). |            |
| <code>--rename</code>                  | Rename a single table when it is selected by the <code>--include-tables</code> option to be restored                                                       | 3.12.0     |
| <code>--sbt-database-name</code>       | Used as a hint to the Media Management Software (MMS) for the selection of media and policies for tape backup.                                             |            |
| <code>--sbt-environment</code>         | Comma separated list of environment variable assignments to be given to the SBT library.                                                                   |            |
| <code>--sbt-lib-path</code>            | Path name of the SBT library used by software that manages tape backups.                                                                                   |            |
| <code>--secure-auth</code>             | Refuse client connecting to server if it uses old (pre-4.1.1) protocol.                                                                                    |            |
| <code>--shared-memory-base-name</code> | It designates the shared-memory name used by a Windows server to permit clients to connect using shared memory (Windows only).                             |            |
| <code>--show-progress</code>           | Instructs mysqlbackup to periodically output short progress reports known as progress indicators on its operation.                                         |            |
| <code>--skip-binlog</code>             | Do not include binary log files during backup, or do not restore binary log files during restore.                                                          |            |
| <code>--skip-messages-logdir</code>    | Disable logging to teelog file.                                                                                                                            |            |
| <code>--skip-relaylog</code>           | Do not include relay log files during backup, or do not restore relay log files during a restore.                                                          |            |
| <code>--skip-unused-pages</code>       | Skip unused pages in tablespaces when backing up InnoDB tables.                                                                                            |            |
| <code>--slave-info</code>              | Capture information needed to set up an identical slave server.                                                                                            |            |
| <code>--sleep</code>                   | Time to sleep in milliseconds after copying each 1MB of data.                                                                                              |            |
| <code>--socket</code>                  | Socket file to use to connect.                                                                                                                             |            |
| <code>--src-entry</code>               | Identifies a file or directory to extract from a single-file backup.                                                                                       |            |

| Format                                | Description                                                                                                          | Introduced |
|---------------------------------------|----------------------------------------------------------------------------------------------------------------------|------------|
| <code>--ssl</code>                    | Enable SSL for connection (automatically enabled with other <code>--ssl-</code> flags).                              |            |
| <code>--ssl-ca</code>                 | CA file in PEM format (implies <code>--ssl</code> ).                                                                 |            |
| <code>--ssl-capath</code>             | CA directory (check OpenSSL docs, implies <code>--ssl</code> ).                                                      |            |
| <code>--ssl-cert</code>               | X509 cert in PEM format (implies <code>--ssl</code> ).                                                               |            |
| <code>--ssl-cipher</code>             | SSL cipher to use (implies <code>--ssl</code> ).                                                                     |            |
| <code>--ssl-key</code>                | X509 key in PEM format (implies <code>--ssl</code> ).                                                                |            |
| <code>--ssl-verify-server-cert</code> | Verify server's "Common Name" in its cert against hostname used when connecting.                                     |            |
| <code>--start-lsn</code>              | Specifies the highest LSN value included in a previous backup.                                                       |            |
| <code>--suspend-at-end</code>         | Pauses the <code>mysqlbackup</code> command when the backup procedure is close to ending.                            |            |
| <code>--trace [105]</code>            | Trace level of messages by <code>mysqlbackup</code> .                                                                |            |
| <code>--uncompress</code>             | Uncompress the compressed backup before applying log.                                                                |            |
| <code>--use-tts</code>                | Enable selective backup of InnoDB tables using transportable tablespaces (TTS).                                      |            |
| <code>--user</code>                   | Database user name to connect.                                                                                       |            |
| <code>--verbose</code>                | Print more verbose information.                                                                                      |            |
| <code>--version</code>                | Display version information.                                                                                         |            |
| <code>--with-timestamp</code>         | Create a subdirectory underneath the backup directory with a name formed from the timestamp of the backup operation. |            |
| <code>--write-threads</code>          | Specifies the number of write-threads for the backup operation.                                                      |            |

## 13.1 Standard Options

The standard options are options of a general nature, or options that are not classified under any other specific option group.

Here is a list of the standard options:

- The following standard options also exist for the `mysql` command. Full descriptions for these options can be found in the MySQL reference manual, accessible through, e.g., [Server Option and Variable Reference](#). These options must be specified ahead of any other `mysqlbackup` options, including the rest of the standard options:

```

--print-defaults      Print the program argument list and exit.
--no-defaults        Don't read default options from any option file.
--defaults-file=PATH  Only read default options from the given file. It has to be the first option to
--defaults-extra-file=PATH Read this file after the global files are read.
--defaults-group-suffix=str Also read option groups with the usual names and a suffix of str.

```

- The following options are also common between `mysqlbackup` and `mysql`, and full descriptions for them can be found in the MySQL reference manual, accessible through, e.g., [Server Option and](#)

**Variable Reference.** However, `mysqlbackup` does not accept any short forms for these options as `mysql` does (for example, you must use `--help` instead of `-h` for `mysqlbackup`):

```
--help      Display help.
--version   Display version information.
--verbose   Print more verbose information.
--debug     Print debug information.
```

- More standard options are available for `mysqlbackup`:

```
--force      Force overwriting of data, log, or image files, depending on the operation.
--trace=level Trace level of messages by mysqlbackup.
```

`--force`: By default, some of the operations halt rather than overwrite any user data or log files when told to write to existing files. `--force` allows the following:

- Overwriting of InnoDB data and log files during the `apply-log` and `apply-incremental-backup` operations.
- When restoring a TTS backup, changing temporarily the value of `innodb_file_format` on the server, in order to allow restores of per-table InnoDB data files regardless of their format.
- Replacing of an image file during an `backup-to-image` or `backup-dir-to-image` operation.
- Overwriting data in a non-empty target directory with backup data during a `copy-back` operation performed on a full backup (the use of `copy-back-and-apply-log` is not recommended in this case)
- Restoring a partial image backup created with MySQL Enterprise Backup 3.11 or earlier; the `--force` option is required, due to a known issue (Bug# 20485910).
- Restoring a backup onto a server where the directory pointed to by the `.bl` file in the backup (a copy of the `.isl` file from the backed-up server) already contains `.ibd` data files.

For all other operations, the `--force` option is rejected with an error message.

```
--trace=level
```

| Command-Line Format | <code>--trace=LEVEL</code> |             |
|---------------------|----------------------------|-------------|
| Permitted Values    | Type                       | enumeration |
|                     | Default                    | 0           |
|                     | Valid Values               | 0           |
|                     |                            | 1           |
|                     |                            | 2           |
|                     |                            | 3           |

Trace level of `mysqlbackup` messages. The permissible levels, in the order of increasing fineness, are:

- 0 - INFO (information, warnings, errors)
- 1 - FINE (verbose option is enabled)
- 2 - FINER (debug option is enabled)

- 3 - FINEST (includes all low level outputs)

## 13.2 Connection Options

When `mysqlbackup` creates a backup, it sends SQL commands to MySQL server using a database connection. The general connection details are the same as described in [Connecting to the MySQL Server](#) in the MySQL Reference Manual.

As part of the `mysqlbackup` invocation, specify the appropriate `--user`, `--password`, `--port`, and/or `--socket` options that are necessary to connect to the MySQL server.

You can specify the following connection-specific options in the `[mysqlbackup]` or `[client]` sections of a MySQL configuration file, or through `mysqlbackup` command-line options. `mysqlbackup` reads your default configuration files and then the `my.cnf` file specified on the command line.



### Note

- `mysqlbackup` reads only `--user`, `--password`, `--port`, and `--socket` options from the `[client]` group, and ignores any other connection options.
- If you do not provide a value for the `--password`, the command prompts for one from the keyboard.
- The `--host` option is allowed in the configuration file for compatibility, but currently it has no effect. `mysqlbackup` always connects to the local server's IP address.

```
Options Common to mysqld
=====

--login-path=name
--port=port-num
--protocol=tcp|socket|pipe|memory
--pipe [ alias for --protocol=pipe ]
--user=name [ short option: -u ]
--host=hostname
--socket=name
--shared-memory-base-name=value [Windows only]
--character-sets-dir=PATH
--default-character-set=VALUE
--secure-auth [ Don't connect to pre-4.1.1 server ]
--password[=value] [ short option: -p ]
--connect_timeout
--ssl [ Enable SSL for connection ]
--ssl-key=file_name
--ssl-cert=file_name
--ssl-ca=file_name
--ssl-capath=directory_name
--ssl-cipher=cipher_list
--ssl-verify-server-cert

Connection Options Specific to mysqlbackup
=====

--no-connection
--connect-if-online
```

Most other connection parameters used by the `mysql` command are recognized, but silently ignored. Unknown connection parameters cause `mysqlbackup` to stop.

The following connections options are specific to `mysqlbackup`:

- `--no-connection`

The `--no-connection` option supersedes the other connection options and uses file-level operations to perform the backup. When you use this option, you must specify in the configuration file or on the command line many options whose values are normally retrieved automatically through the database connection.



#### Warning

This option also turns on the `--no-history-logging` and `--no-locking` options, which might result in inconsistencies in non-InnoDB data if the tables are modified during the backup operation. It might also affect subsequent incremental backups; see the description for the `--incremental-base` option for details.

- `--connect-if-online`

By default, a database connection is used for backup operations both during the initial stage to retrieve source repository configuration, and to lock tables while copying non-InnoDB data. This option allows `mysqlbackup` to make connection attempts in both phases, but continues even if the connection cannot be established. If a connection cannot be established, the processing is the same as with the `--no-connection` option. This option can be useful in emergency situations: for example, when the database server goes down during the backup operation.

## 13.3 Server Repository Options

These repository options specify various parameters related to the database server, from which the data is backed up or to which a backup is restored.

These options are used only with the following operations:

- Backup creation operations: `backup`, `backup-and-apply-log`, `backup-to-image`.
- Restore operations: `copy-back`, `copy-back-and-apply-log`.

When a database connection is available during a backup, the parameters describing the source repository are ignored, overridden by the corresponding values retrieved from the database connection.

For information about how these options are specified for the MySQL server, click the option names to see the descriptions in the MySQL Reference Manual.

- `datadir=PATH`

This is the `datadir` value used by the MySQL instance. The `.frm` files reside here inside subdirectories named after the databases inside the instance.

When a database connection exists, the value is retrieved automatically and overrides any value you specify. This is a crucial parameter for both the MySQL server and MySQL Enterprise Backup.

- `innodb_data_home_dir=PATH`

Specifies the directory where InnoDB data files reside. Usually the same as `datadir`, but can be different.

This parameter, together with `innodb_data_file_path=SIZE`, determines where the InnoDB data files such as `ibdata1`, `ibdata2`, and so on, are situated within the MySQL server.

Typically, you do not need to specify this option, because its value is retrieved automatically using the database connection.

Its value is derived as follows:

- If `innodb_data_home_dir` is not specified, it inherits the value of `datadir`.
- If `innodb_data_home_dir` is a relative path, the path is located relative to (that is, underneath) the `datadir` value.
- An `innodb_data_home_dir` of `" "` refers to the `/` root directory.
- If `innodb_data_home_dir` is an absolute path, its value is used as-is.
- `innodb_data_file_path=VALUE`

Specifies InnoDB data file names and sizes. Examples:

```
ibdata1:32M:ibdata2:32M:autoextend
/abs/path/ibdata1:32M:autoextend
innodb-dir/ibdata1:32M:autoextend
```

When a database connection exists, the value is retrieved automatically and overrides any value you specify.

This parameter together with `innodb_data_home_dir` determines where the InnoDB data files (such as `ibdata1`, `ibdata2`, and so on) reside in server repository.

Typically, you do not need to specify this option, because its value is retrieved automatically using the database connection. If no database connection is available, you must specify it.

Whether the initial file name begins with a `/` character or not, the files are located relative to the `innodb_data_home_dir` value.

- `innodb_log_group_home_dir=PATH`

Specifies where InnoDB logs reside within the server repository. Usually the same as `datadir`, but can be different.

Its value is derived as follows:

- If `innodb_log_group_home_dir` is not specified, it inherits the value of `datadir`.
- If `innodb_log_group_home_dir` is a relative path, the path is taken to be relative to (that is, underneath) the `datadir` value.
- If `innodb_log_group_home_dir` is an absolute path, its value is used as-is.
- `innodb_log_files_in_group=N`

Specifies the number of InnoDB log files before being rotated.

Typically, you do not need to specify this option, because its value is retrieved automatically using the database connection. If no database connection is available, you must specify it.

When a database connection exists, the value is retrieved automatically and overrides any value you specify.



- `innodb_log_file_size=SIZE`

Specifies maximum single InnoDB log file size before switching to next log file. Example: 20M.

Typically, you do not need to specify this option, because its value is retrieved automatically using the database connection. If no database connection is available, you must specify it.

When a database connection exists, the value is retrieved automatically and overrides any value you specify.

- `innodb_page_size=SIZE`

Specifies the page size for all InnoDB tablespaces.

Typically, you do not need to specify this option, because its value is retrieved automatically using the database connection. If no database connection is available, you must specify it.

When a database connection exists, the value is retrieved automatically and overrides any value you specify.

- `innodb_checksum_algorithm=NAME`

Specifies the name of the checksum algorithm used for validating InnoDB tablespaces. Default is `innodb`.

Typically, you do not need to specify this option, because its value is retrieved automatically using the database connection. If no database connection is available, you must specify it.

When a database connection exists, the value is retrieved automatically and overrides any value you specify.

## 13.4 Backup Repository Options

These options specify various parameters related to the backup image or directory, or to how the backup will be restored. Typically, `--backup-image` and `--backup-dir` are the only options from the group that you need to specify when using `mysqlbackup`.

The backup repository options are used with the following operations:

- Backup operations: `backup`, `backup-and-apply-log`, `backup-to-image`.
- Restore operations: `copy-back`, `copy-back-and-apply-log`.

The backup repository options are divided into two groups: the first one determines the structure of the backup, and the second one provides information on the original structure of the data on the backed-up server for future operations on the backup.

The following options determine the structure of the backup:

- `--backup-image=IMAGE`

|                            |                                   |                        |
|----------------------------|-----------------------------------|------------------------|
| <b>Command-Line Format</b> | <code>--backup-image=IMAGE</code> |                        |
| <b>Permitted Values</b>    | <b>Type</b>                       | <code>file name</code> |

Specify the path name of the file used for a single-file backup, restore, or another [single-file operation](#). Note that if `--backup-image` does not give a full path name, `mysqlbackup` will actually take the value of `--backup-image` as a path relative to the directory specified by `--backup-dir`, and thus store

the single-file backup under `--backup-dir` (or, if the `--with-timestamp` option is used, under a subdirectory created under `--backup-dir`, which bears the timestamp in its name).

By default, the single-file backup is streamed to standard output, so that you can pipe it directly to other commands such as a tape backup or an `ssh`-related network command.

You can optionally prefix the image name with `file:` to signify a file I/O (the default). For tape backups, prefix the image name with `sbt:`. See [Section 4.3.1.2, “Backing Up to Tape”](#) for details about tape backups.

- `--backup-dir=PATH`

Same as `--backup-dir`. The [backup directory](#) under which the backup data and metadata are stored, permanently or temporarily. It is a crucial parameter required for most kinds of backup and restore operations.

The option is used differently for different operations and under different situations:

- *For backup to a single file (including incremental, compressed, encrypted, and cloud backups):* Use `--backup-dir` to supply a temporary folder to save the backup metadata (including the `mysqlbackup` message log, the start and end [LSN](#), and so on) and some temporary output. The backup data, together with a copy of the metadata, will be stored in a single file whose name is specified with the `--backup-image` option. Note that, however, if `--backup-image` does not give a full path name, `mysqlbackup` will actually take the value of `--backup-image` as a path relative to the directory specified by `--backup-dir`, and thus store the single-file backup under `--backup-dir` (or, if the `--with-timestamp` option is used, under a subdirectory created under `--backup-dir`, which bears the timestamp in its name).
- *For backup to a directory:* Use `--backup-dir` to specify the directory to store the backup data and metadata (including the `mysqlbackup` message log, the start and end [LSN](#), and so on). The directory specified by `--backup-dir` cannot be a subdirectory of the directory specified by `--datadir`.

When the `--with-timestamp` option is also specified, an additional level of subdirectory, with the timestamp in its name, is created under `--backup-dir` (see description for the `--with-timestamp` option for details). Unless the `--with-timestamp` option is used, the directory specified by `--backup-dir` must be empty, or the backup operation will fail with an error.

- *For restoring a single-file backup (including incremental, compressed, encrypted, and cloud backups):* When using `copy-back-and-apply-log` to restore a single-file backup, use `--backup-dir` to supply a temporary folder to store the temporary data of the restore operation. The directory specified by `--backup-dir` should be empty—if a non-empty directory is used, the restore operation will still be carried out, but the restore data might be corrupted.

For restoring an incremental backup to a single-file, you can use `--incremental-backup-dir` for the same purpose.

- *For restoring a backup directory:* Use `--backup-dir` to specify the location of the backup directory, from which data will be restored to the server.
- `backup_innodb_data_home_dir=PATH`

The directory under which the backup's InnoDB data files are to be stored. Specify the option if you want to put the data files at somewhere other than the default location (which is `backup-dir/datadir`). If the value of the parameter is different from `backup-dir/datadir`, it is stored into the `backup-my.cnf` file as `innodb_data_home_dir` for information, so that `mysqlbackup` can understand the structure of the backup when it performs various operations on the backup. Together

with `backup_innodb_data_file_path` option, it determines the actual file paths of the InnoDB data files inside the backup.

The value for the parameter is derived as follows:

- If `backup_innodb_data_home_dir` is not specified, its value will be `backup-dir/datadir`.
- If `backup_innodb_data_home_dir` is an absolute path, its value is used as-is.
- If `backup_innodb_data_home_dir` is a relative path, the path is taken to be relative to (that is, underneath) `backup-dir`.
- An empty string ("" ) for `backup_innodb_data_home_dir` means the value of `innodb_data_file_path` is to be taken as an absolute path..

This parameter is applicable only for backup operations; during a restore, the InnoDB data files are restored under the data directory specified by `--datadir`, unless another location is specified using the `--innodb_data_home_dir` option during restore.

- `backup_innodb_data_file_path=VALUE`

The InnoDB data file names and sizes. Examples:

```
ibdata1:32M:ibdata2:32M:autoextend
/abs/path/ibdata1:32M:autoextend
innodb-dir/ibdata1:32M:autoextend
```

This parameter, together with `backup_innodb_data_home_dir`, determines where the InnoDB data files are stored within the backup repository. Any file path specified with this option is taken to be relative to the value of the `backup_innodb_data_home_dir` option (that is true even if the file path is specified in the form of an absolute path, like `/abs/path/ibdata1:32M:autoextend`). To specify truly absolute paths for InnoDB data files in the backup with this option, you must set the `backup_innodb_data_home` option to "" [empty string], in addition to using an absolute path for this option.

When the parameter is not specified, it inherits the value from the value of the `innodb_data_file_path` option on the backed-up server. If both the source and destination of the backup attempt to use the same absolute paths that resolves to the same files, the backup is cancelled.

The value of the parameter is stored into the `backup-my.cnf` file as `innodb_data_file_path` for information, so that `mysqlbackup` can understand the structure of the backup when it performs various operations on the backup.

- `backup_innodb_log_group_home_dir=PATH`

The directory under which the backup's InnoDB logs will be stored. Specify this option only if you want to put the logs at somewhere other than the default location (which is `backup-dir/datadir`). If the value of the parameter is different from `backup-dir/datadir`, it is stored in the `backup-my.cnf` file as `innodb_log_group_home_dir` for information, so that `mysqlbackup` can understand the structure of the backup when it performs various operations on the backup. Note that while you can specify a directory for saving the logs, the names of the log files are fixed and not reconfigurable.

This parameter is applicable only for backup operations; during a restore, the InnoDB log files are restored under the data directory specified by `--datadir`, unless another location is specified using the `--innodb_log_group_home_dir` option during restore. The value of the parameter is derived as follows:

- If `backup_innodb_log_group_home_dir` is not specified, its value will be `backup-dir/datadir`.
- If `backup_innodb_log_group_home_dir` is an absolute path, its value is used as-is.
- If `backup_innodb_log_group_home_dir` is a relative path, the path is taken to be relative to (that is, underneath) `backup-dir`.
- An empty string ("" ) for the option produces an error.
- `backup_innodb_undo_directory=PATH`

The relative or absolute directory path where separate tablespaces are created for the InnoDB undo logs during the backup. When unspecified, the option takes up the same value as `backup_innodb_log_group_home_dir`; specify this option only if you want to put the undo logs at some other location. If the value of the parameter is different from `backup-dir/datadir`, it is stored in the `backup-my.cnf` file as `innodb_undo_directory` for information, so that `mysqlbackup` can understand the structure of the backup when it performs various operations on the backup.

This parameter is applicable only for backup operations; during a restore, the InnoDB undo log tablespaces are restored under the data directory specified by `--datadir`, unless another location is specified by the `--innodb_undo_directory` option during restore.

- `--with-timestamp`

Creates a subdirectory underneath the backup directory, with a name formed with the timestamp of the backup operation. It is useful for maintaining a single backup directory containing many backup snapshots put under different subdirectories.

Default: no timestamped subdirectory is created. To reuse the same backup directory for a new backup without using this option, either remove the previous backup files manually or, for a single-file backup, specify the `--force` [105] option to overwrite the old backup file.

The following parameters provide information on the original structure of the data on the backed-up server for future operations on the backup, but do not affect the structure of the backup itself:

- `backup_innodb_log_files_in_group=N`

The number of InnoDB log files in a log group on the restored server. See the description for `innodb_log_files_in_group` in the MySQL server manual. The value for this parameter, saved as `innodb_log_files_in_group` in the `backup-my.cnf` file, is derived as follows:

- Use the `backup_innodb_log_files_in_group` value from command line or configuration file, if specified.
- Else, use the `innodb_log_files_in_group` value from the backed-up server, if it is an online backup.
- Else, use the `innodb_log_files_in_group` value from the `mysqlbackup` command line or configuration file.
- `backup_innodb_log_file_size=SIZE`

The maximum single InnoDB log file size in backup before switching to next log file, on the restored server. See the description for `innodb_log_file_size` in the MySQL server manual. The value for this parameter, saved as `innodb_log_file_size` in the `backup-my.cnf` file, is derived as follows:

- Use the `backup_innodb_log_file_size` value from command line or configuration file, if specified.
- Else, use the `innodb_log_file_size` value from the backed-up server, if it is an online backup.
- Else, use the specified `innodb_log_file_size` value from the `mysqlbackup` command line or configuration file.

