

MySQL Enterprise Backup User's Guide (Version 3.11.1)

Abstract

This is the User's Guide for the MySQL Enterprise Backup product. This manual describes the procedures to back up and restore MySQL databases. It covers techniques for minimizing time and storage overhead during backups, and to keep the database available during backup operations. It illustrates the features and syntax of the `mysqlbackup` command, for example, how to back up selected databases or tables, how to back up only the changes since a previous backup, and how to transfer the backup data efficiently to a different server.

For notes detailing the changes in each release, see the [MySQL Enterprise Backup 3.11 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](#).

Licensing information. This product may include third-party software, used under license. See [this document](#) for licensing information, including licensing information relating to third-party software that may be included in this MySQL Enterprise Backup release.

Document generated on: 2016-08-15 (revision: 6577)

Table of Contents

| | |
|--|-----|
| Preface and Legal Notices | ix |
| I Getting Started with MySQL Enterprise Backup | 1 |
| 1 Introduction to MySQL Enterprise Backup | 5 |
| 1.1 Types of Backups | 5 |
| 1.2 The mysqlbackup Command | 6 |
| 1.3 Overview of Backup Performance and Capacity Considerations | 6 |
| 1.4 Files that Are Backed Up | 7 |
| 1.5 Overview of Restoring a Database | 16 |
| 2 Installing MySQL Enterprise Backup | 17 |
| II Using MySQL Enterprise Backup | 19 |
| 3 Backing Up a Database Server | 23 |
| 3.1 Before the First Backup | 23 |
| 3.1.1 Collect Database Information | 23 |
| 3.1.2 Grant MySQL Privileges to Backup Administrator | 25 |
| 3.1.3 Designate a Location for Backup Data | 25 |
| 3.2 The Typical Backup / Verify / Restore Cycle | 26 |
| 3.2.1 Backing Up an Entire MySQL Instance | 26 |
| 3.2.2 Verifying a Backup | 28 |
| 3.2.3 Restoring a Database at its Original Location | 28 |
| 3.3 Backup Scenarios and Examples | 29 |
| 3.3.1 Making a Full Backup | 29 |
| 3.3.2 Making an Incremental Backup | 30 |
| 3.3.3 Making a Compressed Backup | 33 |
| 3.3.4 Making a Partial Backup | 33 |
| 3.3.5 Making a Single-File Backup | 38 |
| 3.3.6 Making an Optimistic Backup | 41 |
| 3.3.7 Making a Back Up of In-Memory Database Data | 43 |
| 3.3.8 Making Scheduled Backups | 43 |
| 4 Recovering or Restoring a Database | 45 |
| 4.1 Preparing the Backup to be Restored | 45 |
| 4.2 Performing a Restore Operation | 46 |
| 4.3 Point-in-Time Recovery from a Hot Backup | 49 |
| 4.4 Backing Up and Restoring a Single .ibd File | 49 |
| 4.5 Restoring a Backup with a Database Upgrade or Downgrade | 51 |
| 5 mysqlbackup Command Reference | 53 |
| 5.1 mysqlbackup Command-Line Options | 54 |
| 5.1.1 Subcommands | 59 |
| 5.1.2 Standard Options | 66 |
| 5.1.3 Connection Options | 67 |
| 5.1.4 Server Repository Options | 69 |
| 5.1.5 Backup Repository Options | 71 |
| 5.1.6 Metadata Options | 73 |
| 5.1.7 Compression Options | 74 |
| 5.1.8 Incremental Backup Options | 75 |
| 5.1.9 Partial Backup and Restore Options | 78 |
| 5.1.10 Single-File Backup Options | 84 |
| 5.1.11 Performance / Scalability / Capacity Options | 86 |
| 5.1.12 Message Logging Options | 92 |
| 5.1.13 Progress Report Options | 93 |
| 5.1.14 Encryption Options | 96 |
| 5.1.15 Cloud Storage Options | 96 |
| 5.1.16 Options for Special Backup Types | 97 |
| 5.2 Configuration Files and Parameters | 99 |
| 6 Using MySQL Enterprise Backup with Replication | 101 |
| 6.1 Setting Up a New Replication Slave | 101 |

| | |
|---|-----|
| 6.2 Backing up and Restoring a Slave Database | 102 |
| 6.3 Restoring a Master Database | 102 |
| 7 Performance Considerations for MySQL Enterprise Backup | 105 |
| 7.1 Optimizing Backup Performance | 105 |
| 7.2 Optimizing Restore Performance | 108 |
| 8 Encryption for Backups | 111 |
| 9 Using MySQL Enterprise Backup with Media Management Software (MMS) Products | 113 |
| 9.1 Backing Up to Tape with Oracle Secure Backup | 113 |
| 10 Monitoring Backups with MySQL Enterprise Monitor | 115 |
| 11 Troubleshooting for MySQL Enterprise Backup | 117 |
| 11.1 Error codes of MySQL Enterprise Backup | 117 |
| 11.2 Working Around Corruption Problems | 117 |
| 11.3 Using the MySQL Enterprise Backup Logs | 118 |
| 11.4 Using the MySQL Enterprise Backup Manifest | 119 |
| 12 Frequently Asked Questions for MySQL Enterprise Backup | 121 |
| III Appendixes | 123 |
| A MySQL Enterprise Backup Limitations | 127 |
| A.1 Limitations of MySQL Enterprise Backup | 127 |
| B Compatibility Information for MySQL Enterprise Backup | 131 |
| B.1 Cross-Platform Compatibility | 131 |
| B.2 Compatibility with MySQL Versions | 131 |
| B.3 Compatibility with Older Versions of MySQL Enterprise Backup | 131 |
| B.4 Compatibility Notes for Specific MySQL Versions | 131 |
| C Extended Examples | 133 |
| C.1 Sample Directory Structure for Full Backup | 133 |
| C.2 Sample Directory Structure for Compressed Backup | 137 |
| C.3 Sample Directory Structure for Incremental Backup | 137 |
| D MySQL Enterprise Backup Release Notes | 139 |
| MySQL Enterprise Backup Glossary | 141 |
| Index | 153 |