- `backup_innodb_page_size=SIZE`

Specifies, for an offline backup, the page size for all InnoDB tablespaces on the restored server. This option should be specified carefully, because the page size must be the same as that of the backed up MySQL instance, or the backup could become useless. For an online backup, the value is taken from the value of the `innodb_page_size` option on the backed-up server.

Value of this option is saved in the `backup-my.cnf` file, to be used for restoring the database.

- `backup_innodb_undo_logs=NUMBER`

Specifies, for an offline backup, the number of rollback segments in the InnoDB system tablespace on the restored server. This option should be specified carefully, because the value must be the same as that of `innodb_undo_logs` on the backed-up MySQL instance, or the backup could become useless. For an online backup, the value for the parameter is taken from the value of the `innodb_undo_logs` option on the backed-up server.

- `backup_innodb_undo_tablespaces=NUMBER`

Specifies the number of tablespace files that the undo logs are divided between, when you use a non-zero `backup_innodb_undo_logs` setting. This option should be specified carefully, because the value must be the same as that of `innodb_undo_tablespaces` on the backed-up MySQL instance, or the backup could become useless. For an online backup, the value for the parameter is taken from the value of the `innodb_undo_tablespaces` option on the backed-up server. By default, all the undo logs are part of the system tablespace, and the system tablespace will always contain one undo tablespace in addition to those configured by `innodb_undo_tablespaces`.

- `backup_innodb_checksum_algorithm=NAME`

Specifies, for an offline backup, the name of the checksum algorithm used for validating the InnoDB tablespaces on the restored server. This option should be specified carefully, because the checksum algorithm must be the same use on the backed-up MySQL instance, or the backup could become useless. For an online backup, the value is taken from the value of the `innodb_checksum_algorithm` option on the backed-up server.

Default value of the option is “innodb”.

Value of this option is saved in the `backup-my.cnf` file, to be used for restoring the database.

## 13.5 Metadata Options

These options control the generation of metadata about backups. Some metadata is stored in the backup directory, other metadata is stored in tables within the `mysql` database of the backed-up instance.

- `--no-history-logging`

Turns off the recording of backup progress and history in logging tables inside the backed-up database. See [Section 10.4, “Using the MySQL Enterprise Backup Logs”](#) for details about these tables.

Default: history logging is enabled. When `--no-connection` is specified, history logging is automatically disabled. When `--connect-if-online` is specified, history logging only works if a database connection is successfully established during the backup.

- `--comments=STRING`

|                            |                                |                     |
|----------------------------|--------------------------------|---------------------|
| <b>Command-Line Format</b> | <code>--comments=STRING</code> |                     |
| <b>Permitted Values</b>    | <b>Type</b>                    | <code>string</code> |

Specifies a comment string that describes or identifies the backup. Surround multi-word comments with appropriate quotation marks. The string is saved in a file `meta/comments.txt` in the backup. For example: `--comments="Backup of HR data on 2010/12/10"`.

- `--comments-file=PATH`

|                            |                                   |                        |
|----------------------------|-----------------------------------|------------------------|
| <b>Command-Line Format</b> | <code>--comments-file=PATH</code> |                        |
| <b>Permitted Values</b>    | <b>Type</b>                       | <code>file name</code> |

Specifies path to a file containing comments describing the backup. This file is saved as `meta/comments.txt` in the backup. For example: `--comments-file=/path/to/comments.txt`.

This option overrides the `--comments` option if both are specified.

## 13.6 Compression Options

For an overview on backup compression, see [Section 4.3.4, “Making a Compressed Backup”](#).

- `--compress`

Create backup in compressed format. For a regular backup, among all the storage engines supported by MySQL, only data files of the InnoDB format are compressed, and they bear the `.ibz` extension after the compression. Similarly, for a single-image backup, only data files of the InnoDB format inside the backup image are compressed. The binary log and relay log files are compressed and saved with the `.bz` extension when being included in a compressed backup.

Default: compression is disabled.

- `--compress-method=ALGORITHM`

|                            |                                          |                          |
|----------------------------|------------------------------------------|--------------------------|
| <b>Command-Line Format</b> | <code>--compress-method=ALGORITHM</code> |                          |
| <b>Permitted Values</b>    | <b>Type</b>                              | <code>enumeration</code> |
|                            | <b>Default</b>                           | <code>lz4</code>         |
|                            | <b>Valid Values</b>                      | <code>zlib</code>        |
|                            |                                          | <code>lz4</code>         |
|                            |                                          | <code>lzma</code>        |

Specifies the compression algorithm. The supported arguments for the option and the algorithms they represent are:

- `lz4`: LZ4 r109. Out of the three algorithms that are supported, this is the most efficient one, typically taking the shortest backup and restore times with the lowest CPU cost. See [lz4—Extremely Fast Compression algorithm](#) for more details, including a comparison with other compression algorithms.

- **lzma**: LZMA 9.20. Out of the three supported algorithms, this typically provides the highest compression ratio; but it is also far more expensive in terms of CPU cost than the other two options. Thus we do not recommend this for active systems, but only for off-hour or inactive databases, or where I/O rates are extremely low.
- **zlib**: ZLIB v1.2.3. This is in between the other two supported algorithms in terms of both speed and compression ratio. ZLIB was the only compression algorithm available for MySQL Enterprise Backup versions prior to 3.10.

Default: lz4. Explicitly specifying a value for the option through a configuration file or command line automatically enables the `--compress` option.

- `--compress-level=LEVEL`

| Command-Line Format | <code>--compress-level=LEVEL</code> |         |
|---------------------|-------------------------------------|---------|
| Permitted Values    | Type                                | numeric |
|                     | Default                             | 1       |
|                     | Min Value                           | 0       |
|                     | Max Value                           | 9       |

Specifies the level of compression, ranging from “0” to “9”: “0” disables compression; “1” is fastest compression, and “9” is highest (and slowest) compression. The option is only meaningful for compression using the ZLIB or LZMA algorithm; it is ignored when any other algorithms are selected by the `--compress-method` option.

Default: 1 (lowest and fastest compression). Explicitly specifying a non-zero value through a configuration file or command line automatically enables the `--compress` option.

- `--uncompress`

When used with the `apply-log` or `copy-back-and-apply-log` operation, uncompresses the compressed backup before applying the InnoDB log.

## 13.7 Incremental Backup Options

For an overview of incremental backups and usage examples for these options, see [Section 4.3.3, “Making a Differential or Incremental Backup”](#).

To take an incremental backup, specify the `--incremental` or `--incremental-with-redo-log-only`, along with the `--incremental-backup-dir` option. All InnoDB data modified after the specified LSN is copied in the incremental backup. Depending on whether `--incremental` or `--incremental-with-redo-log-only` is used, other options are required or recommended.

- `--incremental`

Specifies that the associated `backup` or `backup-to-image` operation is **incremental**. Also requires either the `--incremental-base` option, or the combination of the `--start-lsn` and `--incremental-backup-dir` options.

Only InnoDB tables are backed up incrementally. By default, all non-InnoDB and `.frm` files are included into the incremental backup and in their fullness. To exclude non-InnoDB data in an incremental backup, use the `--only-innodb` option.



- `--incremental-with-redo-log-only`

Specifies that an [incremental](#) backup is to be created using only the redo log. This alternate type of incremental backup has different performance characteristics and operational limitations comparing to backups created with the `--incremental` option; see [Creating Incremental Backups Using Only the Redo Log](#) for a discussion on their differences.

To use this option, you also need to specify the `--incremental-base` option, or a combination of the `--start-lsn` and `--incremental-backup-dir` options. Just like with the `--incremental` option, only InnoDB tables are backed up incrementally. By default, all non-InnoDB and `.frm` files are included in incremental backup and in their fullness. To exclude non-InnoDB data in an incremental backup, use the `--only-innodb` option.

- `--incremental-base=mode:argument`

|                            |                                               |                     |
|----------------------------|-----------------------------------------------|---------------------|
| <b>Command-Line Format</b> | <code>--incremental-base=mode:argument</code> |                     |
| <b>Permitted Values</b>    | <b>Type</b>                                   | <code>string</code> |

With this option, the `mysqlbackup` retrieves the information needed to perform incremental backups from the metadata inside the backup directory rather than from the `--start-lsn` option. It saves you from having to specify an ever-changing, unpredictable [LSN](#) value when doing a succession of incremental backups. Instead, you specify a way to locate the previous backup directory through the combination of *mode* and *argument* in the option syntax. The alternatives are:

- `dir:directory_path`

You specify the prefix `dir:` followed by a directory path. The path argument points to the directory where the data from the previous backup is stored. With the first incremental backup, you specify the directory holding the full backup; with the second incremental backup, you specify the directory holding the first incremental backup, and so on.

- `history:last_backup`

You specify the prefix `history:` followed by `last_backup`, the only valid argument for this mode. This makes `mysqlbackup` query the `end_lsn` value from the last successful [non-TTS backup](#) as recorded in the `backup_history` table of the server instance that is being backed up.



#### Note

If the last full or partial backup made was a [TTS backup](#), `mysqlbackup` skips it, and keeps searching the backup history until it finds the last [non-TTS backup](#) and then returns its `end_lsn` value.



#### Warning

Do not use the `history:` mode if the previous backup was a full backup taken with the `--no-connection` option, which always turns off the recording of backup history and might cause errors for a subsequent incremental backup using this mode of the `--incremental-base` option.

- `--start-lsn=LSN`

|                            |                              |                      |
|----------------------------|------------------------------|----------------------|
| <b>Command-Line Format</b> | <code>--start-lsn=LSN</code> |                      |
| <b>Permitted Values</b>    | <b>Type</b>                  | <code>numeric</code> |



In an [incremental backup](#), specifies the highest [LSN](#) value included in a previous backup. You can get this value from the output of the previous backup operation, or from the [backup\\_history](#) table's [end\\_lsn](#) column for the previous backup operation. Always used in combination with the [--incremental](#) option; not needed when you use the [--incremental-base](#) option; not recommended when you use the [--incremental-with-redo-log-only](#) mechanism for incremental backups.

- [--incremental-backup-dir=PATH](#)

Specifies the location for data of an incremental backup. When creating or restoring an incremental backup, the option serves the same function as [--backup-dir](#) does for backups and restores in general, and the option can in fact be used interchangeably with [--backup-dir](#). See the description for [--backup-dir](#) for details.

The location specified by [--incremental-backup-dir](#) is the same location you will specify with [--incremental-base](#) when you use that option for a subsequent incremental backup.

## 13.8 Partial Backup and Restore Options



### Note

Since MySQL Enterprise Backup 3.10, the two options [--include-tables](#) and [--exclude-tables](#) have been introduced. These were intended for replacing the older options of [--include](#), [--databases](#), [--databases-list-file](#), and [--only-innodb-with-frm](#), which are incompatible with the new options and will be deprecated in future releases. For references purpose, we have included information on the older options at the end of this section in [Legacy Partial Backup Options](#).

To select specific data to be backed up or restored, use the partial backup and restore options described in this section.

For an overview of partial backup and usage examples on the following options, see [Section 4.3.5, “Making a Partial Backup”](#). See also [Section 5.1.4, “Restoring Backups Created with the --use-tts Option”](#), on selective restore of tables from a backup.

- [--include-tables=REGEXP](#)

|                            |                                         |                        |
|----------------------------|-----------------------------------------|------------------------|
| <b>Command-Line Format</b> | <a href="#">--include-tables=REGEXP</a> |                        |
| <b>Permitted Values</b>    | <b>Type</b>                             | <a href="#">string</a> |

Include for backup or restoration only those tables (both innodb and non-innodb) whose fully qualified names (in the form of [db\\_name.table\\_name](#)) match the regular expression [REGEXP](#). The regular expression syntax used is the extended form specified in the POSIX 1003.2 standard. For example, [--include-tables=^mydb\.t\[12\]\\$](#) matches the tables [t1](#) and [t2](#) in the database [mydb](#). On Unix-like systems, quote the regular expression appropriately to prevent interpretation of shell meta-characters. [mysqlbackup](#) throws an error when the option is used without a regular expression being supplied with it.

While [mysqlbackup](#) understands the MySQL convention of quoting the database or the table name (or both) by backticks (see [Schema Object Names](#)), there is no need to include the backticks in the regular expression for [--include-tables](#).

While the option can be used for different kinds of backups, selective restore is only supported for backups created using [transportable tablespace \(TTS\)](#) (that is, backups created with the [--use-tts](#)

option). The option can also be used with the `backup-dir-to-image` and `image-to-backup-dir` commands to select tables when creating or unpacking a backup image.

The option cannot be used together with the legacy `--include`, `--databases`, `--databases-list-file`, or `--only-innodb-with-frm` option.

When used together with the `--exclude-tables` option, `--include-tables` is applied first, meaning `mysqlbackup` first selects all tables specified by `--include-tables` and then excludes from the set those tables specified by `--exclude-tables`.

- `--exclude-tables=REGEXP`

|                            |                                      |                     |
|----------------------------|--------------------------------------|---------------------|
| <b>Command-Line Format</b> | <code>--exclude-tables=REGEXP</code> |                     |
| <b>Permitted Values</b>    | <b>Type</b>                          | <code>string</code> |

Exclude for backup or restoration all tables (both innodb and non-innodb) whose fully qualified names (in the form of `db_name.table_name`) match the regular expression `REGEXP`. The regular expression syntax is the extended form specified in the POSIX 1003.2 standard. For example, `--exclude-tables=^mydb\.t[12]$` matches the tables `t1` and `t2` in the database `mydb`. On Unix-like systems, quote the regular expression appropriately to prevent interpretation of shell meta-characters. `mysqlbackup` throws an error when the option is used without a regular expression being supplied with it.

While `mysqlbackup` understands the MySQL convention of quoting the database or the table name (or both) by backticks (see [Schema Object Names](#)), there is no need to include the backticks in the regular expression for `--exclude-tables`.

While the option can be used for different kinds of backups, selective restore is only supported for backups created using [transportable tablespaces \(TTS\)](#) (that is, backups created with the `--use-tts` option). The option can also be used with the `backup-dir-to-image` and `image-to-backup-dir` commands to select tables when creating or unpacking a backup image.

The option cannot be used together with the `--include`, `--databases`, `--databases-list-file`, or `--only-innodb-with-frm` option.

When used together with the `--include-tables` option, `--include-tables` is applied first, meaning `mysqlbackup` first select all tables specified by `--include-tables`, and then exclude from the set those tables specified by `--exclude-tables` for backup.

- `--only-known-file-types`

For back up only. By default, all files in the database subdirectories under the data directory of the server are included in the backup (see [Table 1.1, “Types of Files in a Backup”](#) for details). If the `--only-known-file-types` option is specified, `mysqlbackup` only backs up those types of files that are data files for MySQL or its built-in storage engines, which have the following extensions:

- `.ARM`: Archive storage engine metadata
- `.ARZ`: Archive storage engine data
- `.CSM`: CSV storage engine data
- `.CSV`: CSV storage engine data
- `.frm`: table definitions

- `.ibd`: InnoDB tablespace created using the file-per-table mode
- `.MRG`: Merge storage engine references to other tables
- `.MYD`: MyISAM data
- `.MYI`: MyISAM indexes
- `.opt`: database configuration information
- `.par`: partition definitions
- `.TRG`: trigger parameters
- `.TRN`: trigger namespace information
- `--only-innodb`

For back up only. Back up only InnoDB data and log files. All files created by other storage engines are excluded. Typically used when no connection to `mysqld` is allowed or when there is no need to copy MyISAM files.

The option is not compatible with the `--slave-info` option.

Default: backups include files from all storage engines.

- `--use-tts[={with-minimum-locking|with-full-locking}]`

| Command-Line Format | <code>--use-tts[={with-minimum-locking with-full-locking}]</code> |                                   |
|---------------------|-------------------------------------------------------------------|-----------------------------------|
| Permitted Values    | Type                                                              | enumeration                       |
|                     | Default                                                           | <code>with-minimum-locking</code> |
|                     | Valid Values                                                      | <code>with-minimum-locking</code> |
|                     |                                                                   | <code>with-full-locking</code>    |

Enable selective backup of InnoDB tables using [transportable tablespaces \(TTS\)](#). This is to be used in conjunction with the `--include-tables` and `--exclude-tables` options for selecting the InnoDB tables to be backed up by regular expressions. Using [TTS](#) for backups offers the following advantages:

- Backups can be restored to a different server
- The [system tablespace](#) is not backed up, saving disk space and I/O resources
- Data consistency of the tables is managed by MySQL Enterprise Backup

However, the option has the following limitations:

- Supports only MySQL version 5.6 and after (as earlier versions of MySQL do not support [TTS](#))
- Backing up and restoring of partitioned tables are only supported for MySQL 5.7.4 and later. Also, individual partitions cannot be selectively backed up or restored. Tables selected by the `--include-tables` and `--exclude-tables` options are always backed up or restored in full.
- Can only backup tables that are stored in their own individual tablespaces (i.e., tables created with the [innodb\\_file\\_per\\_table](#) option enabled)

- Cannot be used for incremental backups
  - Does not include the binary log or the relay log in the backup
- See also [Section B.1, “Limitations of MySQL Enterprise Backup”](#) for some more minor limitations.

There are two possible values for the option:

- `with-minimum-locking`: Hot copies of the selected tables are backed up, and the tables are then locked in read-only mode while the [redo log](#) (with only the portion containing the relevant changes made after the hot backup) is being included in the backup. Any tables created during the locking phase are ignored.
- `with-full-locking`: The selected tables are locked in read-only mode while they are being backed up. The [redo log](#) is not included in the backup. Any tables created during the locking phase are ignored.



### Note

Due to a known issue, when creating a backup using TTS for a server containing tables with a mix of the Antelope and Barracuda file formats, do NOT apply full locking on the tables.

Default: `with-minimum-locking`

To use the `--use-tts` option, extra privileges are required of the user through which `mysqlbackup` connects to the server; see [Section 4.1.2, “Grant MySQL Privileges to Backup Administrator”](#) for details.

There are some special requirements for restoring backups created with the `--use-tts` option; see [Section 5.1.4, “Restoring Backups Created with the --use-tts Option”](#) for details.

- `--rename=“old_table_name to new_table_name”`

Rename a single table when it is selected by the `--include-tables` or `--exclude-tables` option (or both together) to be restored to a database from a backup created using the `--use-tts` option. The table named `old_table_name` is renamed to `new_table_name`. Note that when using the option:

- The `--include-tables` or `--exclude-tables` option (or both together) must be used in the restore command for the `--rename` option to work, unless there is only one table in the backup. Also, the `--include-tables` or `--exclude-tables` option (or both together) should specify one and only one table for restore when `--rename` is used, or the restore will fail.
- `old_table_name` and `new_table_name` can be fully qualified (containing the database names, in the format of `db_name.tb_name`) or not. Regular expressions are not accepted for `old_table_name` and `new_table_name`.
- The restore fails if `old_table_name` does not match with the table specified using the `--include-tables` or `--exclude-tables` option (or both together), or if `new_table_name` already exists in the target database.
- The requirements listed in [Section 5.1.4, “Restoring Backups Created with the --use-tts Option”](#) apply.

See [Section 5.1.4, “Restoring Backups Created with the --use-tts Option”](#), for more information on selective restores, and for an example of table renaming.

## Legacy Partial Backup Options



### Important

Information in this subsection is only for using the legacy options of `--include`, `--databases`, `--databases-list-file`, and `--only-innodb-with-frm`, which will be deprecated in the upcoming issues. For creating partial backups, it is strongly recommended that the new options of `--include-tables` and `--exclude-tables` be used instead. Note that you cannot combine the legacy and the new partial-backup options in a single command.

Besides the legacy options, some other options are also discussed below, but the information is only for using the options together with the legacy partial-backup options.

For an overview of partial backups and usage examples for these legacy options, see [Making a Partial Backup with the Legacy Options](#).

- `--include=REGEXP`

This option is for filtering InnoDB tables for backup. The InnoDB tables' fully qualified names are checked against the regular expression specified by the option. If the REGEXP matches `db_name.table_name`, the table is included. The regular expression syntax used is the extended form specified in the POSIX 1003.2 standard. For example, `--include=mydb\.t[12]` matches the tables `t1` and `t2` in the database `mydb`. `mysqlbackup` throws an error when the option is used without a regular expression being supplied with it.

This option only applies to InnoDB tables created with the MySQL option `innodb_file_per_table` enabled (which is the default setting for MySQL 5.6 and after), in which case the tables are in separate files that can be included or excluded from the backup. All tables in the InnoDB system tablespace are always backed up.

When no InnoDB table names match the specified regular expression, an error is thrown with a message indicating there are no matches.

Default: Backs up all InnoDB tables.



### Note

This option does not filter non-InnoDB tables, for which options like `--databases` and `--databases-list-file` can be used.



### Important

This option does not filter the `.frm` files associated with InnoDB tables, meaning that regardless of the option's value, all the `.frm` files for all InnoDB tables are always backed up unless they are excluded by other options. Those `.frm` files for InnoDB tables that are not backed up should be deleted before the database backup is restored. See [Making a Partial Backup with the Legacy Options](#) for details.

- `--databases=LIST`

Specifies the list of non-InnoDB tables to back up. The argument specifies a space-separated list of database or table names of the following form:

```
"db_name[.table_name] db_name1[.table_name1] ...".
```

If the specified values do not match any database or table, then no non-InnoDB data files are backed up. See [Making a Partial Backup with the Legacy Options \[50\]](#) for details.

By default, all non-InnoDB tables from all databases are backed up.



#### Note

The option has no filtering effects on the InnoDB data files (`.ibd` files) for the databases or tables it specifies. To filter InnoDB data files, use the `--include` option instead.

- `--databases-list-file=PATH`

Specifies the pathname of a file that lists the non-InnoDB tables to be backed up. The file contains entries for databases or fully qualified table names separated by newline or space. The format of the entries is the same as for the `--databases` option:

```
db_name[.table_name]
db_name1[.table_name1]
...
```

Remove any white spaces surrounding the database or table names, as the white spaces are not removed automatically. Begin a line with the `#` character to include a comment. No regular expressions are allowed.

If the specified entries do not match any database or table, then no non-InnoDB data files are backed up.



#### Note

The option has no filtering effects on the InnoDB data files (`.ibd` files) for the databases or tables it specifies. To filter InnoDB data files, use the `--include` option instead.