List of Tables

| | |
|---|----|
| 1.1 Files in a MySQL Enterprise Backup Output Directory | 7 |
| 3.1 Information Needed to Back Up a Database | 23 |
| 5.1 List of All Options | 55 |

List of Examples

| | |
|---|-----|
| 3.1 Making an Uncompressed Partial Backup of InnoDB Tables | 36 |
| 3.2 Making a Compressed Partial Backup | 37 |
| 3.3 Single-File Backup to Absolute Path | 38 |
| 3.4 Single-File Backup to Relative Path | 38 |
| 3.5 Single-File Backup to Standard Output | 38 |
| 3.6 Convert Existing Backup Directory to Single Image | 38 |
| 3.7 Extract Existing Image to Backup Directory | 38 |
| 3.8 List Single-File Backup Contents | 39 |
| 3.9 Extract Single-File Backup into Current Directory | 39 |
| 3.10 Extract Single-File Backup into a Backup Directory | 39 |
| 3.11 Selective Extract of Single File | 39 |
| 3.12 Selective Extract of Single Directory | 39 |
| 3.13 Dealing with Absolute Path Names | 39 |
| 3.14 Single-File Backup to a Remote Host | 40 |
| 3.15 Single-file Backup to a Remote MySQL Server | 40 |
| 3.16 Stream a Backup Directory to a Remote MySQL Server | 40 |
| 3.17 Creating a Cloud Backup | 41 |
| 3.18 Extract an Existing Image from Cloud Storage to Backup Directory | 41 |
| 3.19 Optimistic Backup Using the Option <code>optimistic-time=YYMMDDHHMMSS</code> | 42 |
| 3.20 Optimistic Backup Using the Option <code>optimistic-time=now</code> | 42 |
| 3.21 Optimistic Backup Using the <code>optimistic-busy-tables</code> Option | 42 |
| 3.22 Optimistic and Partial Backup Using both the <code>optimistic-busy-tables</code> and <code>optimistic-time</code> Options | 43 |
| 4.1 Applying the Log to a Backup | 46 |
| 4.2 Applying the Log to a Compressed Backup | 46 |
| 4.3 Applying an Incremental Backup to a Full Backup | 46 |
| 4.4 Shutting Down and Restoring a Database | 46 |
| 4.5 Restoring a Compressed Backup | 47 |
| 4.6 Restoring an Encrypted Backup Image | 47 |
| 4.7 Restoring an Incremental Backup Image | 47 |
| 4.8 Restoring Selected Tables from a TTS Backup | 48 |
| 4.9 Restoring a Single-file Backup from Cloud Storage to a MySQL Server | 48 |
| 5.1 Simple Backup with Connection Parameters from Default Configuration File | 60 |
| 5.2 Basic Incremental Backup | 61 |
| 5.3 Apply Log to Full Backup | 61 |
| 5.4 Incremental Backup | 78 |
| 5.5 Example <code>backup-my.cnf</code> file | 99 |
| 9.1 Sample <code>mysqlbackup</code> Commands Using MySQL Enterprise Backup with Oracle Secure Backup | 114 |

Preface and Legal Notices

This is the User Manual for the MySQL Enterprise Backup product.

Licensing information. This product may include third-party software, used under license. See [this document](#) for licensing information, including licensing information relating to third-party software that may be included in this MySQL Enterprise Backup release.

Legal Notices

Copyright © 2003, 2016, 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 Types of Backups | 5 |
| 1.2 The mysqlbackup Command | 6 |
| 1.3 Overview of Backup Performance and Capacity Considerations | 6 |
| 1.4 Files that Are Backed Up | 7 |
| 1.5 Overview of Restoring a Database | 16 |
| 2 Installing MySQL Enterprise Backup | 17 |

Chapter 1 Introduction to MySQL Enterprise Backup

Table of Contents

| | |
|--|----|
| 1.1 Types of Backups | 5 |
| 1.2 The mysqlbackup Command | 6 |
| 1.3 Overview of Backup Performance and Capacity Considerations | 6 |
| 1.4 Files that Are Backed Up | 7 |
| 1.5 Overview of Restoring a Database | 16 |

The MySQL Enterprise Backup product performs backup operations for MySQL data. It can back up all kinds of MySQL tables. It has special optimizations for fast and convenient backups of [InnoDB](#) tables. Because of the speed of InnoDB backups, and the reliability and scalability features of InnoDB tables, we recommend that you use InnoDB tables for your most important data.

This book describes the best practices regarding MySQL backups and documents how to use MySQL Enterprise Backup features to implement these practices. This book teaches you:

- Why backups are important.
- The files that make up a MySQL database and the roles they play.
- How to keep the database running during a backup.
- How to minimize the time, CPU overhead, and storage overhead for a backup job. Often, minimizing one of these aspects increases another.
- How to restore your data when disaster strikes. You learn how to verify backups and practice recovery, so that you can stay calm and confident under pressure.
- Other ways to use backup data for day-to-day administration and in deploying new servers.