- `--only-innodb-with-frm[={all|related}]`

Back up only InnoDB data, log files, and the `.frm` files associated with the InnoDB tables.

- `--only-innodb-with-frm=all` includes the `.frm` files for all InnoDB tables in the backup.
- `--only-innodb-with-frm=related`, in combination with the `--include` option, copies only the `.frm` files for the tables that are included in the partial backup.
- `--only-innodb-with-frm` with no argument is the same as `--only-innodb-with-frm=related`.



#### Note

For incremental backups, even only changed `.ibd` files are backed up, `.frm` files associated with *all* specified InnoDB tables are included.

This option saves you having to script the backup step for InnoDB `.frm` files, which you would normally do while the server is put into a read-only state by a `FLUSH TABLES WITH READ LOCK` statement. The `.frm` files are copied without putting the server into a read-only state, so that the backup operation is a true [hot backup](#) and does not interrupt database processing. You must ensure that no `ALTER TABLE` or other DDL statements change `.frm` files for InnoDB tables while the backup is in progress.

If `mysqlbackup` detects changes to any relevant `.frm` files during the backup operation, it halts with an error. If it is not practical to forbid DDL on InnoDB tables during the backup operation, use the `--only-innodb` option instead and use the traditional method of copying the `.frm` files while the server is locked.

All files created by other storage engines are excluded. Typically used when no connection to `mysqld` is allowed or when there is no need to copy MyISAM files, for example, when you are sure there are no DDL changes during the backup. See [Making a Partial Backup with the Legacy Options](#) for instructions and examples.

The option is not compatible with the `--slave-info` option.

Default: backups include files from all storage engines.

- `--use-tts[={with-minimum-locking|with-full-locking}]`

Enable selective backup of InnoDB tables using [transportable tablespaces \(TTS\)](#). This is to be used in conjunction with the `--include` option, which selects the InnoDB tables to be backed up by a regular expression. Using [TTS](#) for backups offers the following advantages:

- Backups can be restored to a different server
- The [system tablespace](#) is not backed up, saving disk space and I/O resources
- Data consistency of the tables is managed by MySQL Enterprise Backup

However, the option has the following limitations:

- Supports only MySQL version 5.6 and after (as earlier versions of MySQL do not support [TTS](#))
- Can only backup tables that are stored in their own individual tablespaces (i.e., tables created with the [innodb\\_file\\_per\\_table](#) option enabled)
- Cannot back up partitioned tables
- Cannot restore tables selectively from the backup
- Cannot be used for incremental backups

There are two possible values for the option:

- `with-minimum-locking`: Hot copies of the selected tables are backed up, and the tables are then locked in read-only mode while the [redo log](#) (with only the portion containing the relevant changes made after the hot backup) is being included in the backup. Any tables created during the locking phase are ignored.
- `with-full-locking`: The selected tables are locked in read-only mode while they are being backed up. The [redo log](#) is not included in the backup. Any tables created during the locking phase are ignored.

Default: back up with minimum locking

There are some special requirements for restoring backups created with the `--use-tts` option; see the [explanations](#) in [Section 5.1, “Performing a Restore Operation”](#) for details.

## 13.9 Single-File Backup Options

These options are associated with single-file backups. You use them in combination with the `mysqlbackup` commands `backup-to-image`, `image-to-backup-dir`, `backup-dir-to-image`, `copy-back-and-apply-log`, `list-image`, and `extract`. For usage examples, see [Section 4.3.1, “Making a Single-File Backup”](#).

- `--backup-image=IMAGE`

See description of the option in [Section 13.4, “Backup Repository Options”](#)

- `--src-entry=PATH`

|                            |                               |                        |
|----------------------------|-------------------------------|------------------------|
| <b>Command-Line Format</b> | <code>--src-entry=PATH</code> |                        |
| <b>Permitted Values</b>    | <b>Type</b>                   | <code>path name</code> |

Identifies a file or directory to extract from a single-file backup. This option is used with the `extract` command. If the argument is a directory, all its files and subdirectory contents are extracted. No pattern matching expression is allowed for the argument. Optionally, you can also specify the `--dst-entry` option to extract the file or directory in a location different from its original path name.

For example: `src-entry=meta/comments.txt` extracts only one file, `comments.txt`, while `src-entry=meta` extracts the entire directory tree for the `meta` subdirectory.

Default: All entries are extracted.

- `--dst-entry=PATH`

|                            |                               |                        |
|----------------------------|-------------------------------|------------------------|
| <b>Command-Line Format</b> | <code>--dst-entry=PATH</code> |                        |
| <b>Permitted Values</b>    | <b>Type</b>                   | <code>path name</code> |

Used with single-file backups to extract a single file or directory to a user-specified path. Use of this option requires specifying the `--src-entry` option. This option specifies the destination path for the entry selected from the backup image by `--src-entry`. The entry could point to a single file or single directory. For example, to retrieve the comments file from a backup image and store it as `/tmp/my-comments.txt`, use a command like the following:

```
mysqlbackup --src-entry=meta/comments.txt \
--dst-entry=/tmp/my-comments.txt \
--backup-image=/var/myimage.bki extract
```

Similarly, to extract all the contents of the `meta` directory in a single-file backup as `/data/my-meta`, use a command like the following:

```
mysqlbackup --src-entry=meta \
--dst-entry=/data/my-meta \
--backup-image=/var/myimage.bki extract
```

The specified path is a simple path name without any wildcard expansion or regular expressions.

Default: By default, original pathnames are used to create files in the local file system.

- `--sbt-database-name=NAME`

|                            |                                       |                     |
|----------------------------|---------------------------------------|---------------------|
| <b>Command-Line Format</b> | <code>--sbt-database-name=NAME</code> |                     |
| <b>Permitted Values</b>    | <b>Type</b>                           | <code>string</code> |



|  |                |       |
|--|----------------|-------|
|  | <b>Default</b> | MySQL |
|--|----------------|-------|

For tape backups, this option can be used as a hint to the Media Management Software (MMS) for the selection of media and policies. This name has nothing to do with MySQL database names. It is a term used by the MMS. See [Section 4.3.1.2, “Backing Up to Tape”](#) for usage details.

- `--sbt-lib-path=PATH`

|                            |                                  |                        |
|----------------------------|----------------------------------|------------------------|
| <b>Command-Line Format</b> | <code>--sbt-lib-path=PATH</code> |                        |
| <b>Permitted Values</b>    | <b>Type</b>                      | <code>file name</code> |

Path name of the SBT library used by the software that manages tape backups. If this is not specified, operating system-specific search methods are used to locate `libobk.so` (UNIX) or `orasbt.dll` (Windows). See [Section 4.3.1.2, “Backing Up to Tape”](#) for usage details.

- `--sbt-environment=VAR=value,...`

|                            |                                                                                  |                     |
|----------------------------|----------------------------------------------------------------------------------|---------------------|
| <b>Command-Line Format</b> | <code>--sbt-environment=VAR1=value1[,VAR2=value2[,...]] SBT API provider)</code> |                     |
| <b>Permitted Values</b>    | <b>Type</b>                                                                      | <code>string</code> |

Passes product-specific environment variables to Oracle Secure Backup or another SBT-compliant backup management product, as an alternative to setting and unsetting environment variables before and after each `mysqlbackup` invocation.

The parameter to this option is a comma-separated list of key-value pairs, using syntax similar to that of the RMAN tool for the Oracle Database. For example, `--sbt-environment=VAR1=val1,VAR2=val2,VAR3=val3`.

Consult the documentation for your backup management product to see which of its features can be controlled through environment variables. For example, the Oracle Secure Backup product [defines environment variables](#) such as `OB_MEDIA_FAMILY`, `OB_DEVICE`, and `OB_RESOURCE_WAIT_TIME`. You might set such variables with the `mysqlbackup` by specifying an option such as `--sbt-environment="OB_MEDIA_FAMILY=my_mf,OB_DEVICE=my_tape"`.

If the argument string contains any whitespace or special characters recognized by the command shell, enclose the entire argument string in quotation marks. To escape an equals sign or comma, use the `\` character. For example, `--sbt-environment="VAR1=multiple words,VAR2=<angle_brackets>,VAR3=2+2\=4"`.

- `--disable-manifest`

Disable generation of [manifest](#) files for a backup operation, which are `backup_create.xml` and `backup_content.xml` present in the `meta` subdirectory.

## 13.10 Performance / Scalability / Capacity Options

These options limit the resources used by the backup process, in order to minimize backup overhead for busy or huge databases, or specify behaviors of the process when encountering resource issues.

- `--number-of-buffers=num_buffers`

|                            |                                         |                      |
|----------------------------|-----------------------------------------|----------------------|
| <b>Command-Line Format</b> | <code>--number-of-buffers=NUMBER</code> |                      |
| <b>Permitted Values</b>    | <b>Type</b>                             | <code>numeric</code> |

|  |                  |    |
|--|------------------|----|
|  | <b>Default</b>   | 14 |
|  | <b>Min Value</b> | 1  |

Specifies the number of buffers, each 16MB in size, to use during multithreaded options.

Use a high number for CPU-intensive processing such as backup, particularly when using compression. Use a low number for disk-intensive processing such as restoring a backup. This value should be at least as high as the number of read threads or write threads, depending on the type of operation.

Default: currently 14.

For compression or incremental backup operations, the buffer size is slightly more than 16MB to accommodate the headers.

One additional buffer is used for single-file incremental backup and single-file compressed backup.

Compressed backup, compressed single-file backup, and uncompress apply-log operations require one additional buffer for each process thread.

If you change the number of read, write, and processing threads, you can experiment with changing this value so that it is slightly larger than the total number of threads specified by those other options. See [Section 7.1, “Optimizing Backup Performance”](#) and [Section 7.2, “Optimizing Restore Performance”](#) for additional advice about recommended combinations of values for this and other performance-related options for various hardware configurations, such as RAID or non-RAID storage devices.

- `--read-threads=num_threads`

|                            |                                    |         |
|----------------------------|------------------------------------|---------|
| <b>Command-Line Format</b> | <code>--read-threads=NUMBER</code> |         |
| <b>Permitted Values</b>    | <b>Type</b>                        | numeric |
|                            | <b>Default</b>                     | 1       |
|                            | <b>Min Value</b>                   | 1       |
|                            | <b>Max Value</b>                   | 15      |

Specifies the number of threads to use for reading data from disk.

Default: currently 1. This default applies to these kinds of operations: `copy-back`, `extract`, and `backup`. If you specify a value of 0, it is silently adjusted to 1. The maximum is 15; if you supply a negative value, it is silently adjusted to 15. For `apply-log` operations, the number of read threads is always 1 regardless of this option setting. See [Section 7.1, “Optimizing Backup Performance”](#) and [Section 7.2, “Optimizing Restore Performance”](#) for advice about recommended combinations of values for `--read-threads`, `--process-threads`, and `--write-threads` for various hardware configurations, such as RAID or non-RAID storage devices.

- `--process-threads=num_threads`

|                            |                                       |
|----------------------------|---------------------------------------|
| <b>Command-Line Format</b> | <code>--process-threads=NUMBER</code> |
|----------------------------|---------------------------------------|

|                         |                  |         |
|-------------------------|------------------|---------|
| <b>Permitted Values</b> | <b>Type</b>      | numeric |
|                         | <b>Default</b>   | 6       |
|                         | <b>Min Value</b> | 1       |
|                         | <b>Max Value</b> | 15      |

Specifies the number of threads to use for processing data, such as compressing or uncompressing backup files.

Default: currently 6. This default applies to these kinds of operations: `extract`, and `backup`. It is ignored when you use any of the options `--incremental-with-redo-log-only`, `apply-incremental-backup`, `copy-back`, or `backup-dir-to-image`.

If you specify a value of 0, it is silently adjusted to 1. The maximum is 15; if you supply a negative value, it is silently adjusted to 15. For `apply-log` operations, the number of process threads is always 1 regardless of this option setting. See [Section 7.1, “Optimizing Backup Performance”](#) and [Section 7.2, “Optimizing Restore Performance”](#) for advice about recommended combinations of values for `--read-threads`, `--process-threads`, and `--write-threads` for various hardware configurations, such as RAID or non-RAID storage devices.

- `--write-threads=num_threads`

|                            |                                     |         |
|----------------------------|-------------------------------------|---------|
| <b>Command-Line Format</b> | <code>--write-threads=NUMBER</code> |         |
| <b>Permitted Values</b>    | <b>Type</b>                         | numeric |
|                            | <b>Default</b>                      | 1       |
|                            | <b>Min Value</b>                    | 1       |
|                            | <b>Max Value</b>                    | 15      |

Specifies the number of threads to use for writing data to disk.

Default: currently 1. This default applies to these kinds of operations: `copy-back`, `extract`, and `backup`. It is ignored when you use any of the single-file backup options `list-image` or `validate`.

If you specify a value of 0, it is silently adjusted to 1. The maximum is 15; if you supply a negative value, it is silently adjusted to 15. For `apply-log` operations, the number of write threads is always 0 regardless of this option setting. See [Section 7.1, “Optimizing Backup Performance”](#) and [Section 7.2, “Optimizing Restore Performance”](#) for advice about recommended combinations of values for `--read-threads`, `--process-threads`, and `--write-threads` for various hardware configurations, such as RAID or non-RAID storage devices.

- `--limit-memory=MB`

|                            |                                |         |
|----------------------------|--------------------------------|---------|
| <b>Command-Line Format</b> | <code>--limit-memory=MB</code> |         |
| <b>Permitted Values</b>    | <b>Type</b>                    | numeric |

|  |                  |                                                                                  |
|--|------------------|----------------------------------------------------------------------------------|
|  | <b>Default</b>   | 100 for <code>apply-log</code> (without uncompression), 300 for other operations |
|  | <b>Min Value</b> | 0                                                                                |
|  | <b>Max Value</b> | 999999                                                                           |

Specify maximum memory in megabytes that can be used by `mysqlbackup`. Formerly applied only to `apply-log` operation, but in MySQL Enterprise Backup 3.8 and higher it applies to all operations. Do not include any suffixes such as `mb` or `kb` in the option value.

Default: 100 for `apply-log` not used with `--uncompress`, 300 for all operations (in megabytes).

The memory limit specified by this option also caps the number of 16MB buffers available for multithreaded processing. For example, with a 300 MB limit, the maximum number of buffers is 18. If additional buffers are required because you increase the values for `--read-threads`, `--process-threads`, `--write-threads`, and/or `--number-of-buffers`, increase the `--limit-memory` value proportionally.

- `--sleep=MS`

|                            |                         |         |
|----------------------------|-------------------------|---------|
| <b>Command-Line Format</b> | <code>--sleep=MS</code> |         |
| <b>Permitted Values</b>    | <b>Type</b>             | numeric |
|                            | <b>Default</b>          | 0       |

Specify the number in milliseconds to sleep after copying a certain amount of data from InnoDB tables. Each block of data is 1024 InnoDB data pages, typically totalling 16MB. This is to limit the CPU and I/O overhead on the database server.

Default: 0 (no voluntary sleeps).

- `--no-locking`

Disables locking during backup of non-InnoDB files, even if a connection is available. Can be used to copy non-InnoDB data with less disruption to normal database processing. There could be inconsistencies in non-InnoDB data if any changes are made while those files are being backed up.

- `--page-reread-time=MS`

|                            |                                    |         |
|----------------------------|------------------------------------|---------|
| <b>Command-Line Format</b> | <code>--page-reread-time=MS</code> |         |
| <b>Permitted Values</b>    | <b>Type</b>                        | numeric |
|                            | <b>Default</b>                     | 100     |

Interval in milliseconds that `mysqlbackup` waits before re-reading a `page` that fails a checksum test. A busy server could be writing a page at the same moment that `mysqlbackup` is reading it. Can be a floating-point number, such as 0.05 meaning 50 microseconds. Best possible resolution is 1 microsecond, but it could be worse on some platforms. Default is 100 milliseconds (0.1 seconds).

- `--page-reread-count=retry_limit`

|                            |                                         |         |
|----------------------------|-----------------------------------------|---------|
| <b>Command-Line Format</b> | <code>--page-reread-count=number</code> |         |
| <b>Permitted Values</b>    | <b>Type</b>                             | numeric |

|  |                |     |
|--|----------------|-----|
|  | <b>Default</b> | 500 |
|--|----------------|-----|

Maximum number of re-read attempts, when a [page](#) fails a checksum test. A busy server could be writing a page at the same moment that [mysqlbackup](#) is reading it. If the same page fails this many checksum tests consecutively, with a pause based on the [--page-reread-time](#) option between each attempt, the backup fails. Default is 500.

- [--on-disk-full={abort|abort\\_and\\_remove|warn}](#)

|                            |                                       |                                  |
|----------------------------|---------------------------------------|----------------------------------|
| <b>Command-Line Format</b> | <a href="#">--on-disk-full=option</a> |                                  |
| <b>Permitted Values</b>    | <b>Type</b>                           | <a href="#">enumeration</a>      |
|                            | <b>Default</b>                        | <a href="#">abort</a>            |
|                            | <b>Valid Values</b>                   | <a href="#">abort</a>            |
|                            |                                       | <a href="#">warn</a>             |
|                            |                                       | <a href="#">abort_and_remove</a> |

Specifies the behavior when a backup process encounters a disk-full condition. This option is only for backup operations ([backup](#), [backup-and-apply-log](#), and [backup-to-image](#)).

- [abort](#): Abort backup, without removing the backup directory. The disk remains full.
- [abort\\_and\\_remove](#): Abort backup and remove the backup directory.
- [warn](#): Write a warning message every 30 seconds and retry backup until disk space becomes available.  
Default: [abort](#).

- [--skip-unused-pages](#)

Skip unused pages in tablespaces when backing up InnoDB tables. This option is applicable to the [backup](#) and [backup-to-image](#) operations, but not to [incremental](#) backups. The option is ignored by the [backup-and-apply-log](#) operation.

Note that backups created with the [--skip-unused-pages](#) option cannot be restored using [copy-back-and-apply-log](#).

Unused pages are free pages often caused by bulk delete of data. By skipping the unused pages during backups, this option can reduce the backup sizes and thus the required disk space and I/O resources for the operations. However, subsequent [apply-log](#) operations on the backups will take more time to complete, as the unused pages are inserted back into the tables during the operations.

- [--skip-binlog](#)

Skip including the binary log files in the backup during a backup operation, or skip copying the binary log files onto a server during a restore operation.

Binary log files, together with the binary log index file, are included by default for all kinds of online backups (full, incremental, compressed, partial, single-file, etc.). See [Table 1.1, “Types of Files in a Backup”](#), for details. Use this option to skip backing up binary logs if resource, performance, or other issues arise. Also, if any binary log files are missing on the server you are backing up, you should use this option to avoid [mysqlbackup](#) throwing an error for the missing files.

The binary log files and the binary log index file, when included in a backup, are always copied into the restored server's data directory during a restore operation; if that is not the behavior you desire, use this option to skip the restoring of the binary log.

- `--skip-relaylog`

When working with a slave server, skip including the relay log files in the backup during a backup operation, or skip copying the relay log files onto a server during a restore operation.

Relay log files, together with the relay log index file and the `master.info` and the `slave.info` files, are included by default for all kinds of online backups (full, incremental, compressed, partial, single-file, etc.) of a slave server. See [Table 1.1, “Types of Files in a Backup”](#), for details. Use this option to skip backing up relay logs if resource, performance, or other issues arise.



#### Note

If a user runs a `FLUSH LOGS` statement while backup is in progress for a slave, the backup process will fail. Use the `--skip-relaylog` option if you expect a `FLUSH LOGS` statement will be run during the backup and it is not necessary to include the relay logs in the backup.

The relay log files and the files backed up together with them, when included in a backup, are always copied into the restored server's data directory during a restore operation; if that is not the behavior you desire, use this option to skip the restoring of the relay log.

- `--log-bin-index[=PATH]`

|                            |                                       |                                           |
|----------------------------|---------------------------------------|-------------------------------------------|
| <b>Command-Line Format</b> | <code>--log-bin-index=FILENAME</code> |                                           |
| <b>Permitted Values</b>    | <b>Type</b>                           | <code>file name</code>                    |
|                            | <b>Default</b>                        | <code>data_dir/host_name-bin.index</code> |

For MySQL 5.5 and earlier, as well as all offline backups: specify the absolute path (including file name and extension) of the index file on the MySQL server that lists all the used binary log files, if it is different from the default path given below, in order to include the binary log files in the backup.

Default: `data_dir/host_name-bin.index`.

- `--relay-log-index[=PATH]`

|                            |                                         |                                                 |
|----------------------------|-----------------------------------------|-------------------------------------------------|
| <b>Command-Line Format</b> | <code>--relay-log-index=FILENAME</code> |                                                 |
| <b>Permitted Values</b>    | <b>Type</b>                             | <code>file name</code>                          |
|                            | <b>Default</b>                          | <code>data_dir/host_name-relay-bin.index</code> |

For offline backups of slave servers only: specify the absolute path (including file name and extension) of the index file on the MySQL server that lists all the used relay log files, if it is different from the default path given below, in order to include the relay log files in the backup.

Default: `data_dir/host_name-relay-bin.index`.

- `--master-info-file[=PATH]`

|                            |                                          |                        |
|----------------------------|------------------------------------------|------------------------|
| <b>Command-Line Format</b> | <code>--master-info-file=FILENAME</code> |                        |
| <b>Permitted Values</b>    | <b>Type</b>                              | <code>file name</code> |

|  |                |                                   |
|--|----------------|-----------------------------------|
|  | <b>Default</b> | <code>data_dir/master.info</code> |
|--|----------------|-----------------------------------|

For offline backups of slave servers only: specify the absolute path (including file name and extension) of the information file in which a slave records information about its master, if it is different from the default path given below, in order to include the information file in the backup.

Default: `data_dir/master.info`.

- `--relaylog-info-file[=PATH]`

|                            |                                            |                                      |
|----------------------------|--------------------------------------------|--------------------------------------|
| <b>Command-Line Format</b> | <code>--relaylog-info-file=FILENAME</code> |                                      |
| <b>Permitted Values</b>    | <b>Type</b>                                | <code>file name</code>               |
|                            | <b>Default</b>                             | <code>data_dir/relay-log.info</code> |

For offline backups of slave servers only: specify the absolute path (including file name and extension) of the information file in which a slave records information about the relay logs, if it is different from the default path given below, in order to include the information file in the backup.

Default: `data_dir/relay-log.info`.

- `--optimistic-time[=DATE-TIME]`

|                            |                                          |                     |
|----------------------------|------------------------------------------|---------------------|
| <b>Command-Line Format</b> | <code>--optimistic-time=DATE-TIME</code> |                     |
| <b>Permitted Values</b>    | <b>Type</b>                              | <code>string</code> |
|                            | <b>Default</b>                           | <code>now</code>    |

Perform an optimistic backup with the value specified with the option as the “optimistic time”—a time after which the tables that have not been modified are taken as “inactive tables.” The “inactive tables” are believed to be unlikely to change during the backup process. The inactive tables are backed up in the optimistic phase of the backup, and all other tables are backed up in the normal phase. See [Section 4.3.6, “Making an Optimistic Backup”](#) for details on the concept, use cases, and command samples for an optimistic backup.

Accepted formats for specifying the option include:

- `now`: This includes all tables into the optimistic phase of the backup process. It is the default value for the option when no value is specified.
- `{Number}{Unit}`: Indicates the optimistic time as a time at a certain duration into the past. `{Unit}` can be any one of `years`, `months`, `hours`, and `minutes`. Some examples for option strings in this format include: `5years`, `2days`, `13months`, `23hours`, and `35minutes`.
- A date-time format in any of the following forms: `YYMMDD`, `YYYYMMDD`, `YYMMDDHHMMSS`, `YYYYMMDDHHMMSSYY-MM-DD`, `YYYY-MM-DD`, `YY-MM-DD`, or `HH.MM.SSYYYYMMDDTHHMMSS`, as designated by the ISO 8601 standard.

When both the `optimistic-time` and the `optimistic-busy-tables` options are used and they come into conflict on determining which tables are to be backed up in the optimistic phase, `optimistic-busy-tables` takes precedence over `optimistic-time`.

- `--optimistic-busy-tables=REGEXP`

|                            |                                              |  |
|----------------------------|----------------------------------------------|--|
| <b>Command-Line Format</b> | <code>--optimistic-busy-tables=REGEXP</code> |  |
|----------------------------|----------------------------------------------|--|

|                  |      |        |
|------------------|------|--------|
| Permitted Values | Type | string |
|------------------|------|--------|

Perform an optimistic backup, using the regular expression specified with the option to select tables that will be skipped in the first phase of an optimistic backup, because they are likely to be modified during the backup process. Tables whose fully qualified names (in the form of `database_name.table_name`) are matched by the regular expression are taken as “busy tables”, which will be backed up in the second or the “normal” phase of the backup. Tables whose fully qualified names are NOT matched by the regular expression are taken as “inactive tables”, which will be backed up in the first or the “optimistic” phase of the backup. See [Section 4.3.6, “Making an Optimistic Backup”](#) for details on the concept, use cases, and command samples for an optimistic backup.

MySQL Enterprise Backup will throw an error if the option is used but no regular expression is supplied with it.

When both the `optimistic-time` and the `optimistic-busy-tables` options are used and they come into conflict on determining which tables are to be “optimistic”, `optimistic-busy-tables` takes precedence over `optimistic-time`.

## 13.11 Message Logging Options

`mysqlbackup` writes important progress and error information to the `stderr` stream. The information is often very valuable for tracking down problems that occur during an operation. Starting from MySQL Enterprise Backup 3.9, the output to the `stderr` stream is also saved to a log file by default (for most `mysqlbackup` operations), so that the error information can be easily accessed in any debug process.

The message logging works like a `tee` process on a Unix-like system, in which the output of a program is split to be both displayed and saved to a file. The log file thus produced is named in the following format: `MEB_timestamp_operation.log`, where `operation` is the `mysqlbackup` operation that was run (e.g., `backup`, `apply-log`, etc.), and `timestamp` is the date and time at which the operation was run. Here are some examples of names for the log files:

```
MEB_2013-06-24.16-32-43_backup.log
MEB_2013-06-28.11-07-18_apply_log.log
MEB_2013-06-29.10-08-06_list_image.log
```

The following options control the message logging function:

- `--skip-messages-logdir`

Skip message logging. Logging is turned on by default (except for the `list-image` and `validate` operations; see the description for the `--messages-logdir` option for details), and it is turned off by this option.

- `--messages-logdir=path`

|                     |                                     |                              |
|---------------------|-------------------------------------|------------------------------|
| Command-Line Format | <code>--messages-logdir=PATH</code> |                              |
| Permitted Values    | Type                                | directory name               |
|                     | Default                             | <code>backup_dir/meta</code> |

Specifies the path name of an existing directory for storing the message log. If the specified directory does not exist, message logging fails and returns an error message. When this option is omitted, the default directory of `backup_dir/meta` is used, where `backup_dir` is the directory specified with the `--backup-dir` option.



**Note**

Use this option to turn on message logging for the `list-image` and `validate` operations. Message logging is turned off by default for the two operations, because they do not modify any files and a message log is usually not required for debugging them. And because the default path name of `backup_dir/meta` is not meaningful for the two operations, this option is required for both turning on message logging and for supplying the path name of a directory in which to save the log file. However, if the `--skip-messages-logdir` option is also specified, it takes precedence and message logging is skipped.

The following are some examples showing how the message logging is controlled.

This creates a log file for the `backup` operation in the directory `/home/backup_dir/meta` due to the default settings:

```
mysqlbackup -uroot --port=3306 --backup-dir=/home/backup_dir backup
```

This skips message logging for the `backup` operation:

```
mysqlbackup -uroot --port=3306 --backup-dir=/home/backup_dir \
--skip-messages-logdir backup
```

This creates a log file for the `apply-log` operation in an existing directory named `/home/teelog_dir`, rather than the default location:

```
mysqlbackup -uroot --port=3306 --backup-dir=/home/backup_dir \
--messages-logdir=/home/teelog_dir apply-log
```

This creates a log file for the `list-image` operation in an existing directory named `/home/teelog_dir`:

```
mysqlbackup -uroot --port=3306 --backup-image=/backup/my.mbi \
--messages-logdir=/home/teelog_dir list-image
```

## 13.12 Progress Report Options

There are two options for controlling the progress reporting function of `mysqlbackup`: `--show-progress` and `--progress-interval`:

- `--show-progress[={stderr|stdout|file:FILENAME|fifo:FIFONAME|table|variable}]`

| Command-Line Format | <code>--show-progress[=destinations]</code> |               |
|---------------------|---------------------------------------------|---------------|
| Permitted Values    | Type                                        | enumeration   |
|                     | Valid Values                                | stderr        |
|                     |                                             | stdout        |
|                     |                                             | file:FILENAME |
|                     |                                             | fifo:FIFONAME |
|                     |                                             | table         |
|                     |                                             | variable      |

The option instructs `mysqlbackup` to periodically output short progress reports known as progress indicators on its operation.

The argument of the option controls the destination to which the progress indicators are sent:

- **stderr**: Progress indicators are sent to the standard error stream. The report is embedded in a time-stamped `mysqlbackup` INFO message. For example:

```
130607 12:22:38 mysqlbackup: INFO: Progress: 191 of 191 MB; state: Completed
```

- **stdout**: Progress indicators are sent to the standard output stream. A single newline character is printed after each progress indicator.
- **file:FILENAME**: Progress indicators are sent to a file. Each new progress report overwrites the file, and the file contains the most recent progress indicator followed by a single newline character.
- **fifo:FIFONAME**: Progress indicators are sent to a file system FIFO. A single newline character is printed after each progress indicator.



### Warning

If there is no process reading the FIFO, the `mysqlbackup` process hangs at the end of the execution.

- **table**: Progress indicators are sent to the `mysql.backup_progress` table. This requires a connection to the MySQL server, and therefore, only works when backing up a running MySQL instance. `mysqlbackup` first adds one row of the progress report to the `mysql.backup_progress` table, and then updates the row afterwards with the latest progress indicator. The progress indicator is stored in the `current_status` column of the table.

If the backup locks the MySQL instance (for example, by issuing a `FLUSH TABLES WITH READ LOCK` statement), the progress reports are not delivered to the `mysql.backup_progress` table until the MySQL instance is unlocked.

- **variable**: Progress indicators are sent to the system variable `backup_progress`.



### Warning

The system variable `backup_progress` is not yet defined for the MySQL Server. Users need to create their own plugin to define the variable. See [The MySQL Plugin API](#) for more information on user plugins.

When there is no argument specified for `--show-progress`, progress indicators are sent to `stderr`.

Progress can be reported to multiple destinations by specifying the `--show-progress` option several times on the command line. For example the following command line reports progress of the backup command to `stderr` and to a file called `meh_output`:

```
mysqlbackup --show-progress --show-progress=file:meh_output --backup-dir=/full-backup
backup
```

The progress indicators are short strings that indicate how far the execution of a `mysqlbackup` operation has progressed. A progress indicator consists of one or more meters that measure the progress of the operation. For example:

```
Progress: 100 of 1450 MB; state: Copying .ibd files
```

This shows that 100 megabytes of a total of 1450 megabytes have been copied or processed so far, and `mysqlbackup` is currently copying InnoDB data files (`.ibd` files).

The progress indicator string begins with `Progress:`, followed by one or more meters measuring the progress. If multiple meters are present, they are separated by semicolons. The different types of meters include:

- Total data meter: It is always the first meter in the progress indicator. It is in the format of:

```
DATA of TOTAL UNIT
```

*DATA* and *TOTAL* are unsigned decimal integers, and *UNIT* is either MB (megabytes), KB (kilobytes), or bytes (1MB=1024KB and 1KB=1024 bytes).

The total data meter has two slightly different meanings depending on the `mysqlbackup` operation:

- The amount of data copied or processed and the total amount of data to be copied or processed by the `mysqlbackup` operation. For example:

```
Progress: 200 of 1450 MB
```

When the operation is for, e.g., `backup`, the indicator means 200MB is copied of 1450MB. But when the operation is for, e.g., `validate` or `incremental`, it means 200MB is processed out of 1450MB.

- Total amount of data copied or processed and an estimate for the total that will be copied by the end of the operation. The estimated total is updated as per the data on the server, as the execution of the command progresses.

For some operations such as `backup`, it is not possible to know exactly at the start of the execution how much data will be copied or processed. Therefore, the total data meter shows the estimated amount of the total data for a backup. The estimate is updated during the execution of the command. For example:

```
Progress: 200 of 1450 MB
```

is followed by:

```
Progress: 200 of 1550 MB
```

when 100MB of data is added on the server.

If the operation is successful, the final progress indicator shows the actual amount of data copied at the end of the operation.

- Compression meter: It indicates the sliding average of the compression ratio, which is defined for each block of data that is compressed as  $(\text{orig\_size} - \text{compressed\_size}) / \text{orig\_size}$ . For example:

```
compression: 40%
```

This means that after compression, the data takes 40% less space (calculated as an average over the last 10 data blocks).

The compression meter is included in the progress indicator if the `--compress` option is enabled for the `mysqlbackup` operation. The value of the compression meter is undefined until at least 10 data blocks have been compressed. The undefined meter value is denoted by the '-' in the meter:

```
compression: -
```

- **State meter:** It is a short description of the major step the command is currently executing. For example:

```
state: Copying InnoDB data
```

```
state: Waiting for locks
```

```
state: Copying system tablespace
```

```
state: Copying .ibd files
```

```
state: Copying non-InnoDB data
```

```
state: Completed
```

Here are some examples of progress indicators with different meters:

```
Progress: 300 of 1540 MB; state: Waiting for locks
```

```
Progress: 400 of 1450 MB; state: Copying InnoDB data; compression: 30%
```

The exact set of meters included in the progress indicator depends on the command and the options used for it.

- `--progress-interval=SECONDS`

| Command-Line Format | <code>--progress-interval=SECONDS</code> |         |
|---------------------|------------------------------------------|---------|
| Permitted Values    | Type                                     | numeric |
|                     | Default                                  | 2       |
|                     | Min Value                                | 1       |
|                     | Max Value                                | 100000  |

Interval between progress reports in seconds. Default value is two seconds. The shortest interval is 1 second and the longest allowed interval is 100000 seconds.

## 13.13 Encryption Options

These options are for creating encrypted single-file backups and for decrypting them. See [Chapter 8, Encryption for Backups](#) for more details and usage examples for the encryption and decryption functions of MySQL Enterprise Backup.

- `--encrypt`

Encrypt the data when creating a backup image by a `backup-to-image` operation, or when packing a backup directory into a single file with the `backup-dir-to-image` command. It cannot be used with the `backup` or `backup-and-apply-log` command.

- `--decrypt`

Decrypt an encrypted backup image when performing an `extract`, `image-to-backup-dir`, or `copy-back-and-apply-log` operation. It is also used for performing a `validate` or `list-image` operation on an encrypted backup image.

The option cannot be used in a `apply-log`, `backup-and-apply-log`, or `copy-back` operation. For restoration using the `copy-back` command, the encrypted backup image has to be unpacked and decrypted first using the `image-to-backup-dir` or `extract` command, together with the `--decrypt` option.

- `--key=STRING`

|                            |                        |                     |
|----------------------------|------------------------|---------------------|
| <b>Command-Line Format</b> | <code>--key=KEY</code> |                     |
| <b>Permitted Values</b>    | <b>Type</b>            | <code>string</code> |

The symmetric key for encryption and decryption of a backup image. It should be a 256-bit key, encoded as a string of 64 hexadecimal digits. See [Chapter 8, Encryption for Backups](#) on how to create a key. The option is incompatible with the `--key-file` option.

- `--key-file=PATH`

|                            |                              |                        |
|----------------------------|------------------------------|------------------------|
| <b>Command-Line Format</b> | <code>--key-file=FILE</code> |                        |
| <b>Permitted Values</b>    | <b>Type</b>                  | <code>file name</code> |

The pathname to file that contains a 256-bit key, encoded as a string of 64 hexadecimal digits, for encryption and decryption of a backup image. The option is incompatible with the `--key` option.

## 13.14 Cloud Storage Options

These options are for using cloud storage for single-file operations. See [Section 4.3.1.3, “Backing Up to Cloud Storage”](#), and [Section 5.1.5, “Restoring a Backup from Cloud Storage to a MySQL Server”](#), for more information and instructions on using cloud storage with MySQL Enterprise Backup.

- `--cloud-service=SERVICE`

Cloud service for data backup or restoration. Currently, there are two types of cloud storage services supported by `mysqlbackup`, represented by the following values for the options:

- `s3`: Amazon Simple Storage Service (S3)
- `openstack`: OpenStack Object Storage (Swift). MySQL Enterprise Backup 4.0 supports the Swift v1.0 API, and also the OpenStack Identity (Keystone) API v2.0 for authentication. Also supports authentication using Swift's TempAuth system.

- `--cloud-trace`

Print trace information for cloud operations. It works independently of `--trace` [105], which specifies the trace level for the non-cloud operations of `mysqlbackup`. Any non-zero value for the option enables the trace function.

Default value is “0.”

- `--cloud-proxy=proxy-url:port`

Proxy address and port number for overriding the environment's default proxy settings for accessing a cloud storage service.

**Note**

The `list-image` operation can be performed on a cloud backup only if the cloud proxy supports HTTP range headers.

- Options used only for Amazon S3 (using them when the argument for `--cloud-service` is anything other than `s3` will cause `mysqlbackup` to throw an error):

- `--cloud-bucket=S3_BUCKET`

The storage bucket on Amazon S3 for the backup image.

In order to perform cloud backups and restores with the bucket, the user identified by the `--cloud-access-key-id` option must have at least the following permissions on the bucket:

- `s3:ListBucket`: For listing information on items in the bucket.
- `s3:ListBucketMultipartUploads`: For listing multipart uploads in progress to the bucket.
- `s3:GetObject`: For retrieving objects from the bucket.
- `s3:PutObject`: For adding objects to the bucket.
- `--cloud-object-key=S3_OBJECT_KEY`

The Amazon S3 object key for the backup image.

- `--cloud-access-key-id=S3_KEY-ID`

AWS access key ID for logging onto Amazon S3.

- `--cloud-secret-access-key=S3_ACCESS-KEY`

AWS secret access key for the AWS access key id specified with `--cloud-access-key-id`.

- `--cloud-aws-region=S3_REGION`

Region for Amazon Web Services that `mysqlbackup` accesses for S3.

- Options used only for OpenStack Swift (using them when the argument for `--cloud-service` is anything other than `openstack` will cause `mysqlbackup` to throw an error):

- `--cloud-container=SWIFT_CONTAINER`

The Swift container for the backup image.

- `--cloud-object=SWIFT_OBJECT`

The Swift object for the backup image. Note that names of objects in a container have to be unique.

- `--cloud-user-id=SWIFT_OR_KEYSTONE_USER`

User ID for accessing Swift. The user credentials are authenticated using the Swift TempAuth identity system when the `--cloud-tempauth-url` [139] option is used and by the OpenStack Keystone identity service when the `--cloud-identity-url` [139] option is used.

- `--cloud-password=SWIFT_OR_KEYSTONE_PASSWORD`

Password for the user specified by `--cloud-user-id` [138]. The user credentials are authenticated using the Swift TempAuth identity system when the `--cloud-tempauth-url` [139] option is used and by the OpenStack Keystone identity service when the `--cloud-identity-url` [139] option is used.

- `--cloud-tempauth-url=SWIFT_TEMPAUTH-URL`

The TempAuth URL for authenticating user credentials. Either this option or `--cloud-identity-url` [139] (but not both) should be used when accessing a Swift service.

- `--cloud-identity-url=SWIFT_KEystone-URL`

The URL of the Keystone identity service, when it is used for authenticating user credentials. Either this option or `--cloud-tempauth-url` [139] (but not both) should be used when accessing a Swift service.

- `--cloud-tenant=SWIFT_KEystone-TENANT`

The Keystone tenant for the user specified by `--cloud-user-id` [138], when the Keystone identity service is used for authenticating user credentials.

- `--cloud-region=SWIFT_KEystone-REGION`

The Keystone region for the user specified by `--cloud-user-id` [138], when the Keystone identity service is used for authenticating user credentials.

## 13.15 Options for Special Backup Types

These options are for backing up database servers that play specific roles in replication, or contain certain kinds of data that require special care in backing up.

- `--slave-info`

When backing up a replication slave server, this option captures information needed to set up an identical slave server. It creates a file `meta/ibbackup_slave_info` inside the backup directory, containing a `CHANGE MASTER` statement with the binary log position and name of the binary log file of the master server. This information is also printed in the `mysqlbackup` output. To set up a new slave for this master, restore the backup data on another server, start a slave server on the backup data, and issue a `CHANGE MASTER` command with the binary log position saved in the `ibbackup_slave_info` file. See [Section 6.1, “Setting Up a New Replication Slave”](#) for instructions.



### Note

Only use this option when backing up a slave server. Its behavior is undefined when used on a master or non-replication server.

This option is not compatible with the `--no-locking` option; using both options together will make `mysqlbackup` throw an error.

This option is not compatible with the `--only-innodb` or `--only-innodb-with-frm` options.

- `--suspend-at-end`

This option pauses the `mysqlbackup` command when the backup procedure is close to ending. It creates a file called `ibbackup_suspended` in the backup log group home directory and waits until you

delete that file before proceeding. This option is useful to customize locking behavior and backup of non-InnoDB files through custom scripting.

All tables are locked before suspending, putting the database into a read-only state, unless you turn off locking with the `--no-locking` or `--no-connection` option. The `--only-innodb` and `--only-innodb-with-frm` options also prevent the locking step. Because locking all tables could be problematic on a busy server, you might use a combination of `--only-innodb` and `--suspend-at-end` to back up only certain non-InnoDB tables.

- `--exec-when-locked="utility arg1 arg2 ..."`

|                            |                                                         |                     |
|----------------------------|---------------------------------------------------------|---------------------|
| <b>Command-Line Format</b> | <code>--exec-when-locked="utility arg1 arg2 ..."</code> |                     |
| <b>Permitted Values</b>    | <b>Type</b>                                             | <code>string</code> |

You can use this option to run a script that backs up any information that is not included as part of the usual backup. For example, with `--exec-when-locked`, you can use `mysqldump` to back up tables from the MEMORY storage engine, which are not on disk.

Set any variable you want to use within your script before you run `mysqlbackup`. In the following example, the `BACKUP_DIR` environment variable is set to point to the current backup directory (quotes are used for the argument of `--exec-when-locked`, to prevent premature expansion of the variable `BACKUP_DIR`):

On Unix or Linux systems:

```
export BACKUP_DIR=path_to_backupdir
mysqlbackup --exec-when-locked='mysqldump mydb t1 > $BACKUP_DIR/t1.sql' other_options mysqlbackup_command
```

Or on Windows systems:

```
set BACKUP_DIR=path_to_backupdir
mysqlbackup --exec-when-locked="mysqldump mydb t1 > %BACKUP_DIR%/t1.sql" other_options mysqlbackup_command
```

If the utility cannot be executed or returns a non-zero exit status, the whole backup process is cancelled. If you also use the `--suspend-at-end` option, the utility specified by `--exec-when-locked` is executed after the suspension is lifted.



---

## Chapter 14 Configuration Files and Parameters

You can specify `mysqlbackup` options either on the command line or as configuration parameters inside a configuration file. This section describes the use of configuration files.

In general, `mysqlbackup` follows the `mysql` style of processing configuration options: `[mysqlbackup]` and `[client]` group options are passed as command-line options. Any command-line options that you specify when you run `mysqlbackup` override the values from the configuration file, and in the case of duplicate options, the last instance takes precedence. `mysqlbackup` also reads options in the `[mysqld]` group to detect parameters related to the source repository when no connection to `mysqld` is available.

The underscore characters in the `mysqlbackup` option names can be replaced with dashes and treated as synonyms, similar to `mysqld` parameters that use this same convention. (See [Using Options on the Command Line](#) in the MySQL Reference Manual for details.) The MySQL server's reference manual typically lists the parameter names with underscores, to match the output of the `SHOW VARIABLES` statement.

### Options Files

`mysqlbackup` reads the location of the MySQL data to be backed up from the following sources (listed in order of priority):

- The connection information from the running database, whenever possible. Thus, in most cases, you can avoid specifying most options on the command line or in a configuration file.
- Parameters you specify on the `mysqlbackup` command line. You can specify certain options for individual backup jobs this way.
- The MySQL configuration file (by default, `my.cnf` on Unix and `my.ini` on Windows). The parameters are searched for first under the `[mysqlbackup]` group, then under the `[client]` group. You can put common parameters that apply to most of your backup jobs into the configuration file.