1.1 Types of Backups

The various kinds of backup techniques are classified on a scale ranging from hot (the most desirable) to cold (the most disruptive). Your goal is to keep the database system, and associated applications and web sites, operating and responsive while the backup is in progress.

[Hot backups](#) are performed while the database is running. This type of backup does not block normal database operations. It captures even changes that occur while the backup is happening. For these reasons, hot backups are desirable when your database “grows up”: when the data is large enough that the backup takes significant time, and when your data is important enough to your business so that you must capture every last change, without taking your application, web site, or web service offline.

MySQL Enterprise Backup does a hot backup of all InnoDB tables. MyISAM and other non-InnoDB tables are backed up last, using the [warm backup](#) technique: the database continues to run, but the system is in a read-only state during that phase of the backup.

You can also perform [cold backups](#) while the database is stopped. To avoid service disruption, you would typically perform such a backup from a replication slave, which can be stopped without taking down the entire application or web site.

Points to Remember

To back up as much data as possible during the hot backup phase, you can designate InnoDB as the default storage engine for new tables, or convert existing tables to use the InnoDB storage engine. (In MySQL 5.5 and higher, InnoDB is now the default storage engine for new tables.)

During hot and warm backups, information about the structure of the database is retrieved automatically through a database connection. For a cold backup, you must specify file locations through configuration files or command-line options.

1.2 The mysqlbackup Command

When using the MySQL Enterprise Backup product, you primarily work with the `mysqlbackup` command. Based on the options you specify, this command performs all the different types of backup operations, and restore operations too. `mysqlbackup` can do other things that you would otherwise code into your own backup scripts, such as creating a timestamped subdirectory for each backup, compressing the backup data, and packing the backup data into a single file for easy transfer to another server.

For usage information about `mysqlbackup` features, see [Chapter 3, Backing Up a Database Server](#). For option syntax, see [Chapter 5, mysqlbackup Command Reference](#).

1.3 Overview of Backup Performance and Capacity Considerations

In your backup strategy, performance and storage space are key aspects. You want the backup to complete quickly, with little CPU overhead on the database server. 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. All of these aspects are controlled by options of the `mysqlbackup` command.

Sometimes you must balance the different kinds of overhead -- CPU cycles, storage space, and network traffic. Always be aware how much time it takes to restore the data during planned maintenance or when disaster strikes. For example, here are factors to consider for some of the key MySQL Enterprise Backup features:

- **Parallel backups** are the default in MySQL Enterprise Backup 3.8, a major performance improvement over earlier MySQL Enterprise Backup releases. The read, process and write are the primary sub-operations of all MEB operations. For example, in a backup operation, MySQL Enterprise Backup first reads the data from the disk, then processes this data, writes the data to disk, and reads the data again for verification. MySQL Enterprise Backup ensures that these sub-operations are independent of each other and run in parallel to gain performance improvement. Read, process and write sub-operations are performed in parallel using multiple threads of the same kind: multiple read threads, multiple process threads, and multiple write threads, resulting in better performance. The performance improvement is typically greater when RAID arrays are used as both source and target devices, and for compressed backups which can use more CPU cycles in parallel.

Parallel backup employs block-level parallelism, using blocks of 16MB. Different threads can read, process, and write different 16MB chunks within a single file. Parallel backup improves the performance of operations whether the instance contains a single huge `system tablespace`, or many smaller tablespaces (represented by `.ibd` files created in the `innodb_file_per_table` mode).

- **Incremental backups** are faster than full backups, save storage space on the database server, and save on network traffic to transfer the backup data on a different server. Incremental backup requires additional processing to make the backup ready to restore, which you can perform on a different system to minimize CPU overhead on the database server.
- **Compressed backups** save on storage space for InnoDB tables, and network traffic to transfer the backup data on a different server. They do impose more CPU overhead than uncompressed backups. During restore, you need the compressed and uncompressed data at the same time, so take into account this additional storage space and the time to uncompress the data.

In addition to compressing data within InnoDB tables, compressed backups also skip unused space within InnoDB tablespace files. Uncompressed backups include this unused space.

- When space is limited, or you have a storage device such as tape that is cheap, large, but also slow, the performance and space considerations are different. Rather than aiming for the fastest possible backup, you want to avoid storing an intermediate copy of the backup data on the database server. MySQL Enterprise Backup can produce a single-file backup and stream that file directly to the other server or device. Since the backup data is never saved to the local system, you avoid the space overhead on the database server. You also avoid the performance overhead of saving a set of backup files and then bundling them into an archive for transport to another server or storage device. For details, see [Section 3.3.5.1, “Streaming the Backup Data to Another Device or Server”](#).

When streaming backup data to tape, you typically do not compress the backup, because the CPU overhead on the database server to do the compression is more expensive than the additional storage space on the tape device. When streaming backup data to another server, you might compress 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. Or, you might leave the backup data uncompressed on the destination server so that it is ready to be restored on short notice.

For disaster recovery, when speed to restore the data is critical, you might prefer to have critical backup data already prepared and uncompressed, so that the restore operation involves as few steps as possible.

It is during a disaster recovery that speed is most critical. For example, although a [logical backup](#) performed with the `mysqldump` command might take about the same time as a [physical backup](#) with the MySQL Enterprise Backup product (at least for a small database), the MySQL Enterprise Backup restore operation is typically faster. Copying the actual data files back to the data directory skips the overhead of inserting rows and updating indexes that comes from replaying the SQL statements from `mysqldump` output.