### Configuration Files Stored with the Backup Data

Each set of backup data includes a configuration file, `backup-my.cnf`, containing a minimal set of configuration parameters. The `mysqlbackup` command generates this file to record the settings that apply to this backup data. Subsequent operations, such as the `apply-log` process, read options from this file to determine how the backup data is structured.

#### Example 14.1 Example `backup-my.cnf` file

Here is an example `backup-my.cnf` file generated by `mysqlbackup`:

```
[mysqld]
innodb_data_file_path=ibdata1:256M;ibdata2:256M:autoextend
innodb_log_file_size=256M
innodb_log_files_in_group=3
```

All paths in the generated `backup-my.cnf` file point to a single backup directory. For ease of verification and maintenance, you typically store all data for a backup inside a single directory rather than scattered among different directories.

During a backup, the configuration parameters that are required for later stages (such as the restore operation) are recorded in the `backup-my.cnf` file that is generated in the backup directory. Only the minimally required parameters are stored in `backup-my.cnf`, to allow you to restore the backup

to a different environment without having to make extensive changes to that file. For example, the `innodb_data_home_dir` and `innodb_log_group_home_dir` options are omitted from the `backup-my.cnf` file when their values are the same as the value for `backup-dir`.

---

## Part IV Appendixes

---

---

---

---

## Table of Contents

|                                                                                                                                |     |
|--------------------------------------------------------------------------------------------------------------------------------|-----|
| A Frequently Asked Questions for MySQL Enterprise Backup .....                                                                 | 147 |
| B MySQL Enterprise Backup Limitations .....                                                                                    | 149 |
| B.1 Limitations of MySQL Enterprise Backup .....                                                                               | 149 |
| C Compatibility Information for MySQL Enterprise Backup .....                                                                  | 153 |
| C.1 Cross-Platform Compatibility .....                                                                                         | 153 |
| C.2 Compatibility with MySQL Versions .....                                                                                    | 153 |
| C.3 Compatibility with Older MySQL Enterprise Backup .....                                                                     | 153 |
| D Licenses for Third-Party Components .....                                                                                    | 155 |
| D.1 cURL (libcurl) License .....                                                                                               | 155 |
| D.2 GNU General Public License Version 3.0, 29 June 2007 and GCC Runtime Library<br>Exception Version 3.1, 31 March 2009 ..... | 155 |
| D.3 GNU Standard C++ Library (libstdc++) License .....                                                                         | 167 |
| D.4 Google Controlling Master Thread I/O Rate Patch License .....                                                              | 168 |
| D.5 Google SMP Patch License .....                                                                                             | 169 |
| D.6 LZ4 License .....                                                                                                          | 170 |
| D.7 OpenSSL v1.0 License .....                                                                                                 | 170 |
| D.8 Percona Multiple I/O Threads Patch License .....                                                                           | 172 |
| D.9 RE2 License .....                                                                                                          | 172 |
| D.10 RegEX-Spencer Library License .....                                                                                       | 173 |
| D.11 RFC 3174 - US Secure Hash Algorithm 1 (SHA1) License .....                                                                | 173 |
| D.12 zlib License .....                                                                                                        | 174 |
| MySQL Enterprise Backup Glossary .....                                                                                         | 175 |



---

# Appendix A Frequently Asked Questions for MySQL Enterprise Backup

This section lists some common questions about MySQL Enterprise Backup, with answers and pointers to further information.

## Questions

- [A.1: \[147\]](#) What versions of the MySQL server does MySQL Enterprise Backup 4.0 support?
- [A.2: \[147\]](#) What is the big `ibdata` file that is in all the backups?
- [A.3: \[147\]](#) Can I back up non-InnoDB data with MySQL Enterprise Backup?
- [A.4: \[148\]](#) What happens if the apply-log or apply-incremental-backup step is interrupted?
- [A.5: \[148\]](#) Why is the option `--defaults-file` not recognized?
- [A.6: \[148\]](#) Can I back up a database on one OS platform and restore it on another one using MySQL Enterprise Backup?
- [A.7: \[148\]](#) What if I have included the binary log or relay log in my backup but do not want to restore it?
- [A.8: \[148\]](#) What would happen if I start a server directly using a raw directory backup, without running either the copy-back or the apply-log operation?

## Questions and Answers

### A.1: What versions of the MySQL server does MySQL Enterprise Backup 4.0 support?

See [Section C.2, “Compatibility with MySQL Versions”](#) for details of compatibility between different releases of MySQL Enterprise Backup and MySQL Server.

### A.2: What is the big `ibdata` file that is in all the backups?

You might find your backup data taking more space than expected because of a large file with a name such as `ibdata1`. This file represents the InnoDB [system tablespace](#), which grows but never shrinks as a database operates, and is included in every full and incremental backup. To reduce the space taken up by this file in your backup data:

- After doing a [full backup](#), do a succession of [incremental backups](#), which take up less space. The `ibdata1` file in the incremental backups is typically much smaller, containing only the portions of the system tablespace that changed since the full backup.
- Set the configuration option `innodb_file_per_table=1` before creating your biggest or most active InnoDB tables. Those tables are split off from the system tablespaces into separate `.ibd` files; the tables can then be individually included or excluded from backups, and disk space is freed when the tables are dropped or truncated.
- If your system tablespace is very large because you created a high volume of InnoDB data before turning on the `innodb_file_per_table` setting, you might use `mysqldump` to create a dump of your entire server instance, then turn on `innodb_file_per_table` before re-creating the database, so that all the table data is kept outside the system tablespace.

### A.3: Can I back up non-InnoDB data with MySQL Enterprise Backup?

---

While MySQL Enterprise Backup can back up non-InnoDB data (like MYISAM tables), the MySQL server to be backed up must support InnoDB (i.e., the backup process will fail if the server was started up with the `--innodb=OFF` or `--skip-innodb` option), and the server must contain at least one InnoDB table.

**A.4: What happens if the `apply-log` or `apply-incremental-backup` step is interrupted?**

If `mysqlbackup` is interrupted during the `apply-log` or `apply-incremental-backup` stage, the backup data is OK. The file operations performed by those options can be performed multiple times without harming the consistency of the backup data. Just run the same `mysqlbackup` command again, and when it completes successfully, all the necessary changes are present in the backup data.

**A.5: Why is the option `--defaults-file` not recognized?**

When you specify the `--defaults-file` option, it must be the first option going after `mysqlbackup`. Otherwise, the error message makes it look as if the option name is not recognized.

**A.6: Can I back up a database on one OS platform and restore it on another one using MySQL Enterprise Backup?**

See [Section C.1, “Cross-Platform Compatibility”](#) for details.

**A.7: What if I have included the binary log or relay log in my backup but do not want to restore it?**

If you want to skip the restore of the binary log, relay log, or both during a restore, use the `--skip-binlog` option, the `--skip-relaylog` option, or both with your `copy-back` or `copy-back-and-apply-log` command.

**A.8: What would happen if I start a server directly using a raw directory backup, without running either the `copy-back` or the `apply-log` operation?**

*This should never be attempted.* Not only would the server crash, but the backup would likely get corrupted and become unusable. This is because the directory backup contains metadata created by `mysqlbackup` that the MySQL server would not understand; also, the raw backup might be inconsistent and need to be brought up-to-date by an `apply-log operation`, so that changes made to the database during the backup process can be applied.



---

## Appendix B MySQL Enterprise Backup Limitations

### Table of Contents

|                                                  |     |
|--------------------------------------------------|-----|
| B.1 Limitations of MySQL Enterprise Backup ..... | 149 |
|--------------------------------------------------|-----|

Please refer to the [MySQL Enterprise Backup 4.0 Release Notes](#) for a list of fixed bugs for `mysqlbackup`.

### B.1 Limitations of MySQL Enterprise Backup

- In Linux, Unix, and OS X systems, `mysqlbackup` does not record file ownership or permissions of the files that are backed up. Upon restore, these files might have different ownership (for example, being owned by `root` now rather than `mysql`). They might also have different read/write permissions (for example, being readable by anyone rather than just the file owner). When planning your backup strategy, survey the files in the MySQL data directory to ensure they have consistent owner and permission settings. When executing a restore operation, use an appropriate combination of `su`, `umask`, `chown`, and `chmod` on the restored files to set up the same owners and privileges as on the original files.
- In some cases, backups of non-transactional tables such as `MyISAM` tables could contain additional uncommitted data. If `autocommit` is turned off, and both `InnoDB` tables and non-transactional tables are modified within the same transaction, data can be written to the non-transactional table before the binary log position is updated. The binary log position is updated when the transaction is committed, but the non-transactional data is written immediately. If the backup occurs while such a transaction is open, the backup data contains the updates made to the non-transactional table.
- If the `mysqlbackup` process is interrupted by, for example, a Unix `kill -9` command, a `FLUSH TABLES WITH READ LOCK` operation might remain running. In this case, use the `KILL QUERY` statement from the `mysql` command line to kill the `FLUSH TABLES WITH READ LOCK` statement. This issue is more likely to occur if the `FLUSH TABLES` operation is stalled by a long-running query or transaction. Refer to [Section 7.1, “Optimizing Backup Performance”](#) for guidelines about backup timing and performance.
- Do not run the DDL operations `ALTER TABLE`, `TRUNCATE TABLE`, `OPTIMIZE TABLE`, `REPAIR TABLE`, `RESTORE TABLE` or `CREATE INDEX` while a backup operation is going on. The resulting backup might become corrupted.

The only `ALTER TABLE` operations that can be safely run in parallel with a backup are those that do not influence the physical representation of records on disk, such as changing column names or default column values.

- When statement-based `binary log format` is used on the MySQL server (which is the default behavior), if you take a backup when there are temporary tables in the database and they have been used to update or insert into normal tables during the backup, application of the MySQL binary log to a backup could then fail—that is, you might not be able to roll forward the backup to a particular point in time using the MySQL binary log. This is because temporary tables are not copied to the backup, as the physical filenames `#sql*.frm` do not correspond to the logical table names that MySQL writes to the binary log. To avoid the problem, use row-based or mixed format for the binary log by setting the value for the `--binlog-format` option to “ROW” or “MIXED” on the server.
- The `engines` column in the `mysql.backup_history` table does not correctly reflect the storage engines of the backed-up databases.
- Hot backups for large databases with heavy writing workloads (say, in the order of gigabytes per minute) can take a very long time to complete, due to the huge redo log files that are generated on the server

when the backup is in progress. And if the redo log files grow faster than they can be processed by `mysqlbackup`, the backup operation can actually fail when `mysqlbackup` cannot catch up with the redo log cycles and LSNs get overwritten by the server before they are read by `mysqlbackup`. The problem is intensified when the I/O resources available for reading and writing the redo logs are scarce during the backup process. However, if only a small number of tables of the database are modified frequently, the Optimistic Backup feature might alleviate the problem. See [Section 4.3.6, “Making an Optimistic Backup”](#) for details.

- Compressed InnoDB tables from MySQL server 5.6.10 and earlier cannot be restored with MySQL Enterprise Backup 3.9.0 or later, due to a known issue with the InnoDB storage engine (see Bug# 72851 on the MySQL Bug System).
- While it is possible to backup to or restore from a Network Attached Storage (NAS) device using MySQL Enterprise Backup, due to networking issues that might arise, the consistency of the backups and the performance of the backup or restore operations might be compromised.
- When creating a backup using [transportable tablespace \(TTS\)](#) for a server containing tables with a mix of the Antelope and Barracuda file formats, do not apply full locking on the tables (that is, do not specify `--use-tts=with-full-locking`). Instead, just specify `--use-tts` or `--use-tts=with-minimum-locking`, both of which will apply minimum locking to the tables.
- Backup of a partitioned table using [transportable tablespace \(TTS\)](#) would fail when any (or all) of its partitions was created in a shared tablespace.
- A `copy-back` process of a partitioned table backed up using [transportable tablespace \(TTS\)](#) would fail if any of the partitions was created outside of the backed-up server's data directory, even when the `--force [105]` option is used.
- If Data Definition Language (DDL) statements are executed on the server when a backup is being created using [transportable tablespace \(TTS\)](#), the backup might fail. It is because tables not being backed up are not locked during the backup process, but `mysqlbackup` still checks the status of those tables at the end of the process, and an error might occur if definitions of those tables have been changed. To avoid the problem, do not perform any DDL operation, especially `DROP TABLE`, when a TTS backup is going on.
- If a table containing full-text search (FTS) index is backed up using [transportable tablespace \(TTS\)](#), after it is restored, the FTS index will be corrupted. Users will need to recreate the index with the following command:

```
mysql> ALTER TABLE mytable ENGINE = INNODB;
```

Then, check that there are no more errors with the table:

```
mysql> CHECK TABLE mytable;
```

- Tables created on the MySQL server with the [ANSI\\_QUOTES](#) SQL mode cannot be backed up using [transportable tablespace \(TTS\)](#).
- Currently, MySQL Enterprise Backup cannot backup a multi-source replication slave.
- MySQL Enterprise Backup does not include the `.pem` files from the server into the backup. The files are part of the server instance when SSL connections are enabled.
- When using the `backup-and-apply-log` command without a connection to the MySQL server, `mysqlbackup` cannot know the correct binary log file name and binary log position for the backup; yet,

at the end of the `backup-and-apply-log` operation, `mysqlbackup` still prints out some values for the binary log file name and position, which are random in nature and should be ignored by users.

- When backing up an instance of MySQL 5.7, if a `CREATE INDEX` statement with `ALGORITHM = INPLACE` is issued when the backup process is going on, because the statement will not go into the redo log of the MySQL 5.7 server (see [Sorted Index Builds](#) for details), it cannot be recorded in the backup, and the index will not be recreated by `mysqlbackup` when the backup is restored.



---

# Appendix C Compatibility Information for MySQL Enterprise Backup

## Table of Contents

|                                                            |     |
|------------------------------------------------------------|-----|
| C.1 Cross-Platform Compatibility .....                     | 153 |
| C.2 Compatibility with MySQL Versions .....                | 153 |
| C.3 Compatibility with Older MySQL Enterprise Backup ..... | 153 |

This section describes information related to compatibility issues for MySQL Enterprise Backup releases.

### C.1 Cross-Platform Compatibility

MySQL Enterprise Backup is cross-platform compatible when running on the Linux and Windows operating systems: backups on a Linux machine can be restored on a Windows machine, and vice versa. However, to avoid data transfer problems arising from letter cases of database or table names, the variable `lower_case_table_names` must be properly configured on the MySQL servers. For details, see [Identifier Case Sensitivity](#).

### C.2 Compatibility with MySQL Versions

Currently, MySQL Enterprise Backup 4.0 only supports MySQL 5.7. For MySQL 5.6 and earlier, use MySQL Enterprise Backup 3.12.

### C.3 Compatibility with Older MySQL Enterprise Backup

Currently, MySQL Enterprise Backup 4.0 is incompatible with earlier MySQL Enterprise Backup versions.



---

## Appendix D Licenses for Third-Party Components

### Table of Contents

|                                                                                                                                |     |
|--------------------------------------------------------------------------------------------------------------------------------|-----|
| D.1 cURL (libcurl) License .....                                                                                               | 155 |
| D.2 GNU General Public License Version 3.0, 29 June 2007 and GCC Runtime Library Exception<br>Version 3.1, 31 March 2009 ..... | 155 |
| D.3 GNU Standard C++ Library (libstdc++) License .....                                                                         | 167 |
| D.4 Google Controlling Master Thread I/O Rate Patch License .....                                                              | 168 |
| D.5 Google SMP Patch License .....                                                                                             | 169 |
| D.6 LZ4 License .....                                                                                                          | 170 |
| D.7 OpenSSL v1.0 License .....                                                                                                 | 170 |
| D.8 Percona Multiple I/O Threads Patch License .....                                                                           | 172 |
| D.9 RE2 License .....                                                                                                          | 172 |
| D.10 RegEX-Spencer Library License .....                                                                                       | 173 |
| D.11 RFC 3174 - US Secure Hash Algorithm 1 (SHA1) License .....                                                                | 173 |
| D.12 zlib License .....                                                                                                        | 174 |

### D.1 cURL (libcurl) License

The following software may be included in this product:

```
cURL (libcurl)

Use of any of this software is governed by the terms of the license below:

COPYRIGHT AND PERMISSION NOTICE

Copyright (c) 1996 - 2014, Daniel Stenberg, <daniel@haxx.se>.
All rights reserved.

Permission to use, copy, modify, and distribute this software for any purpose
with or without fee is hereby granted, provided that the above copyright
notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR
IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY,
FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY
RIGHTS. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR
ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT
OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR
THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Except as contained in this notice, the name of a copyright holder shall not
be used in advertising or otherwise to promote the sale, use or other
dealings in this Software without prior written authorization of the copyright
holder.
```

### D.2 GNU General Public License Version 3.0, 29 June 2007 and GCC Runtime Library Exception Version 3.1, 31 March 2009

```
GNU GENERAL PUBLIC LICENSE
Version 3, 29 June 2007

Copyright (C) 2007 Free Software Foundation, Inc. <http://fsf.org/>
Everyone is permitted to copy and distribute verbatim copies of
```

this license document, but changing it is not allowed.

#### Preamble

The GNU General Public License is a free, copyleft license for software and other kinds of works.

The licenses for most software and other practical works are designed to take away your freedom to share and change the works. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change all versions of a program--to make sure it remains free software for all its users. We, the Free Software Foundation, use the GNU General Public License for most of our software; it applies also to any other work released this way by its authors. You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for them if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs, and that you know you can do these things.

To protect your rights, we need to prevent others from denying you these rights or asking you to surrender the rights. Therefore, you have certain responsibilities if you distribute copies of the software, or if you modify it: responsibilities to respect the freedom of others.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must pass on to the recipients the same freedoms that you received. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

Developers that use the GNU GPL protect your rights with two steps: (1) assert copyright on the software, and (2) offer you this License giving you legal permission to copy, distribute and/or modify it.

For the developers' and authors' protection, the GPL clearly explains that there is no warranty for this free software. For both users' and authors' sake, the GPL requires that modified versions be marked as changed, so that their problems will not be attributed erroneously to authors of previous versions.

Some devices are designed to deny users access to install or run modified versions of the software inside them, although the manufacturer can do so. This is fundamentally incompatible with the aim of protecting users' freedom to change the software. The systematic pattern of such abuse occurs in the area of products for individuals to use, which is precisely where it is most unacceptable. Therefore, we have designed this version of the GPL to prohibit the practice for those products. If such problems arise substantially in other domains, we stand ready to extend this provision to those domains in future versions of the GPL, as needed to protect the freedom of users.

Finally, every program is threatened constantly by software patents. States should not allow patents to restrict development and use of software on general-purpose computers, but in those that do, we wish to avoid the special danger that patents applied to a free program could make it effectively proprietary. To prevent this, the GPL assures that patents cannot be used to render the program non-free.



The precise terms and conditions for copying, distribution and modification follow.

#### TERMS AND CONDITIONS

##### 0. Definitions.

"This License" refers to version 3 of the GNU General Public License.

"Copyright" also means copyright-like laws that apply to other kinds of works, such as semiconductor masks.

"The Program" refers to any copyrightable work licensed under this License. Each licensee is addressed as "you". "Licensees" and "recipients" may be individuals or organizations.

To "modify" a work means to copy from or adapt all or part of the work in a fashion requiring copyright permission, other than the making of an exact copy. The resulting work is called a "modified version" of the earlier work or a work "based on" the earlier work.

A "covered work" means either the unmodified Program or a work based on the Program.

To "propagate" a work means to do anything with it that, without permission, would make you directly or secondarily liable for infringement under applicable copyright law, except executing it on a computer or modifying a private copy. Propagation includes copying, distribution (with or without modification), making available to the public, and in some countries other activities as well.

To "convey" a work means any kind of propagation that enables other parties to make or receive copies. Mere interaction with a user through a computer network, with no transfer of a copy, is not conveying.

An interactive user interface displays "Appropriate Legal Notices" to the extent that it includes a convenient and prominently visible feature that (1) displays an appropriate copyright notice, and (2) tells the user that there is no warranty for the work (except to the extent that warranties are provided), that licensees may convey the work under this License, and how to view a copy of this License. If the interface presents a list of user commands or options, such as a menu, a prominent item in the list meets this criterion.

##### 1. Source Code.

The "source code" for a work means the preferred form of the work for making modifications to it. "Object code" means any non-source form of a work.

A "Standard Interface" means an interface that either is an official standard defined by a recognized standards body, or, in the case of interfaces specified for a particular programming language, one that is widely used among developers working in that language.

The "System Libraries" of an executable work include anything, other than the work as a whole, that (a) is included in the normal form of packaging a Major Component, but which is not part of that Major Component, and (b) serves only to enable use of the work with that Major Component, or to implement a Standard Interface for which an implementation is available to the public in source code form. A "Major Component", in this context, means a major essential component (kernel, window system, and so on) of the specific operating system (if any) on which the executable work runs, or a compiler used to produce the work, or an object code interpreter used to run it.

The "Corresponding Source" for a work in object code form means all the source code needed to generate, install, and (for an executable work) run the object code and to modify the work, including scripts to control those activities. However, it does not include the work's System Libraries, or general-purpose tools or generally available free programs which are used unmodified in performing those activities but which are not part of the work. For example, Corresponding Source includes interface definition files associated with source files for the work, and the source code for shared libraries and dynamically linked subprograms that the work is specifically designed to require, such as by intimate data communication or control flow between those subprograms and other parts of the work.

The Corresponding Source need not include anything that users can regenerate automatically from other parts of the Corresponding Source.

The Corresponding Source for a work in source code form is that same work.

## 2. Basic Permissions.

All rights granted under this License are granted for the term of copyright on the Program, and are irrevocable provided the stated conditions are met. This License explicitly affirms your unlimited permission to run the unmodified Program. The output from running a covered work is covered by this License only if the output, given its content, constitutes a covered work. This License acknowledges your rights of fair use or other equivalent, as provided by copyright law.

You may make, run and propagate covered works that you do not convey, without conditions so long as your license otherwise remains in force. You may convey covered works to others for the sole purpose of having them make modifications exclusively for you, or provide you with facilities for running those works, provided that you comply with the terms of this License in conveying all material for which you do not control copyright. Those thus making or running the covered works for you must do so exclusively on your behalf, under your direction and control, on terms that prohibit them from making any copies of your copyrighted material outside their relationship with you.

Conveying under any other circumstances is permitted solely under the conditions stated below. Sublicensing is not allowed; section 10 makes it unnecessary.

## 3. Protecting Users' Legal Rights From Anti-Circumvention Law.

No covered work shall be deemed part of an effective technological measure under any applicable law fulfilling obligations under article 11 of the WIPO copyright treaty adopted on 20 December 1996, or similar laws prohibiting or restricting circumvention of such measures.

When you convey a covered work, you waive any legal power to forbid circumvention of technological measures to the extent such circumvention is effected by exercising rights under this License with respect to the covered work, and you disclaim any intention to limit operation or modification of the work as a means of enforcing, against the work's users, your or third parties' legal rights to forbid circumvention of technological measures.

## 4. Conveying Verbatim Copies.

You may convey verbatim copies of the Program's source code as you

receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice; keep intact all notices stating that this License and any non-permissive terms added in accord with section 7 apply to the code; keep intact all notices of the absence of any warranty; and give all recipients a copy of this License along with the Program.

You may charge any price or no price for each copy that you convey, and you may offer support or warranty protection for a fee.

#### 5. Conveying Modified Source Versions.

You may convey a work based on the Program, or the modifications to produce it from the Program, in the form of source code under the terms of section 4, provided that you also meet all of these conditions:

- a) The work must carry prominent notices stating that you modified it, and giving a relevant date.
- b) The work must carry prominent notices stating that it is released under this License and any conditions added under section 7. This requirement modifies the requirement in section 4 to "keep intact all notices".
- c) You must license the entire work, as a whole, under this License to anyone who comes into possession of a copy. This License will therefore apply, along with any applicable section 7 additional terms, to the whole of the work, and all its parts, regardless of how they are packaged. This License gives no permission to license the work in any other way, but it does not invalidate such permission if you have separately received it.
- d) If the work has interactive user interfaces, each must display Appropriate Legal Notices; however, if the Program has interactive interfaces that do not display Appropriate Legal Notices, your work need not make them do so.

A compilation of a covered work with other separate and independent works, which are not by their nature extensions of the covered work, and which are not combined with it such as to form a larger program, in or on a volume of a storage or distribution medium, is called an "aggregate" if the compilation and its resulting copyright are not used to limit the access or legal rights of the compilation's users beyond what the individual works permit. Inclusion of a covered work in an aggregate does not cause this License to apply to the other parts of the aggregate.

#### 6. Conveying Non-Source Forms.

You may convey a covered work in object code form under the terms of sections 4 and 5, provided that you also convey the machine-readable Corresponding Source under the terms of this License, in one of these ways:

- a) Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by the Corresponding Source fixed on a durable physical medium customarily used for software interchange.
- b) Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by a written offer, valid for at least three years and valid for as long as you offer spare parts or customer support for that product model, to give anyone who possesses the object code either (1) a copy of the Corresponding Source for all the software in the product that is covered by this License, on a

lasting physical medium customarily used for software interchange, for a price no more than your reasonable cost of physically performing this conveying of source, or (2) access to copy the Corresponding Source from a network server at no charge.

c) Convey individual copies of the object code with a copy of the written offer to provide the Corresponding Source. This alternative is allowed only occasionally and noncommercially, and only if you received the object code with such an offer, in accord with subsection 6b.

d) Convey the object code by offering access from a designated place (gratis or for a charge), and offer equivalent access to the Corresponding Source in the same way through the same place at no further charge. You need not require recipients to copy the Corresponding Source along with the object code. If the place to copy the object code is a network server, the Corresponding Source may be on a different server (operated by you or a third party) that supports equivalent copying facilities, provided you maintain clear directions next to the object code saying where to find the Corresponding Source. Regardless of what server hosts the Corresponding Source, you remain obligated to ensure that it is available for as long as needed to satisfy these requirements.

e) Convey the object code using peer-to-peer transmission, provided you inform other peers where the object code and Corresponding Source of the work are being offered to the general public at no charge under subsection 6d.

A separable portion of the object code, whose source code is excluded from the Corresponding Source as a System Library, need not be included in conveying the object code work.

A "User Product" is either (1) a "consumer product", which means any tangible personal property which is normally used for personal, family, or household purposes, or (2) anything designed or sold for incorporation into a dwelling. In determining whether a product is a consumer product, doubtful cases shall be resolved in favor of coverage. For a particular product received by a particular user, "normally used" refers to a typical or common use of that class of product, regardless of the status of the particular user or of the way in which the particular user actually uses, or expects or is expected to use, the product. A product is a consumer product regardless of whether the product has substantial commercial, industrial or non-consumer uses, unless such uses represent the only significant mode of use of the product.

"Installation Information" for a User Product means any methods, procedures, authorization keys, or other information required to install and execute modified versions of a covered work in that User Product from a modified version of its Corresponding Source. The information must suffice to ensure that the continued functioning of the modified object code is in no case prevented or interfered with solely because modification has been made.

If you convey an object code work under this section in, or with, or specifically for use in, a User Product, and the conveying occurs as part of a transaction in which the right of possession and use of the User Product is transferred to the recipient in perpetuity or for a fixed term (regardless of how the transaction is characterized), the Corresponding Source conveyed under this section must be accompanied by the Installation Information. But this requirement does not apply if neither you nor any third party retains the ability to install modified object code on the User Product (for example, the work has been installed in ROM).

The requirement to provide Installation Information does not include a requirement to continue to provide support service, warranty, or updates for a work that has been modified or installed by the recipient, or for the User Product in which it has been modified or installed. Access to a network may be denied when the modification itself materially and adversely affects the operation of the network or violates the rules and protocols for communication across the network.

Corresponding Source conveyed, and Installation Information provided, in accord with this section must be in a format that is publicly documented (and with an implementation available to the public in source code form), and must require no special password or key for unpacking, reading or copying.

#### 7. Additional Terms.

"Additional permissions" are terms that supplement the terms of this License by making exceptions from one or more of its conditions. Additional permissions that are applicable to the entire Program shall be treated as though they were included in this License, to the extent that they are valid under applicable law. If additional permissions apply only to part of the Program, that part may be used separately under those permissions, but the entire Program remains governed by this License without regard to the additional permissions.

When you convey a copy of a covered work, you may at your option remove any additional permissions from that copy, or from any part of it. (Additional permissions may be written to require their own removal in certain cases when you modify the work.) You may place additional permissions on material, added by you to a covered work, for which you have or can give appropriate copyright permission.

Notwithstanding any other provision of this License, for material you add to a covered work, you may (if authorized by the copyright holders of that material) supplement the terms of this License with terms:

- a) Disclaiming warranty or limiting liability differently from the terms of sections 15 and 16 of this License; or
- b) Requiring preservation of specified reasonable legal notices or author attributions in that material or in the Appropriate Legal Notices displayed by works containing it; or
- c) Prohibiting misrepresentation of the origin of that material, or requiring that modified versions of such material be marked in reasonable ways as different from the original version; or
- d) Limiting the use for publicity purposes of names of licensors or authors of the material; or
- e) Declining to grant rights under trademark law for use of some trade names, trademarks, or service marks; or
- f) Requiring indemnification of licensors and authors of that material by anyone who conveys the material (or modified versions of it) with contractual assumptions of liability to the recipient, for any liability that these contractual assumptions directly impose on those licensors and authors.

All other non-permissive additional terms are considered "further restrictions" within the meaning of section 10. If the Program as you received it, or any part of it, contains a notice stating that it is governed by this License along with a term that is a further

restriction, you may remove that term. If a license document contains a further restriction but permits relicensing or conveying under this License, you may add to a covered work material governed by the terms of that license document, provided that the further restriction does not survive such relicensing or conveying.

If you add terms to a covered work in accord with this section, you must place, in the relevant source files, a statement of the additional terms that apply to those files, or a notice indicating where to find the applicable terms.

Additional terms, permissive or non-permissive, may be stated in the form of a separately written license, or stated as exceptions; the above requirements apply either way.

#### 8. Termination.

You may not propagate or modify a covered work except as expressly provided under this License. Any attempt otherwise to propagate or modify it is void, and will automatically terminate your rights under this License (including any patent licenses granted under the third paragraph of section 11).

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, you do not qualify to receive new licenses for the same material under section 10.

#### 9. Acceptance Not Required for Having Copies.

You are not required to accept this License in order to receive or run a copy of the Program. Ancillary propagation of a covered work occurring solely as a consequence of using peer-to-peer transmission to receive a copy likewise does not require acceptance. However, nothing other than this License grants you permission to propagate or modify any covered work. These actions infringe copyright if you do not accept this License. Therefore, by modifying or propagating a covered work, you indicate your acceptance of this License to do so.

#### 10. Automatic Licensing of Downstream Recipients.

Each time you convey a covered work, the recipient automatically receives a license from the original licensors, to run, modify and propagate that work, subject to this License. You are not responsible for enforcing compliance by third parties with this License.

An "entity transaction" is a transaction transferring control of an organization, or substantially all assets of one, or subdividing an organization, or merging organizations. If propagation of a covered work results from an entity transaction, each party to that transaction who receives a copy of the work also receives whatever licenses to the work the party's predecessor in interest had or could

give under the previous paragraph, plus a right to possession of the Corresponding Source of the work from the predecessor in interest, if the predecessor has it or can get it with reasonable efforts.

You may not impose any further restrictions on the exercise of the rights granted or affirmed under this License. For example, you may not impose a license fee, royalty, or other charge for exercise of rights granted under this License, and you may not initiate litigation (including a cross-claim or counterclaim in a lawsuit) alleging that any patent claim is infringed by making, using, selling, offering for sale, or importing the Program or any portion of it.

#### 11. Patents.

A "contributor" is a copyright holder who authorizes use under this License of the Program or a work on which the Program is based. The work thus licensed is called the contributor's "contributor version".

A contributor's "essential patent claims" are all patent claims owned or controlled by the contributor, whether already acquired or hereafter acquired, that would be infringed by some manner, permitted by this License, of making, using, or selling its contributor version, but do not include claims that would be infringed only as a consequence of further modification of the contributor version. For purposes of this definition, "control" includes the right to grant patent sublicenses in a manner consistent with the requirements of this License.

Each contributor grants you a non-exclusive, worldwide, royalty-free patent license under the contributor's essential patent claims, to make, use, sell, offer for sale, import and otherwise run, modify and propagate the contents of its contributor version.

In the following three paragraphs, a "patent license" is any express agreement or commitment, however denominated, not to enforce a patent (such as an express permission to practice a patent or covenant not to sue for patent infringement). To "grant" such a patent license to a party means to make such an agreement or commitment not to enforce a patent against the party.

If you convey a covered work, knowingly relying on a patent license, and the Corresponding Source of the work is not available for anyone to copy, free of charge and under the terms of this License, through a publicly available network server or other readily accessible means, then you must either (1) cause the Corresponding Source to be so available, or (2) arrange to deprive yourself of the benefit of the patent license for this particular work, or (3) arrange, in a manner consistent with the requirements of this License, to extend the patent license to downstream recipients. "Knowingly relying" means you have actual knowledge that, but for the patent license, your conveying the covered work in a country, or your recipient's use of the covered work in a country, would infringe one or more identifiable patents in that country that you have reason to believe are valid.

If, pursuant to or in connection with a single transaction or arrangement, you convey, or propagate by procuring conveyance of, a covered work, and grant a patent license to some of the parties receiving the covered work authorizing them to use, propagate, modify or convey a specific copy of the covered work, then the patent license you grant is automatically extended to all recipients of the covered work and works based on it.

A patent license is "discriminatory" if it does not include within the scope of its coverage, prohibits the exercise of, or is conditioned on the non-exercise of one or more of the rights that are specifically

granted under this License. You may not convey a covered work if you are a party to an arrangement with a third party that is in the business of distributing software, under which you make payment to the third party based on the extent of your activity of conveying the work, and under which the third party grants, to any of the parties who would receive the covered work from you, a discriminatory patent license (a) in connection with copies of the covered work conveyed by you (or copies made from those copies), or (b) primarily for and in connection with specific products or compilations that contain the covered work, unless you entered into that arrangement, or that patent license was granted, prior to 28 March 2007.

Nothing in this License shall be construed as excluding or limiting any implied license or other defenses to infringement that may otherwise be available to you under applicable patent law.

#### 12. No Surrender of Others' Freedom.

If conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot convey a covered work so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not convey it at all. For example, if you agree to terms that obligate you to collect a royalty for further conveying from those to whom you convey the Program, the only way you could satisfy both those terms and this License would be to refrain entirely from conveying the Program.

#### 13. Use with the GNU Affero General Public License.

Notwithstanding any other provision of this License, you have permission to link or combine any covered work with a work licensed under version 3 of the GNU Affero General Public License into a single combined work, and to convey the resulting work. The terms of this License will continue to apply to the part which is the covered work, but the special requirements of the GNU Affero General Public License, section 13, concerning interaction through a network will apply to the combination as such.

#### 14. Revised Versions of this License.

The Free Software Foundation may publish revised and/or new versions of the GNU General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies that a certain numbered version of the GNU General Public License "or any later version" applies to it, you have the option of following the terms and conditions either of that numbered version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of the GNU General Public License, you may choose any version ever published by the Free Software Foundation.

If the Program specifies that a proxy can decide which future versions of the GNU General Public License can be used, that proxy's public statement of acceptance of a version permanently authorizes you to choose that version for the Program.

Later license versions may give you additional or different permissions. However, no additional obligations are imposed on any author or copyright holder as a result of your choosing to follow a later version.

#### 15. Disclaimer of Warranty.



THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

#### 16. Limitation of Liability.

IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MODIFIES AND/OR CONVEYS THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

#### 17. Interpretation of Sections 15 and 16.

If the disclaimer of warranty and limitation of liability provided above cannot be given local legal effect according to their terms, reviewing courts shall apply local law that most closely approximates an absolute waiver of all civil liability in connection with the Program, unless a warranty or assumption of liability accompanies a copy of the Program in return for a fee.

#### END OF TERMS AND CONDITIONS

#### How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively state the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

```
<one line to give the program's name and a brief idea of what it does.>
Copyright (C) <year> <name of author>
```

This program is free software: you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation, either version 3 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program. If not, see  
<<http://www.gnu.org/licenses/>>.

Also add information on how to contact you by electronic and paper mail.

If the program does terminal interaction, make it output a short

notice like this when it starts in an interactive mode:

```
<program> Copyright (C) <year> <name of author>
This program comes with ABSOLUTELY NO WARRANTY; for details type 'show w'.
This is free software, and you are welcome to redistribute it
under certain conditions; type 'show c' for details.
```

The hypothetical commands 'show w' and 'show c' should show the appropriate parts of the General Public License. Of course, your program's commands might be different; for a GUI interface, you would use an "about box".

You should also get your employer (if you work as a programmer) or school, if any, to sign a "copyright disclaimer" for the program, if necessary. For more information on this, and how to apply and follow the GNU GPL, see <http://www.gnu.org/licenses/>.

The GNU General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Lesser General Public License instead of this License. But first, please read <http://www.gnu.org/philosophy/why-not-lgpl.html>.

==

==

GCC RUNTIME LIBRARY EXCEPTION

Version 3.1, 31 March 2009

Copyright © 2009 Free Software Foundation, Inc. <http://fsf.org/>

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

This GCC Runtime Library Exception ("Exception") is an additional permission under section 7 of the GNU General Public License, version 3 ("GPLv3"). It applies to a given file (the "Runtime Library") that bears a notice placed by the copyright holder of the file stating that the file is governed by GPLv3 along with this Exception.

When you use GCC to compile a program, GCC may combine portions of certain GCC header files and runtime libraries with the compiled program. The purpose of this Exception is to allow compilation of non-GPL (including proprietary) programs to use, in this way, the header files and runtime libraries covered by this Exception.

#### 0. Definitions.

A file is an "Independent Module" if it either requires the Runtime Library for execution after a Compilation Process, or makes use of an interface provided by the Runtime Library, but is not otherwise based on the Runtime Library.

"GCC" means a version of the GNU Compiler Collection, with or without modifications, governed by version 3 (or a specified later version) of the GNU General Public License (GPL) with the option of using any subsequent versions published by the FSF.

"GPL-compatible Software" is software whose conditions of propagation, modification and use would permit combination with GCC in accord with the license of GCC.

"Target Code" refers to output from any compiler for a real or virtual target processor architecture, in executable form or suitable for input to an assembler, loader, linker and/or execution phase. Notwithstanding that, Target Code does not include data in any

format that is used as a compiler intermediate representation, or used for producing a compiler intermediate representation.

The "Compilation Process" transforms code entirely represented in non-intermediate languages designed for human-written code, and/or in Java Virtual Machine byte code, into Target Code. Thus, for example, use of source code generators and preprocessors need not be considered part of the Compilation Process, since the Compilation Process can be understood as starting with the output of the generators or preprocessors.

A Compilation Process is "Eligible" if it is done using GCC, alone or with other GPL-compatible software, or if it is done without using any work based on GCC. For example, using non-GPL-compatible Software to optimize any GCC intermediate representations would not qualify as an Eligible Compilation Process.

#### 1. Grant of Additional Permission.

You have permission to propagate a work of Target Code formed by combining the Runtime Library with Independent Modules, even if such propagation would otherwise violate the terms of GPLv3, provided that all Target Code was generated by Eligible Compilation Processes. You may then convey such a combination under terms of your choice, consistent with the licensing of the Independent Modules.

#### 2. No Weakening of GCC Copyleft.

The availability of this Exception does not imply any general presumption that third-party software is unaffected by the copyleft requirements of the license of GCC.

==

## D.3 GNU Standard C++ Library (libstdc++) License

The following software may be included in this product: GNU Standard C++ Library (libstdc++)

This component is licensed under [Section D.2, "GNU General Public License Version 3.0, 29 June 2007 and GCC Runtime Library Exception Version 3.1, 31 March 2009"](#).

Additional notices:

==

Copyright (c) 1994  
Hewlett-Packard Company

Permission to use, copy, modify, distribute and sell this software and its documentation for any purpose is hereby granted without fee, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation. Hewlett-Packard Company makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

==

==

Copyright (c) 1996,1997  
Silicon Graphics Computer Systems, Inc.

Permission to use, copy, modify, distribute and sell this software and its documentation for any purpose is hereby granted without fee, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear

```
in supporting documentation. Silicon Graphics makes no
representations about the suitability of this software for any
purpose. It is provided "as is" without express or implied
warranty.
==
==
shared_count.hpp
@ Copyright (c) 2001, 2002, 2003 Peter Dimov and Multi Media Ltd.

shared_ptr.hpp
Copyright (C) 1998, 1999 Greg Colvin and Beman Dawes.
Copyright (C) 2001, 2002, 2003 Peter Dimov

weak_ptr.hpp
Copyright (C) 2001, 2002, 2003 Peter Dimov

enable_shared_from_this.hpp
Copyright (C) 2002 Peter Dimov

Distributed under the Boost Software License, Version 1.0.

Boost Software License - Version 1.0 - August 17th, 2003

Permission is hereby granted, free of charge, to any person or
organization obtaining a copy of the software and accompanying
documentation covered by this license (the "Software") to use,
reproduce, display, distribute, execute, and transmit the Software,
and to prepare derivative works of the Software, and to permit
third-parties to whom the Software is furnished to do so, all subject
to the following:

The copyright notices in the Software and this entire statement,
including the above license grant, this restriction and the following
disclaimer, must be included in all copies of the Software, in whole
or in part, and all derivative works of the Software, unless such
copies or derivative works are solely in the form of machine-executable
object code generated by a source language processor.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND,
EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF
MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND
NON-INFRINGEMENT. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR ANYONE
DISTRIBUTING THE SOFTWARE BE LIABLE FOR ANY DAMAGES OR OTHER
LIABILITY, WHETHER IN CONTRACT, TORT OR OTHERWISE, ARISING FROM,
OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER
DEALINGS IN THE SOFTWARE.
==
==
Copyright (C) 2004 Ami Tavory and Vladimir Dreizin, IBM-HRL.

Permission to use, copy, modify, sell, and distribute this software
is hereby granted without fee, provided that the above copyright
notice appears in all copies, and that both that copyright notice
and this permission notice appear in supporting documentation. None
of the above authors, nor IBM Haifa Research Laboratories, make any
representation about the suitability of this software for any
purpose. It is provided "as is" without express or implied warranty.
==
```

## D.4 Google Controlling Master Thread I/O Rate Patch License

The following software may be included in this product:

Google Controlling master thread I/O rate patch

```
Copyright (c) 2009, Google Inc.  
All rights reserved.
```

```
Redistribution and use in source and binary forms, with or without  
modification, are permitted provided that the following conditions  
are met:
```

- \* Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- \* Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- \* Neither the name of the Google Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

```
THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS  
"AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT  
LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS  
FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE  
COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT,  
INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING,  
BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES;  
LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER  
CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT  
LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN  
ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE  
POSSIBILITY OF SUCH DAMAGE.
```

## D.5 Google SMP Patch License

The following software may be included in this product:

### Google SMP Patch

```
Google SMP patch
```

```
Copyright (c) 2008, Google Inc.  
All rights reserved.
```

```
Redistribution and use in source and binary forms, with or without  
modification, are permitted provided that the following conditions  
are met:
```

- \* Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- \* Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- \* Neither the name of the Google Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

```
THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS  
"AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT  
LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS  
FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE  
COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT,  
INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING,  
BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES;  
LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER  
CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT  
LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN  
ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE
```

POSSIBILITY OF SUCH DAMAGE.

## D.6 LZ4 License

The following software may be included in this product:

```
LZ4 - Fast LZ compression algorithm

Copyright (C) 2011-2013, Yann Collet.
BSD 2-Clause License (http://www.opensource.org/licenses/bsd-license.php)

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions are
met:

    * Redistributions of source code must retain the above copyright
    notice, this list of conditions and the following disclaimer.
    * Redistributions in binary form must reproduce the above
    copyright notice, this list of conditions and the following disclaimer
    in the documentation and/or other materials provided with the
    distribution.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS
"AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT
LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR
A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT
OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL,
SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT
LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE,
DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY
THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT
(INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE
OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

You can contact the author at :
- LZ4 source repository : http://code.google.com/p/lz4/
- LZ4 public forum : https://groups.google.com/forum/#!forum/lz4c
```

## D.7 OpenSSL v1.0 License

The following software may be included in this product:

OpenSSL v1.0

```
LICENSE ISSUES
=====

The OpenSSL toolkit stays under a dual license, i.e. both the conditions of
the OpenSSL License and the original SSLeay license apply to the toolkit. See
below for the actual license texts. Actually both licenses are BSD-style Open
Source licenses. In case of any license issues related to OpenSSL please
contact openssl-core@openssl.org.

OpenSSL License
-----
/ =====
Copyright (c) 1998-2008 The OpenSSL Project.
All rights reserved.
Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice,
this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice,
```

this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

3. All advertising materials mentioning features or use of this software must display the following acknowledgment: "This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (Link1 /)"

4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact openssl-core@openssl.org.