To minimize any impact on server performance on Linux and Unix systems, MySQL Enterprise Backup writes the backup data without storing it in the operating system's disk cache, by using the `posix_fadvise()` system call. This technique minimizes any slowdown following the backup operation, by preventing frequently accessed data from being flushed from the disk cache by the large one-time read operation for the backup data.

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

1.4 Files that Are Backed Up


DBA and development work typically involves logical structures such as tables, rows, columns, the data dictionary, and so on. For backups, you must understand the physical details of how these structures are represented by files.

Table 1.1 Files in a MySQL Enterprise Backup Output Directory

| 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 <code>apply-log</code> step applies the same changes to the corresponding backup files. |
| <code>*.ibd</code> | InnoDB file-per-table tablespaces, each containing a single InnoDB table and associated indexes. | Used for tables created under the <code>innodb_file_per_table</code> . Because the original files might change while the backup is in progress, the <code>apply-log</code> step applies the same changes to the corresponding backup files. |


| File Name, Pattern, or Extension | Relation to Original Data Files | Notes |
|----------------------------------|---|--|
| *.ibz | Compressed form of InnoDB data files from the MySQL data directory. | Produced instead of <code>.ibd</code> files in a compressed backup. The <code>ibdata*</code> files representing the InnoDB system tablespace also receive this extension in a compressed backup. The <code>.ibz</code> files are uncompressed for the apply-log 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 without changes. |
| *.MYD | MyISAM table data. | The database is put into a read-only state while these files are copied. These files are copied without changes. |
| *.MYI | MyISAM index data. | The database is put into a read-only state while these files are copied. These files are copied without changes. |
| *.CSM | Metadata for CSV tables. | These files are copied without changes. The <code>backup_history</code> and <code>backup_progress</code> tables created by <code>mysqlbackup</code> use the CSV format, so the backup always includes some files with this extension. |
| *.CSV | Data for CSV tables. | These files are copied without changes. The <code>backup_history</code> and <code>backup_progress</code> tables created by <code>mysqlbackup</code> 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 without changes. |
| *.TRG | Trigger parameters. | The database is put into a read-only state while these files are copied. These files are copied without changes. |
| *.TRN | Trigger namespace information. | The database is put into a read-only state while these files are copied. These files are copied without changes. |
| *.opt | Database configuration information. | The database is put into a read-only state while these files are copied. These files are copied without changes. |


| File Name, Pattern, or Extension | Relation to Original Data Files | Notes |
|-------------------------------------|--|--|
| <code>*.par</code> | Definitions for partitioned tables. | The database is put into a read-only state while these files are copied. These files are copied without changes. |
| <code>*.ARM</code> | Archive storage engine metadata. | The database is put into a read-only state while these files are copied. These files are copied without changes. |
| <code>*.ARZ</code> | Archive storage engine data. | The database is put into a read-only state while these files are copied. These files are copied without changes. |
| <code>backup-my.cnf</code> | 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. |
| <code>ibbackup_logfile</code> | 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 database 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. |
| <code>ibbackup_redo_log_only</code> | Used instead of <code>ibbackup_logfile</code> for incremental backups taken with the <code>--incremental-with-redo-log-only</code> option. | |
| <code>ib_logfile*</code> | Created in the backup directory during the apply-log 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 apply-log 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. | 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> , to act like a symbolic link pointing to the tablespace file. (See Creating a File-Per-Table Tablespace Outside the Data Directory for details.) The <code>.bl</code> files might or might not be turned back into <code>.isl</code> files during the <code>copy-back</code> operation. If the specified |

| File Name, Pattern, or Extension | Relation to Original Data Files | Notes |
|---|--|--|
| | | directory does not exist on the server where the backup is restored, the <code>mysqlbackup</code> command attempts to create it. If the directory cannot be created, the restore operation fails. Thus, if you use the <code>DATA DIRECTORY</code> clause to put tables in different locations, and 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. |
| 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 whenever you intend to keep more than one set of backup data available under the same main backup directory. |
| <code>datadir</code> directory | A subdirectory that stores all the data files and database subdirectories from the original MySQL instance. | Created under the backup directory by the <code>mysqlbackup</code> command. |
| 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 Section 4.3, “Point-in-Time Recovery from a Hot Backup” 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.</p> <div>  <div> Note <p>Due to some known issues, users should always use the <code>--skip-binlog</code> option</p> </div> </div> |

| File Name, Pattern, or Extension | Relation to Original Data Files | Notes |
|----------------------------------|--|--|
| | | when creating offline backups, or when creating an incremental backup that is based on a full backup created with the <code>--no-locking</code> option. See Appendix A, MySQL Enterprise Backup Limitations for details. |
| 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. | 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. |
| 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 Slave Status Logs , 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 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. |
| <i>image file</i> | A single-file backup produced by the <code>backup-to-image</code> option, | If your backup data directory consists only of zero-byte files, |

| File Name, Pattern, or Extension | Relation to Original Data Files | Notes |
|--|---|--|
| | with a name specified by the <code>--backup-image</code> option. | with a single giant data file in the top-level directory, you have a single-file backup. You can move the image file without losing or damaging the contents inside it, then unpack it with the <code>mysqlbackup</code> command using the <code>extract</code> option 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 the <code>mysqlbackup</code> command. 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 the <code>mysqlbackup</code> command only. Prior to MySQL Enterprise Backup 3.6, this information was in a file named <code>ibbackup_binlog_info</code> . | The <code>mysqlbackup</code> command 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 Section 11.4, “Using the MySQL Enterprise Backup Manifest” . | 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 about this file, see Section 11.4, “Using the MySQL Enterprise Backup Manifest” . | This file is not modified once it is created. You can prevent this file from being generated 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 | This file is not modified once created. You can prevent this file from being generated by |

| File Name, Pattern, or Extension | Relation to Original Data Files | Notes |
|----------------------------------|---|---|
| | 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 Section 11.4, “Using the MySQL Enterprise Backup Manifest” . | 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 Replication with Global Transaction Identifiers for details. When you back up a server with GTIDs enabled, the file <code>gtid_executed.sql</code> is created in the backup directory. Edit and execute this file after restoring the backup data on a slave server. |
| <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. | <p>During a <code>copy-back</code> or <code>copy-back-and-apply-log</code> operation, the server repository options 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>  <div> Warning <p>When using the file to restart the target server, change parameters like <code>--tmpdir</code>, <code>--</code></p> </div> </div> |

| File Name, Pattern, or Extension | Relation to Original Data Files | Notes |
|----------------------------------|--|---|
| | | <p><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.</p> |
| <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>  <p>Warning</p> <p>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</p> </div> |

| File Name, Pattern, or Extension | Relation to Original Data Files | Notes |
|----------------------------------|---------------------------------|---|
| | | absolute filepath to avoid the accidental usage of the wrong file locations by the target server. |

InnoDB Data

Data managed by the InnoDB storage engine is always backed up. The primary InnoDB-related data files that are backed up include the `ibdata*` files that represent the [system tablespace](#) and possibly the data for some user tables; any `.ibd` files, containing data from user tables created with the `file-per-table` setting enabled; 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, 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 to their original location, or for some other use, 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, the `mysqlbackup` command automatically copies the `.frm` files during a backup operation and restores the files during a restore operation.

Data from MyISAM and Other Storage Engines

The `mysqlbackup` command can also back up the `.MYD` files, `.MYI` files, and associated `.frm` files for MyISAM tables. The same applies to files with other extensions, as shown in [this list](#).



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. This phase is a [warm backup](#): 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.



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 associated data.

Generated Files Included in the Backup

The backup data includes 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 `my.cnf` file in the meantime.
- `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 all the files in the backup directory, see [Table 1.1, “Files in a MySQL Enterprise Backup Output Directory”](#).

Single-File Backups

Depending on your workflow, you might perform a single-file backup rather than the typical backup that produces a separate file for every file in the original instance. The single-file format is easier to transfer to a different system, compress and uncompress, and ensure that no backed-up files are deleted later by mistake. It is just as fast as a multi-file backup to do a full restore; restoring individual files can be slower than in a multi-file backup. For instructions, see [Section 3.3.5, “Making a Single-File Backup”](#).

1.5 Overview of Restoring a Database

To initiate the restore process, you run the `mysqlbackup` command with the `copy-back` or the `copy-back-and-apply-log` subcommand. You can restore all the data for a MySQL server: multiple databases, each containing multiple tables. Or, you can restore selected databases, tables, or both.

To repair a problem such as data corruption, you restore the data back to its original location on the original server machine. You might restore to a different server machine or a different location to set up a new replication slave with the data from a master server, or to clone a database for reporting purposes.

See [Chapter 4, *Recovering or Restoring a Database*](#) for instructions on restoring databases.

Chapter 2 Installing MySQL Enterprise Backup

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

Optional: You can also install the MySQL Enterprise Backup product on computers other than the database server, only to run `mysqlbackup` with the `apply-log` option. See [Section 5.1.1.2, “Apply-Log Operations for Existing Backup Data”](#) for information about bringing backup data to a separate server and running the “apply log” step there.