5. Products derived from this software may not be called "OpenSSL" nor may "OpenSSL" appear in their names without prior written permission of the OpenSSL Project.

6. Redistributions of any form whatsoever must retain the following acknowledgment: "This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (Link2 /)"

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT ``AS IS'' AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

=====

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com). This product includes software written by Tim Hudson (tjh@cryptsoft.com).

Original SSLeay License

-----

/ Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com)  
All rights reserved.

This package is an SSL implementation written by Eric Young (eay@cryptsoft.com). The implementation was written so as to conform with Netscapes SSL. This library is free for commercial and non-commercial use as long as the following conditions are aheared to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, lhash, DES, etc., code; not just the SSL code. The SSL documentation included with this distribution is covered by the same copyright terms except that the holder is Tim Hudson (tjh@cryptsoft.com). Copyright remains Eric Young's, and as such any Copyright notices in the code are not to be removed. If this package is used in a product, Eric Young should be given attribution as the author of the parts of the library used. This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: 1. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer. 2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. 3. All advertising materials mentioning features or use of this software must display the following acknowledgement: "This product includes cryptographic software written by Eric Young (eay@cryptsoft.com)" The word 'cryptographic' can be left out if the routines from the library being used are not cryptographic related :-). 4. If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include an acknowledgement: "This product includes software written by Tim Hudson (tjh@cryptsoft.com)"

THIS SOFTWARE IS PROVIDED BY ERIC YOUNG ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND

```
ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT
(INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF
THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE. The
license and distribution terms for any publically available version or
derivative of this code cannot be changed. i.e. this code cannot simply be
copied and put under another distribution license [including the GNU Public
License.]
```

## D.8 Percona Multiple I/O Threads Patch License

The following software may be included in this product:

Percona Multiple I/O threads patch

```
Copyright (c) 2008, 2009 Percona Inc
All rights reserved.
```

```
Redistribution and use of this software in source and binary forms,
with or without modification, are permitted provided that the
following conditions are met:
```

- \* Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- \* Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- \* Neither the name of Percona Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission of Percona Inc.

```
THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS
"AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT
LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS
FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE
COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT,
INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING,
BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES;
LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER
CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN
ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE
POSSIBILITY OF SUCH DAMAGE.
```

## D.9 RE2 License

The following software may be included in this product:

RE2

```
Copyright (c) 2009 The RE2 Authors. All rights reserved.
```

```
Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions are met:
```

- \* Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- \* Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- \* Neither the name of Google Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.



THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The RE2 Authors are:  
Google Inc.  
Samsung Electronics  
Stefano Rivera <stefano.rivera@gmail.com>

## D.10 RegEX-Spencer Library License

The following software may be included in this product:

Henry Spencer's Regular-Expression Library (RegEX-Spencer)

Copyright 1992, 1993, 1994, 1997 Henry Spencer. All rights reserved.  
This software is not subject to any license of the American Telephone and Telegraph Company or of the Regents of the University of California.

Permission is granted to anyone to use this software for any purpose on any computer system, and to alter it and redistribute it, subject to the following restrictions:

1. The author is not responsible for the consequences of use of this software, no matter how awful, even if they arise from flaws in it.
2. The origin of this software must not be misrepresented, either by explicit claim or by omission. Since few users ever read sources, credits must appear in the documentation.
3. Altered versions must be plainly marked as such, and must not be misrepresented as being the original software. Since few users ever read sources, credits must appear in the documentation.
4. This notice may not be removed or altered.

## D.11 RFC 3174 - US Secure Hash Algorithm 1 (SHA1) License

The following software may be included in this product:

RFC 3174 - US Secure Hash Algorithm 1 (SHA1)

RFC 3174 - US Secure Hash Algorithm 1 (SHA1)

Copyright (C) The Internet Society (2001). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this

document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

#### Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

## D.12 zlib License

The following software may be included in this product:

zlib

Oracle gratefully acknowledges the contributions of Jean-loup Gailly and Mark Adler in creating the zlib general purpose compression library which is used in this product.

```
zlib.h -- interface of the 'zlib' general purpose compression library
Copyright (C) 1995-2004 Jean-loup Gailly and Mark Adler
```

```
zlib.h -- interface of the 'zlib' general purpose compression library
version 1.2.3, July 18th, 2005
Copyright (C) 1995-2005 Jean-loup Gailly and Mark Adler
```

```
zlib.h -- interface of the 'zlib' general purpose compression library
version 1.2.5, April 19th, 2010
Copyright (C) 1995-2010 Jean-loup Gailly and Mark Adler
```

This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software. Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
3. This notice may not be removed or altered from any source distribution.

Jean-loup Gailly jloup@gzip.org  
Mark Adler madler@alumni.caltech.edu

---

# MySQL Enterprise Backup Glossary

These terms are commonly used in information about the MySQL Enterprise Backup product.

## A

### .ARM file

Metadata for ARCHIVE tables. Contrast with **.ARZ file**. Files with this extension are always included in backups produced by the `mysqlbackup` command of the **MySQL Enterprise Backup** product.  
See Also [.ARZ file](#).

### .ARZ file

Data for ARCHIVE tables. Contrast with **.ARM file**. Files with this extension are always included in backups produced by the `mysqlbackup` command of the **MySQL Enterprise Backup** product.  
See Also [.ARM file](#).

### Antelope

The code name for the original InnoDB **file format**. It supports the **redundant** and **compact** row formats, but not the newer **dynamic** and **compressed** row formats available in the **Barracuda** file format.

If your application could benefit from InnoDB table **compression**, or uses BLOBs or large text columns that could benefit from the dynamic row format, you might switch some tables to Barracuda format. You select the file format to use by setting the `innodb_file_format` option before creating the table.

See Also [Barracuda](#), [compression](#), [file format](#).

### apply

The operation that transforms a **raw backup** into a **prepared backup** by incorporating changes that occurred while the backup was running, using data from the **log**.

See Also [log](#), [prepared backup](#), [raw backup](#).

## B

### .bz file

When `mysqlbackup` performs a compressed backup for a server that has binary logging enabled, it transforms each binary log file and relay log file (for a **slave** server in a **replication** setting) to a `binary-or-relay-log-file-name.bz` file. The **.bz** files are uncompressed at the time of restore.

See Also [binary log](#), [.bz file](#), [compression](#), [compression level](#), [.ibz file](#), [relay log](#).

### backup

The process of copying some or all table data and metadata from a MySQL instance, for safekeeping. Can also refer to the set of copied files. This is a crucial task for DBAs. The reverse of this process is the **restore** operation.

With MySQL, **physical backups** are performed by the **MySQL Enterprise Backup** product, and **logical backups** are performed by the `mysqldump` command. These techniques have different characteristics in terms of size and representation of the backup data, and speed (especially speed of the restore operation).

Backups are further classified as **hot**, **warm**, or **cold** depending on how much they interfere with normal database operation. (Hot backups have the least interference, cold backups the most.)

See Also [cold backup](#), [hot backup](#), [logical backup](#), [mysqldump](#), [physical backup](#), [warm backup](#).

### backup directory

The directory under which the backup data and metadata are stored, permanently or temporarily. It is used in most kinds of backup and restore operations, including single-file backups and restores. See the description of the `--backup-dir` option on how the backup directory is used for different purposes and for different operations.

---

backup repository

Contrast with **server repository**.

See Also [repository](#), [server repository](#).

backup-my.cnf

A small **configuration file** generated by **MySQL Enterprise Backup**, containing a minimal set of configuration parameters. This file records the settings that apply to this backup data. Subsequent operations, such as the **apply** process, read options from this file to determine how the backup data is structured. This file always has the extension `.cnf`, rather than `.cnf` on Unix-like systems and `.ini` on Windows systems.

See Also [apply](#), [configuration file](#).

Barracuda

The code name for an InnoDB **file format** that supports compression for table data. This file format was first introduced in the InnoDB Plugin. It supports the **compressed** row format that enables InnoDB table compression, and the **dynamic** row format that improves the storage layout for BLOB and large text columns. You can select it through the `innodb_file_format` option.

Because the InnoDB **system tablespace** is stored in the original **Antelope** file format, to use the Barracuda file format you must also enable the **file-per-table** setting, which puts newly created tables in their own tablespaces separate from the system tablespace.

The **MySQL Enterprise Backup** product version 3.5 and above supports backing up tablespaces that use the Barracuda file format.

See Also [Antelope](#), [file format](#), [MySQL Enterprise Backup](#), [row format](#), [system tablespace](#).

binary log

A file containing a record of all statements that attempt to change table data. These statements can be replayed to bring slave servers up to date in a **replication** scenario, or to bring a database up to date after restoring table data from a backup. The binary logging feature can be turned on and off, although Oracle recommends always enabling it if you use replication or perform backups.

You can examine the contents of the binary log, or replay those statements during replication or recovery, by using the `mysqlbinlog` command. For full information about the binary log, see [The Binary Log](#). For MySQL configuration options related to the binary log, see [Binary Logging Options and Variables](#).

For the **MySQL Enterprise Backup** product, the file name of the binary log and the current position within the file are important details. To record this information for the master server when taking a backup in a replication context, you can specify the `--slave-info` option.

Prior to MySQL 5.0, a similar capability was available, known as the update log. In MySQL 5.0 and higher, the binary log replaces the update log.

The binary log, if enabled on the server, is backed up by default.

See Also [binlog](#), [relay log](#), [replication](#).

binlog

An informal name for the **binary log** file. For example, you might see this abbreviation used in e-mail messages or forum discussions.

See Also [binary log](#).

## C

cold backup

A **backup** taken while the database is shut down. For busy applications and web sites, this might not be practical, and you might prefer a **warm backup** or a **hot backup**.

---

See Also [backup](#), [connection](#), [hot backup](#), [warm backup](#).

#### compression

A technique that produces smaller **backup** files, with size reduction influenced by the **compression level** setting. Suitable for keeping multiple sets of non-critical backup files. (For recent backups of critical data, you might leave the data uncompressed, to allow fast restore speed in case of emergency.)

**MySQL Enterprise Backup** can apply compression to the contents of **InnoDB** tables during the backup process, turning the **.ibd** files into **.ibz** files.

Compression adds CPU overhead to the backup process, and requires additional time and disk space during the **restore** process.

See Also [backup](#), [compression level](#), [.ibd file](#), [.ibz file](#), [InnoDB](#), [restore](#).

#### compression level

A setting that determines how much **compression** to apply to a compressed backup. This setting ranges from 0 (none), 1 (default level when compression is enabled) to 9 (maximum). The amount of compression for a given compression level depends on the nature of your data values. Higher compression levels do impose additional CPU overhead, so ideally you use the lowest value that produces a good balance of compression with low CPU overhead.

See Also [compression](#).

#### configuration file

The file that holds the startup options of the MySQL server and related products and components. Often referred to by its default file name, **my.cnf** on Linux, Unix, and OS X systems, and **my.ini** on Windows systems. The **MySQL Enterprise Backup** stores its default configuration settings in this file, under a [\[mysqlbackup\]](#) section. For convenience, MySQL Enterprise Backup can also read settings from the [\[client\]](#) section, for configuration options that are common between MySQL Enterprise Backup and other programs that connect to the MySQL server.

See Also [my.cnf](#), [my.ini](#).

#### connection

The mechanism used by certain backup operations to communicate with a running MySQL **server**. For example, the [mysqlbackup](#) command can log into the server being backed up to insert and update data in the **progress table** and the **history table**. A **hot backup** typically uses a database connection for convenience, but can proceed anyway if the connection is not available. A **warm backup** always uses a database connection, because it must put the server into a read-only state. A **cold backup** is taken while the MySQL server is shut down, and so cannot use any features that require a connection.

See Also [cold backup](#), [history table](#), [hot backup](#), [progress table](#), [server](#), [warm backup](#).

#### crash recovery

The cleanup activities for InnoDB tables that occur when MySQL is started again after a crash. Changes that were committed before the crash, but not yet written to the tablespace files, are reconstructed from the **doublewrite buffer**. When the database is shut down normally, this type of activity is performed during shutdown by the **purge** operation.

## D

#### data dictionary

A set of tables, controlled by the InnoDB storage engine, that keeps track of InnoDB-related objects such as tables, indexes, and table columns. These tables are part of the InnoDB **system tablespace**.

Because the **MySQL Enterprise Backup** product always backs up the system tablespace, all backups include the contents of the data dictionary.

See Also [hot backup](#), [MySQL Enterprise Backup](#), [system tablespace](#).

---

## database

A set of tables and related objects owned by a MySQL user. Equivalent to “schema” in Oracle Database terminology. **MySQL Enterprise Backup** can perform a **partial backup** that includes some databases and not others. The full set of databases controlled by a MySQL server is known as an **instance**.

See Also [instance](#), [partial backup](#).

## differential backup

A backup that captures only the data changed since the last full backup. It has the potential to be smaller and faster than a **full backup**, but is usually bigger and takes longer to create than an **incremental backup**.

See [Section 4.3.3, “Making a Differential or Incremental Backup”](#) for usage details. Related `mysqlbackup` options are `--incremental`, `--incremental-with-redo-log-only`, `--incremental-backup-dir`, `--incremental-base`, and `--start-lsn`.

See Also [full backup](#), [incremental backup](#).

## downtime

A period when the database is unresponsive. The database might be entirely shut down, or in a read-only state when applications are attempting to insert, update, or delete data. The goal for your backup strategy is to minimize downtime, using techniques such as **hot backup** for InnoDB tables, **cold backup** using **slave** servers in a **replication** configuration, and minimizing the duration of the **suspend** stage where you run customized backup logic while the MySQL server is **locked**.

See Also [cold backup](#), [hot backup](#), [InnoDB](#), [locking](#), [replication](#), [slave](#), [suspend](#).

# E

## exclude

In a **partial backup**, to select a set of tables, databases, or a combination of both to be omitted from the backup. Contrast with **include**.

See Also [partial backup](#).

## extract

The operation that retrieves some content from an **image** file produced by a **single-file backup**. It can apply to a single file (unpacked to an arbitrary location) or to the entire backup (reproducing the original directory structure of the backup data). These two kinds of extraction are performed by the `mysqlbackup` options `extract` and `image-to-backup-dir`, respectively.

See Also [image](#), [single-file backup](#).

# F

## .frm file

A file containing the metadata, such as the table definition, of a MySQL table.

For backups, you must always keep the full set of `.frm` files along with the backup data to be able to restore tables that are altered or dropped after the backup.

Although each InnoDB table has a `.frm` file, InnoDB maintains its own table metadata in the system tablespace; the `.frm` files are not needed for InnoDB to operate on InnoDB tables.

These files are backed up by the **MySQL Enterprise Backup** product. These files must not be modified by an `ALTER TABLE` operation while the backup is taking place, which is why backups that include non-InnoDB tables perform a `FLUSH TABLES WITH READ LOCK` operation to freeze such activity while backing up the `.frm` files. Restoring a backup can result in `.frm` files being created, changed, or removed to match the state of the database at the time of the backup.

## file format

The format used by InnoDB for its data files named `ibdata1`, `ibdata2`, and so on. Each file format supports one or more row formats.

---

See Also [Antelope](#), [Barracuda](#), [ibdata file](#), [row format](#).

#### full backup

A **backup** that includes all the **tables** in each MySQL database, and all the databases in a MySQL instance. Contrast with **partial backup** and **incremental backup**. Full backups take the longest, but also require the least amount of followup work and administration complexity. Thus, even when you primarily do partial or incremental backups, you might periodically do a full backup.

See Also [backup](#), [incremental backup](#), [partial backup](#), [table](#).

## H

#### history table

The table `mysql.backup_history` that holds details of completed **backup** operations. While a backup job is running, the details (especially the changing status value) are recorded in the **progress table**.

See Also [backup](#), [progress table](#).

#### hot backup

A backup taken while the MySQL **instance** and is running and applications are reading and writing to it. Contrast with **warm backup** and **cold backup**.

A hot backup involves more than simply copying data files: it must include any data that was inserted or updated while the backup was in process; it must exclude any data that was deleted while the backup was in process; and it must ignore any changes started by **transactions** but not committed.

The Oracle product that performs hot backups, of **InnoDB** tables especially but also tables from MyISAM and other storage engines, is **MySQL Enterprise Backup**.

The hot backup process consists of two stages. The initial copying of the InnoDB data files produces a **raw backup**. The **apply** step incorporates any changes to the database that happened while the backup was running. Applying the changes produces a **prepared backup**; these files are ready to be restored whenever necessary.

A **full backup** consists of a hot backup phase that copies the InnoDB data, followed by a **warm backup** phase that copies any non-InnoDB data such as MyISAM tables and **.frm** files.

See Also [apply](#), [cold backup](#), [.frm file](#), [full backup](#), [InnoDB](#), [instance](#), [prepared backup](#), [raw backup](#), [warm backup](#).

## I

#### .ibd file

Each InnoDB **tablespace** created using the **file-per-table** setting has a filename with a **.ibd** extension. This extension does not apply to the **system tablespace**, which is made up of files named `ibdata1`, `ibdata2`, and so on.

See Also [.ibz file](#), [system tablespace](#), [tablespace](#).

#### .ibz file

When the **MySQL Enterprise Backup** product performs a **compressed backup**, it transforms each **tablespace** file that is created using the **file-per-table** setting from a **.ibd** extension to a **.ibz** extension.

The compression applied during backup is distinct from the **compressed row format** that keeps table data compressed during normal operation. An InnoDB tablespace that is already in compressed row format is not compressed a second time, but is, nevertheless, still saved as an **.ibz** file in the compressed backup.

See Also [.bz file](#), [compression](#), [compression level](#), [.ibd file](#), [.ibz file](#), [MySQL Enterprise Backup](#), [tablespace](#).

#### ibdata file

A set of files with names such as `ibdata1`, `ibdata2`, and so on, that make up the InnoDB **system tablespace**. These files contain metadata about InnoDB tables, and can contain some or all of the table and index data

---

also (depending on whether the **file-per-table option** is in effect when each table is created). For backward compatibility these files always use the **Antelope** file format.  
See Also [Antelope](#), [system tablespace](#).

#### image

The file produced as part of a **single-file backup** operation. It can be a real file that you store locally, or standard output (specified as `-`) when the backup data is **streamed** directly to another command or remote server. This term is referenced in several `mysqlbackup` options such as `backup-dir-to-image` and `image-to-backup-dir`.  
See Also [single-file backup](#), [streaming](#).

#### include

In a **partial backup**, to select a set of tables, databases, or a combination of both to be backed up. Contrast with **exclude**.  
See Also [partial backup](#).

#### incremental backup

A backup that captures only data changed since the previous backup. It has the potential to be smaller and faster than a **full backup**. The incremental backup data must be merged with the contents of the previous backup before it can be restored. See [Section 4.3.3, “Making a Differential or Incremental Backup”](#) for usage details. Related `mysqlbackup` options are `--incremental`, `--incremental-with-redo-log-only`, `--incremental-backup-dir`, `--incremental-base`, and `--start-lsn`.  
See Also [full backup](#).

#### InnoDB

The type of MySQL **table** that works best with **MySQL Enterprise Backup**. These tables can be backed up using the **hot backup** technique that avoids interruptions in database processing. For this reason, and because of the higher reliability and concurrency possible with InnoDB tables, most deployments should use InnoDB for the bulk of their data and their most important data. In MySQL 5.5 and higher, the `CREATE TABLE` statement creates InnoDB tables by default.  
See Also [hot backup](#), [table](#).

#### instance

The full contents of a MySQL server, possibly including multiple **databases**. A **backup** operation can back up an entire instance, or a **partial backup** can include selected databases and tables.  
See Also [database](#), [partial backup](#).

## L

#### locking

See Also [suspend](#), [warm backup](#).

#### log

Several types of log files are used within the MySQL Enterprise Backup product. The most common is the InnoDB **redo log** that is consulted during **incremental backups**.  
See Also [incremental backup](#), [redo log](#).

#### log sequence number

See [LSN](#).

#### logical backup

A **backup** that reproduces table structure and data, without copying the actual data files. For example, the `mysqldump` command produces a logical backup, because its output contains statements such as `CREATE TABLE` and `INSERT` that can re-create the data. Contrast with **physical backup**.  
See Also [backup](#), [physical backup](#).



---

## LSN

Acronym for **log sequence number**. This arbitrary, ever-increasing value represents a point in time corresponding to operations recorded in the **redo log**. (This point in time is regardless of transaction boundaries; it can fall in the middle of one or more transactions.) It is used internally by InnoDB during **crash recovery** and for managing the buffer pool.

In the **MySQL Enterprise Backup** product, you can specify an LSN to represent the point in time from which to take an **incremental backup**. The relevant LSN is displayed by the output of the `mysqlbackup` command. Once you have the LSN corresponding to the time of a full backup, you can specify that value to take a subsequent incremental backup, whose output contains another LSN for the next incremental backup.

See Also [crash recovery](#), [hot backup](#), [incremental backup](#), [redo log](#).

## M

### .MRG file

A file containing references to other tables, used by the **MERGE** storage engine. Files with this extension are always included in backups produced by the `mysqlbackup` command of the **MySQL Enterprise Backup** product.

### .MYD file

A file that MySQL uses to store data for a MyISAM table.

See Also [.MYI file](#).

### .MYI file

A file that MySQL uses to store indexes for a MyISAM table.

See Also [.MYD file](#).

### manifest

The record of the environment (for example, command-line arguments) and data files involved in a backup, stored in the files `meta/backup_create.xml` and `meta/backup_content.xml`, respectively. This data can be used by management tools during diagnosis and troubleshooting procedures.

### master

In a **replication** configuration, a database server that sends updates to a set of **slave** servers. It typically dedicates most of its resources to write operations, leaving user queries to the slaves. With **MySQL Enterprise Backup**, typically you perform backups on the slave servers rather than the master, to minimize any slowdown of the overall system.