The MySQL Enterprise Backup product 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
```

The `mysqlbackup` command 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` command is installed in the directory `/opt/mysql/meb-3.11`. You must add this directory to your `$PATH` setting.

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 the `mysqlbackup` command 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` command from any directory.

Verify the installation by selecting the menu item **Start > Programs > MySQL Enterprise Backup 3.11 > MySQL Enterprise Backup Command Line**. The menu item displays version information and opens a command prompt for running the `mysqlbackup` command.

Part II Using MySQL Enterprise Backup

Table of Contents

| | | |
|--------|---|-----|
| 3 | Backing Up a Database Server | 23 |
| 3.1 | Before the First Backup | 23 |
| 3.1.1 | Collect Database Information | 23 |
| 3.1.2 | Grant MySQL Privileges to Backup Administrator | 25 |
| 3.1.3 | Designate a Location for Backup Data | 25 |
| 3.2 | The Typical Backup / Verify / Restore Cycle | 26 |
| 3.2.1 | Backing Up an Entire MySQL Instance | 26 |
| 3.2.2 | Verifying a Backup | 28 |
| 3.2.3 | Restoring a Database at its Original Location | 28 |
| 3.3 | Backup Scenarios and Examples | 29 |
| 3.3.1 | Making a Full Backup | 29 |
| 3.3.2 | Making an Incremental Backup | 30 |
| 3.3.3 | Making a Compressed Backup | 33 |
| 3.3.4 | Making a Partial Backup | 33 |
| 3.3.5 | Making a Single-File Backup | 38 |
| 3.3.6 | Making an Optimistic Backup | 41 |
| 3.3.7 | Making a Back Up of In-Memory Database Data | 43 |
| 3.3.8 | Making Scheduled Backups | 43 |
| 4 | Recovering or Restoring a Database | 45 |
| 4.1 | Preparing the Backup to be Restored | 45 |
| 4.2 | Performing a Restore Operation | 46 |
| 4.3 | Point-in-Time Recovery from a Hot Backup | 49 |
| 4.4 | Backing Up and Restoring a Single .ibd File | 49 |
| 4.5 | Restoring a Backup with a Database Upgrade or Downgrade | 51 |
| 5 | mysqlbackup Command Reference | 53 |
| 5.1 | mysqlbackup Command-Line Options | 54 |
| 5.1.1 | Subcommands | 59 |
| 5.1.2 | Standard Options | 66 |
| 5.1.3 | Connection Options | 67 |
| 5.1.4 | Server Repository Options | 69 |
| 5.1.5 | Backup Repository Options | 71 |
| 5.1.6 | Metadata Options | 73 |
| 5.1.7 | Compression Options | 74 |
| 5.1.8 | Incremental Backup Options | 75 |
| 5.1.9 | Partial Backup and Restore Options | 78 |
| 5.1.10 | Single-File Backup Options | 84 |
| 5.1.11 | Performance / Scalability / Capacity Options | 86 |
| 5.1.12 | Message Logging Options | 92 |
| 5.1.13 | Progress Report Options | 93 |
| 5.1.14 | Encryption Options | 96 |
| 5.1.15 | Cloud Storage Options | 96 |
| 5.1.16 | Options for Special Backup Types | 97 |
| 5.2 | Configuration Files and Parameters | 99 |
| 6 | Using MySQL Enterprise Backup with Replication | 101 |
| 6.1 | Setting Up a New Replication Slave | 101 |
| 6.2 | Backing up and Restoring a Slave Database | 102 |
| 6.3 | Restoring a Master Database | 102 |
| 7 | Performance Considerations for MySQL Enterprise Backup | 105 |
| 7.1 | Optimizing Backup Performance | 105 |
| 7.2 | Optimizing Restore Performance | 108 |
| 8 | Encryption for Backups | 111 |
| 9 | Using MySQL Enterprise Backup with Media Management Software (MMS) Products | 113 |
| 9.1 | Backing Up to Tape with Oracle Secure Backup | 113 |
| 10 | Monitoring Backups with MySQL Enterprise Monitor | 115 |
| 11 | Troubleshooting for MySQL Enterprise Backup | 117 |

| | |
|---|-----|
| 11.1 Error codes of MySQL Enterprise Backup | 117 |
| 11.2 Working Around Corruption Problems | 117 |
| 11.3 Using the MySQL Enterprise Backup Logs | 118 |
| 11.4 Using the MySQL Enterprise Backup Manifest | 119 |
| 12 Frequently Asked Questions for MySQL Enterprise Backup | 121 |

Chapter 3 Backing Up a Database Server

Table of Contents

| | |
|--|----|
| 3.1 Before the First Backup | 23 |
| 3.1.1 Collect Database Information | 23 |
| 3.1.2 Grant MySQL Privileges to Backup Administrator | 25 |
| 3.1.3 Designate a Location for Backup Data | 25 |
| 3.2 The Typical Backup / Verify / Restore Cycle | 26 |
| 3.2.1 Backing Up an Entire MySQL Instance | 26 |
| 3.2.2 Verifying a Backup | 28 |
| 3.2.3 Restoring a Database at its Original Location | 28 |
| 3.3 Backup Scenarios and Examples | 29 |
| 3.3.1 Making a Full Backup | 29 |
| 3.3.2 Making an Incremental Backup | 30 |
| 3.3.3 Making a Compressed Backup | 33 |
| 3.3.4 Making a Partial Backup | 33 |
| 3.3.5 Making a Single-File Backup | 38 |
| 3.3.6 Making an Optimistic Backup | 41 |
| 3.3.7 Making a Back Up of In-Memory Database Data | 43 |
| 3.3.8 Making Scheduled Backups | 43 |

This section describes the different kinds of backups that MySQL Enterprise Backup can create and the techniques for producing them, with examples showing the relevant syntax for the `mysqlbackup` command. It also includes a full syntax reference for the `mysqlbackup` command.

3.1 Before the First Backup

The best practices for backups involve planning and strategies. This section outlines some of the preparation needed to put such plans and strategies in place.

3.1.1 Collect Database Information

Before backing up a particular database server for the first time, gather some information and decide on some directory names, as outlined in the following table.

Table 3.1 Information Needed to Back Up a Database

| Information to Gather | Where to Find It | How Used |
|----------------------------------|---|---|
| Path to MySQL configuration file | Default system locations, hardcoded application default locations, or from <code>--defaults-file</code> option in <code>mysqld</code> startup script. | This is the preferred way to convey database configuration information to the <code>mysqlbackup</code> command, using the <code>--defaults-file</code> option. When connection and data layout information is available from the configuration file, you can skip most of the other choices listed below. |
| MySQL port | MySQL configuration file or <code>mysqld</code> startup script. | Used to connect to the database instance during backup operations. Specified via the <code>--port</code> option of <code>mysqlbackup</code> . <code>--port</code> is not needed if available from MySQL configuration file. Not needed |

| Information to Gather | Where to Find It | How Used |
|--|--|--|
| | | when doing an offline (cold) backup, which works directly on the files using OS-level file permissions. |
| 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. Taken from 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 | You choose this. See Section 3.1.3, “Designate a Location for Backup Data” for details. | By default, this directory must be empty for <code>mysqlbackup</code> to write data into it, to avoid overwriting old backups or mixing up data from different backups. Use the <code>--with-timestamp</code> option to automatically create a subdirectory with a unique name, when storing multiple sets of backup data under the same main directory. |
| 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 Section A.1, “Limitations of MySQL Enterprise Backup” 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 Used |
|--------------------------------------|--|---|
| Rate at which redo data is generated | Calculated from the values of the InnoDB logical sequence number at different points in time. Use the technique explained for the --incremental-with-redo-log-only option. | Only needed if you perform incremental backups using the --incremental-with-redo-log-only option rather than the --incremental option. The size of the InnoDB redo log and the rate of generation for redo data dictate how often you must perform incremental backups. |

3.1.2 Grant MySQL Privileges to Backup Administrator

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

The minimum privileges for the MySQL user that [mysqlbackup](#) connects 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 [binlog](#) position, which is stored with the backup.

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';
```