See Also [replication](#), [slave](#).

### media management software

A class of software programs for managing backup media, such as libraries of tape backups. One example is **Oracle Secure Backup**. Abbreviated **MMS**.

See Also [Oracle Secure Backup](#).

### my.cnf

The typical name for the MySQL **configuration file** on Linux, Unix, and OS X systems.

See Also [configuration file](#), [my.ini](#).

### my.ini

The typical name for the MySQL **configuration file** on Windows systems.

See Also [configuration file](#), [my.cnf](#).

### MyISAM

A MySQL storage engine, formerly the default for new tables. In MySQL 5.5 and higher, **InnoDB** becomes the default storage engine. MySQL Enterprise Backup can back up both types of tables, and tables from other

---

storage engines also. The backup process for InnoDB tables (**hot backup**) is less disruptive to database operations than for MyISAM tables (**warm backup**).  
See Also [hot backup](#), [InnoDB](#), [warm backup](#).

#### MySQL Enterprise Backup

A licensed products that performs **hot backups** of MySQL databases. It offers the most efficiency and flexibility when backing up **InnoDB** tables; it can also back up MyISAM and other kinds of tables. It is included as part of the MySQL Enterprise Edition subscription.  
See Also [Barracuda](#), [hot backup](#), [InnoDB](#).

#### mysqlbackup

The primary command of the **MySQL Enterprise Backup** product. Different options perform **backup** and **restore** operations.  
See Also [backup](#), [restore](#).

#### mysqldump

A MySQL command that performs **logical backups**, producing a set of SQL commands to recreate tables and data. Suitable for smaller backups or less critical data, because the **restore** operation takes longer than with a **physical backup** produced by **MySQL Enterprise Backup**.  
See Also [logical backup](#), [physical backup](#), [restore](#).

## N

#### non-TTS backup

A backup that is NOT created using [transportable tablespace \(TTS\)](#), that is, not with the `--use-tts` option.  
See Also [transportable tablespace](#).

## O

#### .opt file

A file containing database configuration information. Files with this extension are always included in backups produced by the backup operations of the **MySQL Enterprise Backup** product.

#### offline

A type of operation performed while the database server is stopped. With the **MySQL Enterprise Backup** product, the main offline operation is the **restore** step. You can optionally perform a **cold backup**, which is another offline operation. Contrast with **online**.  
See Also [cold backup](#), [online](#), [restore](#).

#### online

A type of operation performed while the database server is running. A **hot backup** is the ideal example, because the database continues to run and no read or write operations are blocked. For that reason, sometimes “hot backup” and “online backup” are used as synonyms. A **cold backup** is the opposite of an online operation; by definition, the database server is shut down while the backup happens. A **warm backup** is also a kind of online operation, because the database server continues to run, although some write operations could be blocked while a warm backup is in progress. Contrast with **offline**.  
See Also [cold backup](#), [hot backup](#), [offline](#), [warm backup](#).

#### optimistic backup

Optimistic backup is a feature introduced in MySQL Enterprise Backup 3.11 for improving performance for backing up and restoring huge databases in which only a small number of tables are modified frequently. An optimistic backup consists of two phases: (1) the optimistic phase in which tables that are unlikely to be modified during the backup process (identified by the user with the `optimistic-time` option or, by exclusion, with the `optimistic-busy-tables` option) are backed up without locking the MySQL instance; (2) a normal phase,

---

in which tables that are not backed up in the first phase are being backed up in a manner similar to how they are processed in an ordinary backup: the InnoDB files are copied first, and then other relevant files and copied or processed with the MySQL instance locked. The redo logs, undo logs, and the system tablespace are also backed up in this phase. See [Section 4.3.6, “Making an Optimistic Backup”](#) for details.

#### Oracle Secure Backup

An Oracle product for managing **backup** media, and so classified as **media management software (MMS)**. Abbreviated **OSB**. For **MySQL Enterprise Backup**, OSB is typically used to manage tape backups. See Also [backup](#), [media management software](#), [OSB](#).

#### OSB

Abbreviation for **Oracle Secure Backup**, a **media management software** product (**MMS**). See Also [Oracle Secure Backup](#).

## P

#### .par file

A file containing partition definitions. Files with this extension are always included in backups produced by the [mysqlbackup](#) command of the **MySQL Enterprise Backup** product.

#### parallel backup

The default processing mode in MySQL Enterprise Backup 3.8 and higher, employing multiple threads for different classes of internal operations (read, process, and write). See [Section 1.2, “Overview of Backup Types”](#) for an overview, [Section 13.10, “Performance / Scalability / Capacity Options”](#) for the relevant [mysqlbackup](#) options, and [Chapter 7, Performance Considerations for MySQL Enterprise Backup](#) for performance guidelines and tips.

#### partial backup

A **backup** that contains some of the **tables** in a MySQL database, or some of the databases in a MySQL instance. Contrast with **full backup**. Related [mysqlbackup](#) options are [--include-tables](#), [--exclude-tables](#), [--use-tts](#), [--only-known-file-types](#), and [--only-innodb](#). See Also [backup](#), [database](#), [full backup](#), [partial restore](#), [table](#).

#### partial restore

A **restore** operation that applies to one or more **tables** or **databases**, but not the entire contents of a MySQL server. The data being restored could come from either a **partial backup** or a **full backup**. Related [mysqlbackup](#) options are [--include-tables](#), [--exclude-tables](#), and [--rename](#). See Also [database](#), [full backup](#), [partial backup](#), [restore](#), [table](#).

#### physical backup

A **backup** that copies the actual data files. For example, the **MySQL Enterprise Backup** command produces a physical backup, because its output contains data files that can be used directly by the [mysqld](#) server. Contrast with **logical backup**. See Also [backup](#), [logical backup](#).

#### point in time

The time corresponding to the end of a **backup** operation. A **prepared backup** includes all the changes that occurred while the backup operation was running. **Restoring** the backup brings the data back to the state at the moment when the backup operation completed. See Also [backup](#), [prepared backup](#), [restore](#).

#### prepared backup

The set of backup data that is entirely consistent and ready to be restored. It is produced by performing the **apply** operation on the **raw backup**. See Also [apply](#), [raw backup](#).

---

progress table

The table `mysql.backup_progress` that holds details of running **backup** operations. When a backup job finishes, the details are recorded in the **history table**.

See Also [backup](#), [history table](#).

## R

raw backup

The initial set of backup data, not yet ready to be restored because it does not incorporate changes that occurred while the backup was running. The **apply** operation transforms the backup files into a **prepared backup** that is ready to be restored.

See Also [apply](#), [prepared backup](#).

redo log

A set of files, typically named `ib_logfile0` and `ib_logfile1`, that record statements that attempt to change data in InnoDB tables. These statements are replayed automatically to correct data written by incomplete transactions, on startup following a crash. The passage of data through the redo logs is represented by the ever-increasing **LSN** value. The 4GB limit on maximum size for the redo log is raised in MySQL 5.6.

See Also [LSN](#).

regular expression

Some MySQL Enterprise Backup features use POSIX-style regular expressions, for example to specify tables, databases, or both to **include** or **exclude** from a **partial backup**. Regular expressions require escaping for dots in filenames, because the dot is the single-character wildcard; no escaping is needed for forward slashes in path names. When specifying regular expressions on the command line, surround them with quotation marks as appropriate for the shell environment, to prevent expansion of characters such as asterisks by the shell wildcard mechanism.

See Also [exclude](#), [include](#), [partial backup](#).

relay log

A record on a **slave** server for the events read from the **binary log** of the master and written by the slave I/O thread. The relay log, like the **binary log**, consists of a set of numbered files containing events that describe database changes, and an index file that contains the names of all used relay log files. For more information on relay log, see [The Slave Relay Log](#). The relay log on a server is backed up by default.

See Also [binary log](#), [replication](#).

replication

A common configuration for MySQL deployments, with data and DML operations from a **master** server synchronized with a set of **slave** servers. With MySQL **Enterprise Backup**, you might take a backup on one server, and restore on a different system to create a new slave server with the data already in place. You might also back up data from a slave server rather than the master, to minimize any slowdown of the overall system.

See Also [master](#), [slave](#).

repository

We distinguish between the **server repository** and the **backup repository**.

See Also [backup repository](#), [server repository](#).

restore

The converse of the **backup** operation. The data files from a **prepared backup** are put back into place to repair a data issue or bring the system back to an earlier state.

See Also [backup](#), [prepared backup](#).

row format

The disk storage format for a row from an InnoDB table. As InnoDB gains new capabilities such as compression, new row formats are introduced to support the resulting improvements in storage efficiency and performance.

---

Each table has its own row format, specified through the [ROW\\_FORMAT](#) option. To see the row format for each InnoDB table, issue the command [SHOW TABLE STATUS](#). Because all the tables in the system tablespace share the same row format, to take advantage of other row formats typically requires setting the [innodb\\_file\\_per\\_table](#) option, so that each table is stored in a separate tablespace.

## S

### SBT

Acronym for **system backup to tape**.

See Also [system backup to tape](#).

### selective backup

Another name for **partial backup**

See Also [partial backup](#), [selective restore](#).

### selective restore

Another name for **partial restore**

See Also [partial restore](#), [selective backup](#).

### server

A MySQL **instance** controlled by a [mysqld](#) daemon. A physical machine can host multiple MySQL servers, each requiring its own **backup** operations and schedule. Some backup operations communicate with the server through a **connection**.

See Also [connection](#), [instance](#).

### server repository

Contrast with **backup repository**.

See Also [backup repository](#), [repository](#).

### single-file backup

A backup technique that packs all the backup data into one file (the backup **image**), for ease of storage and transfer. The **streaming** backup technique requires using a single-file backup.

See Also [image](#), [streaming](#).

### slave

In a **replication** configuration, a database server that receives updates from a **master** server. Typically used to service user queries, to minimize the query load on the master. With **MySQL Enterprise Backup**, you might take a backup on one server, and restore on a different system to create a new slave server with the data already in place. You might also back up data from a slave server rather than the master, to minimize any slowdown of the overall system.

See Also [master](#), [replication](#).

### streaming

A backup technique that transfers the data immediately to another server, rather than saving a local copy. Uses mechanisms such as Unix pipes. Requires a **single-file backup**, with the destination file specified as `-` (standard output).

See Also [single-file backup](#).

### suspend

An optional stage within the backup where the MySQL Enterprise Backup processing stops, to allow for user-specific operations to be run. The [mysqlbackup](#) command has options that let you specify commands to be run while the backup is suspended. Most often used in conjunction with backups of **InnoDB** tables only, where you might do your own scripting for handling **.frm files**.

See Also [.frm file](#), [InnoDB](#).

---

system backup to tape

An API for **media management software**. Abbreviated **SBT**. Several [mysqlbackup](#) options (with **sbt** in their names) pass information to **media management software** products such as **Oracle Secure Backup**.

See Also [Oracle Secure Backup](#), [SBT](#).

system tablespace

By default, this single data file stores all the table data for a database, as well as all the metadata for InnoDB-related objects (the **data dictionary**).

Turning on the **innodb\_file\_per\_table** option causes each newly created table to be stored in its own **tablespace**, reducing the size of, and dependencies on, the system tablespace.

Keeping all table data in the system tablespace has implications for the **MySQL Enterprise Backup** product (backing up one large file rather than several smaller files), and prevents you from using certain InnoDB features that require the newer **Barracuda** file format. on the

See Also [Barracuda](#), [data dictionary](#), [file format](#), [ibdata file](#), [tablespace](#).

## T

.TRG file

A file containing **trigger** parameters. Files with this extension are always included in backups produced by the [mysqlbackup](#) command of the **MySQL Enterprise Backup** product.

table

Although a table is a distinct, addressable object in the context of SQL, for **backup** purposes we are often concerned with whether the table is part of the **system tablespace**, or was created under the **file-per-table** setting and so resides in its own **tablespace**.

See Also [backup](#), [system tablespace](#), [tablespace](#).

tablespace

For **InnoDB** tables, the file that holds the data and indexes for a table. Can be either the **system tablespace** containing multiple tables, or a table created with the **file-per-table** setting that resides in its own tablespace file.

See Also [InnoDB](#), [system tablespace](#).

transportable tablespace

A feature that allows a **tablespace** to be moved from one instance to another. Traditionally, this has not been possible for InnoDB tablespaces because all table data was part of the **system tablespace**. In MySQL 5.6 and higher, the [FLUSH TABLES ... FOR EXPORT](#) syntax prepares an InnoDB table for copying to another server; running [ALTER TABLE ... DISCARD TABLESPACE](#) and [ALTER TABLE ... IMPORT TABLESPACE](#) on the other server brings the copied data file into the other instance. A separate [.cfg](#) file, copied along with the **.ibd file**, is used to update the table metadata (for example the **space ID**) as the tablespace is imported. See [Copying File-Per-Table Tablespaces to Another Server](#) for usage information.

Use the [--use-tts](#) option to create a backup with transportable tablespace. See also [Section 5.1.4, "Restoring Backups Created with the --use-tts Option"](#).

See Also [partial backup](#).

TTS

Short form for **transportable tablespace**.

See Also [partial backup](#), [transportable tablespace](#).

---

## W

### warm backup

A **backup** taken while the database is running, but that restricts some database operations during the backup process. For example, tables might become read-only. For busy applications and web sites, you might prefer a **hot backup**.

See Also [backup](#), [cold backup](#), [hot backup](#).





---

# Index

## Symbols

--exclude-tables option, 46  
--include-tables option, 46  
.frm file, 46

## A

Antelope, 89, 175  
apply, 175  
apply-incremental-backup option, 93  
--apply-log option, 93  
.ARM file, 175  
.ARZ file, 175

## B

backup, 175  
backup directory, 110, 175  
backup option, 92  
backup repository, 176  
backup-and-apply-log option, 92  
--backup-dir option, 110  
backup-dir-to-image option, 97  
backup-image option, 109  
backup-my.cnf, 176  
backup-my.cnf file, 7  
backup-to-image option, 92, 123  
backups  
    cold, 5  
    compressed, 5, 45, 50, 59, 114  
    controlling overhead, performance, and scalability, 125  
    encrypted, 73  
    full, 40  
    hot, 5  
    incremental, 5, 41, 115  
    InnoDB tables only, 89  
    logical, 5  
    message logging, 132  
    monitoring, 79  
    optimistic, 50  
    parallel, 5  
    partial, 46, 117  
    physical, 5  
    prepared, 7, 59  
    preparing to restore, 59  
    progress report, 133  
    raw, 7, 59  
    scheduled, 52  
    single-file, 5, 36  
    streaming, 5, 38  
    to cloud, 39  
    to tape, 39

        with Oracle Secure Backup, 75  
    troubleshooting, 79  
    uncompressed, 5, 49  
    verifying, 32  
    warm, 5  
backup\_content.xml, 7  
backup\_content.xml file, 82  
backup\_create.xml, 7  
backup\_create.xml file, 82  
BACKUP\_HISTORY table, 80  
BACKUP\_PROGRESS table, 80  
backup\_variables.txt file, 7  
Barracuda, 89, 176  
benchmarking, 67  
binary log, 60, 176  
binlog, 176  
.bz file, 175

## C

cloud backups, 39  
    --cloud-access-key-id option, 138  
    --cloud-aws-region option, 138  
    --cloud-identity-url, 139  
    --cloud-object, 138  
    --cloud-password, 139  
    --cloud-region, 139  
    --cloud-tempauth-url, 139  
    --cloud-tenant, 139  
    --cloud-user-id, 138  
    --counter-container, 138  
    --cloud-bucket option, 138  
    --cloud-object-key option, 138  
    --cloud-proxy option, 137  
    --cloud-secret-access-key option, 138  
    --cloud-service option, 137  
    --cloud-trace option, 137  
cold backup, 5, 176  
command-line tools, 5  
    --comments option, 114  
    --comments-file option, 114  
comments.txt file, 7, 114  
    --compress option, 45, 114  
    --compress-level option, 45, 115  
    --compress-method option, 114  
compressed backups, 5, 45, 50, 59, 114  
compression, 177  
compression level, 177  
configuration file, 177  
configuration options, 141  
connection, 177  
connection options, 106  
copy-back option, 55, 95  
copy-back-and-apply-log option, 33, 94

---

- corruption problems, 80
- crash recovery, 59, 177
- cron jobs, 52
- .CSM file, 7
- .CSV file, 7

## D

- data dictionary, 177
- database, 178
- databases option, 121
- databases-list-file option, 122
- datadir directory, 7
- datadir option, 107
- data\_home\_dir option, 107
- decrypt option, 136
- decryption, 73
- differential backup, 178
- disable-manifest option, 125
- disk storage for backup data, 5, 38
- distributed file system, 53
- downtime, 178
- dst-entry option, 124

## E

- encrypt option, 136
- encrypted backups, 73
- encryption, 73
- error codes, 79
- exclude, 178
- exclude-tables option, 118
- exec-when-locked option, 140
- extract, 178
- extract option, 97, 123

## F

- FAQ, 147
- file format, 178
- files backed up, 7
- frequently asked questions, 147
- .frm file, 7, 178
- full backup, 40, 179

## G

- GRANT statement, 29

## H

- history table, 179
- hot backup, 5, 179

## I

- ibbackup\_logfile file, 7
- .ibd file, 7, 179
- ibdata file, 7, 179

- ibreset command, 80
- .bz file, 7
- .ibz file, 7, 179
- ib\_logfile file, 7
- image, 180
- image-to-backup-dir option, 97, 123, 123
- image\_files.xml file, 7, 82
- include, 180
- include option, 46, 121
- include-tables option, 117
- incremental backup, 5, 115, 180
- incremental option, 115
- incremental-backup-dir option, 117
- incremental-base option, 116
- incremental-with-redo-log-only option, 116
- InnoDB, 180
- InnoDB tables, 5, 7, 89, 89
  - compressed backup feature, 45
  - incremental backup feature, 41
- installing MySQL Enterprise Backup, 19
- instance, 180

## K

- key option, 137
- key-file option, 137

## L

- limit-memory option, 128
- list-image option, 97, 123
- locking, 180
- log, 7, 92, 180
- log-bin-index, 130
- logical backup, 5, 180
- logs
  - of backup operations, 80
- LSN, 41, 115, 181

## M

- manifest, 7, 82, 125, 181
- master, 65, 181
- master-info-file, 131
- media management software, 181
- media management software (MMS) products, 75
- MEMORY tables, 52
- message logging, 132
- meta directory, 7
- MMS products, 75
- monitoring backup jobs, 79
- .MRG file, 181
- my.cnf, 181
- my.ini, 181
- .MYD file, 7
- .MYD file, 181

---

- .MYI file, 7
- .MYI file, 181
- MyISAM, 181
- MyISAM tables, 89
- MySQL Enterprise Backup, 182
- MySQL Enterprise Monitor, 79
- mysqlbackup, 89, 182
  - and media management software (MMS) products, 75
  - configuration options, 141
  - examples, 36
  - files produced, 7
  - modes of operation, 91
  - options, 99
  - overview, 5
  - required privileges, 29
  - using, 27
- mysqlbinlog command, 60
- mysqldump, 52, 182

## N

- no-history-logging option, 113
- no-locking option, 128
- non-TTS backup, 182
- number-of-buffers option, 126

## O

- offline, 182
- on-disk-full option, 129
- online, 182
- only-innodb option, 119
- rename option, 120
- only-innodb option, 122
- only-known-file-types option, 118
- .opt file, 7, 182
- optimistic backup, 50, 131, 132, 182
- optimistic-busy-tables, 132
- optimistic-time, 131
- options, mysqlbackup, 99
  - connection, 106
  - for cloud storage, 137
  - for compression, 114
  - for controlling backup overhead, performance, and scalability, 125
  - for controlling message logging, 132
  - for controlling progress reporting, 133
  - for encryption, 136
  - for generating metadata, 113
  - for incremental backups, 115
  - for partial backups, 117
  - for single-file backups, 123
  - for special types of backups, 139
  - in configuration files, 141
  - layout of backup files, 109

- layout of database files, 107
- modes of operation, 91
- options in common with mysql, 104
- standard options, 104
- Oracle Secure Backup, 183
- OSB, 183

## P

- page-reread-count option, 129
- page-reread-time option, 128
- .par file, 7, 183
- parallel backup, 67, 70, 183
- parallel backups, 5
- partial backup, 46, 117, 183
- partial restore, 183
- performance
  - of backups, 67
  - of restores, 70
- performance of backup operations, 5
- physical backup, 5, 183
- point in time, 183
- point-in-time recovery, 60
- posix\_fadvise() system call, 5
- prepared backup, 7, 59, 183
- privileges, 29
- process-threads option, 127
- progress indicator, 133
- progress table, 184
- progress-interval, 136

## R

- RAID, 67, 70
- raw backup, 7, 59, 184
- read-threads option, 126
- redo log, 184
- regular expression, 184
- relay log, 184
- relay-log-index, 130
- relaylog-info-file, 131
- replication, 63, 64, 65, 184
- repository, 184
- restore, 184
- restoring a backup, 55
  - at original location, 33
  - backup created with the --use-tts option, 58
  - examples, 55
  - mysqlbackup options, 93
  - point-in-time recovery, 60
  - preparation, 59
- row format, 184

## S

- SBT, 185

---

- sbt-database-name option, 125
- sbt-environment option, 125
- sbt-lib-path option, 125
- selective backup, 185
- selective restore, 185
- server, 185
- server repository, 185
- show-progress, 133
- single-file backup, 5, 36, 96, 123, 185
- skip-binlog, 129
- skip-relaylog, 130
- skip-unused-pages, 129
- slave, 63, 64, 185
- slave-info option, 139
- sleep option, 128
- space for backup data, 5
- src-entry option, 124
- start-lsn option, 117
- storage access network, 53
- streaming, 38, 185
- streaming backups, 5
- suspend, 185
- suspend-at-end option, 139
- system backup to tape, 186
- system tablespace, 7, 186

## T

- table, 186
- tablespace, 186
- tape backups, 39, 75
- third-party contributions, 155
- transportable tablespace, 186
- .TRG file, 7
- .TRG file, 186
- .TRN file, 7
- troubleshooting for backups, 79
- TTS, 186

## U

- uncompress option, 115
- uncompressed backups, 5, 49
- use-tts option, 119

## V

- validate option, 96
- validating a backup, 95
- verifying a backup, 32

## W

- warm backup, 5, 187
- what is new, 21
- with-timestamp option, 112
- write-threads option, 127