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

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

```
GRANT LOCK TABLES, SELECT, CREATE, ALTER ON *.* TO 'mysqlbackup'@'localhost';
```

3.1.3 Designate a Location for Backup Data

All backup-related operations either create new files or reference existing files underneath a specified directory that holds backup data. Choose this directory in advance, 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 invocations of the [mysqlbackup](#) command.

Once you establish a regular backup schedule with automated jobs, it is preferable to keep each backup within a timestamped subdirectory underneath the main backup directory. To make the [mysqlbackup](#) command create these subdirectories automatically, specify the [--with-timestamp](#) option each time you run [mysqlbackup](#).

For one-time backup operations, for example when cloning a database to set up a replication slave, you might specify a new directory each time, or specify the `--force` option of `mysqlbackup` to overwrite older backup files.

3.2 The Typical Backup / Verify / Restore Cycle

To illustrate the basic steps in making and using a backup, the following examples show how to do a full backup, examine the data files in the backup directory, and then restore the backup to correct an issue with corruption or lost data.

3.2.1 Backing Up an Entire MySQL Instance

In this example, we specify all required options on the command line for illustration purposes. After testing and standardizing the backup procedure, we could move some options to the MySQL configuration file. The options specify connection information for the database and the location to store the backup data. The final option `backup` specifies the type of operation, because `mysqlbackup` can perform several kinds of backup, restore, and pack/unpack operations.

For this example, we specify the final option as `backup-and-apply-log`. This option performs an extra stage after the initial backup, to bring all InnoDB tables up-to-date with any changes that occurred during the backup operation, so that the backup is immediately ready to be restored. For backups of huge or busy databases, you might split up these stages to minimize load on the database server. That is, run `mysqlbackup` first with the `backup` option, transfer the backup to another server, then run `mysqlbackup` with the `apply-log` option to perform the final processing.

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 instance, allowing you to monitor long-running backups and view the results of 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 is just made.

```
$ ./mysqlbackup --user=root -p --backup-dir=/home/admin/backups backup-and-apply-log
```

```
MySQL Enterprise Backup version 3.11.1 Linux-2.6.18-274.el5-i686 [2014/09/04]
Copyright (c) 2003, 2014, Oracle and/or its affiliates. All Rights Reserved.
```

```
mysqlbackup: INFO: Starting with following command line ...
./mysqlbackup --user=root -p --backup-dir=/home/admin/backups
backup-and-apply-log
```

```
mysqlbackup: INFO:
Enter password:
mysqlbackup: INFO: MySQL server version is '5.6.17-log'.
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-and-apply-log' run mysqlbackup
prints "mysqlbackup completed OK!".
```

```
140904 12:34:54 mysqlbackup: INFO: MEB logfile created at /home/admin/backups/meta/MEB_2014-09-04.12-34-54_
```

```
-----
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 = innodb
innodb_undo_directory = /var/lib/mysql/
innodb_undo_tablespace = 0
```

```

innodb_undo_logs = 128

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

datadir = /home/admin/backups/datadir
innodb_data_home_dir = /home/admin/backups/datadir
innodb_data_file_path = ibdata1:12M:autoextend
innodb_log_group_home_dir = /home/admin/backups/datadir
innodb_log_files_in_group = 2
innodb_log_file_size = 50331648
innodb_page_size = 16384
innodb_checksum_algorithm = innodb
innodb_undo_directory = /home/admin/backups/datadir
innodb_undo_tablespace = 0
innodb_undo_logs = 128

mysqlbackup: INFO: Unique generated backup id for this is 14098484948398421

mysqlbackup: INFO: Creating 14 buffers each of size 16777216.
140904 12:34:57 mysqlbackup: INFO: Full Backup operation starts with following threads
      1 read-threads      6 process-threads      1 write-threads
140904 12:34:57 mysqlbackup: INFO: System tablespace file format is Antelope.
140904 12:34:57 mysqlbackup: INFO: Starting to copy all innodb files...
140904 12:34:57 mysqlbackup: INFO: Found checkpoint at lsn 1611237.
140904 12:34:57 mysqlbackup: INFO: Starting log scan from lsn 1610752.
140904 12:34:57 mysqlbackup: INFO: Copying log...
140904 12:34:57 mysqlbackup: INFO: Log copied, lsn 1611237.
140904 12:34:57 mysqlbackup: INFO: Copying /var/lib/mysql/ibdata1 (Antelope file format).
140904 12:34:59 mysqlbackup: INFO: Completing the copy of innodb files.
140904 12:34:59 mysqlbackup: INFO: Starting to copy Binlog files...
140904 12:34:59 mysqlbackup: INFO: Preparing to lock tables: Connected to mysqld server.
140904 12:34:59 mysqlbackup: INFO: Starting to lock all the tables...
140904 12:35:00 mysqlbackup: INFO: All tables are locked and flushed to disk
140904 12:35:00 mysqlbackup: INFO: Opening backup source directory '/var/lib/mysql/'
140904 12:35:00 mysqlbackup: INFO: Starting to backup all non-innodb files in
      subdirectories of '/var/lib/mysql/'
140904 12:35:00 mysqlbackup: INFO: Copying the database directory 'performance_schema'
140904 12:35:01 mysqlbackup: INFO: Completing the copy of all non-innodb files.
140904 12:35:01 mysqlbackup: INFO: Completed the copy of binlog files...
140904 12:35:02 mysqlbackup: INFO: A copied database page was modified at 1611237.
      (This is the highest lsn found on page)
      Scanned log up to lsn 1611237.
      Was able to parse the log up to lsn 1611237.
      Maximum page number for a log record 0
140904 12:35:02 mysqlbackup: INFO: All tables unlocked
140904 12:35:02 mysqlbackup: INFO: All MySQL tables were locked for 2.530 seconds.
140904 12:35:03 mysqlbackup: INFO: Reading all global variables from the server.
140904 12:35:03 mysqlbackup: INFO: Completed reading of all global variables from the server.
140904 12:35:03 mysqlbackup: INFO: Creating server config files server-my.cnf and server-all.cnf in /ho
140904 12:35:03 mysqlbackup: INFO: Full Backup operation completed successfully.
140904 12:35:03 mysqlbackup: INFO: Backup created in directory '/home/admin/backups'
140904 12:35:03 mysqlbackup: INFO: MySQL binlog position: filename mysqld-bin.000001, position 120

-----
Parameters Summary
-----

Start LSN          : 1610752
End LSN            : 1611237
-----

mysqlbackup: INFO: Creating 14 buffers each of size 65536.
140904 12:35:03 mysqlbackup: INFO: Apply-log operation starts with following threads
      1 read-threads      1 process-threads
mysqlbackup: INFO: Using up to 100 MB of memory.
140904 12:35:03 mysqlbackup: INFO: ibbackup_logfile's creation parameters:
      start lsn 1610752, end lsn 1611237,
      start checkpoint 1611237.
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
mysqlbackup: INFO: InnoDB: Setting log file size to 50331648

```

```
mysqlbackup: INFO: InnoDB: Setting log file size to 50331648
140904 12:35:14 mysqlbackup: INFO: We were able to parse ibbackup_logfile up to
lsn 1611237.
mysqlbackup: INFO: Last MySQL binlog file position 0 120, file name mysqld-bin.000001:120
140904 12:35:14 mysqlbackup: INFO: The first data file is '/home/admin/backups/datadir/ibdata1'
and the new created log files are at '/home/admin/backups/datadir'
140904 12:35:14 mysqlbackup: INFO: Apply-log operation completed successfully.
140904 12:35:14 mysqlbackup: INFO: Full backup prepared for recovery successfully.

mysqlbackup completed OK!
```

Now the backup subdirectory is created under the `backup-dir` we specified. The directory name for each new backup is formed from the date and the clock time when the backup run was started, in the local time zone. The backup directory contains the backed-up `ibdata` files and `ibbackup_logfile`. Each subdirectory corresponds to a MySQL database, and contains copies of `.frm`, `.MYD`, `.MYI`, and similar files. For an example of the layout of such a backup directory, see [Section C.1, “Sample Directory Structure for Full Backup”](#).

3.2.2 Verifying a Backup

To verify the backup, restore the backup data on a different server and run the MySQL daemon (`mysqld`) on the new data directory. Then you can execute `SHOW` statements to verify the database and table structure, and execute queries to verify the number of rows, latest updates, and so on.

This is the same general technique to use when you intend to put the backup data to some other use. For example, you might set up a replication slave by making a backup of the master server, or turn a backup into a new MySQL instance for running report queries.



Note

Always do verification against restored data, rather than running `mysqld` with `datadir` pointing to the backup directory. The SQL statements you use to verify the data change the underlying [logical sequence number](#), which would interfere with using the backup directory for subsequent incremental backups.

If you did the backup with the `backup-and-apply-log` option as in the previous example, the backup data is fully consistent and ready to verify. If you only ran the first stage by using the `backup` option, run `mysqlbackup` a second time with the `apply-log` option before doing this verification. (Typically, you run this second phase on the other server after transferring the backup data there, to minimize the load on the original database server.)

See [Chapter 4, *Recovering or Restoring a Database*](#) for the procedure to restore the database files on a different server.

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, combined with the `backup-my.cnf` file in the backup directory, which contains only the small subset of parameters required by `mysqlbackup`. Create a new configuration file by concatenating those two files into a new one, and use that configuration file on the server where you do the verification. Edit the resulting 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.

3.2.3 Restoring a Database at its Original Location

To restore a MySQL instance from a backup:

- Shut down the database server using your usual technique, such as the command `mysqladmin shutdown`.
- Make sure the backup data is fully consistent, by either using the `backup-and-apply-log` option to perform the backup, or running `mysqlbackup` with the `apply-log` option after the initial backup.

- Use the `mysqlbackup` command with the `copy-back` option. This operation copies 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 5.1.1.3, “Restore an Existing Backup”](#).

```
$ mysqlbackup --defaults-file=path_to_my.cnf \  
  --datadir=path_to_data_directory \  
  --backup-dir=path_to_backup_directory copy-back  
...many lines of output...  
mysqlbackup: Finished copying backup files.  
  
101208 16:48:13 mysqlbackup: mysqlbackup completed OK!
```

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

3.3 Backup Scenarios and Examples

All of the following tasks and examples make use of the `mysqlbackup` command. For detailed syntax information, see [Chapter 5, *mysqlbackup Command Reference*](#).

3.3.1 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 do one [full backup](#), you might do [incremental backups](#) (which are smaller and faster) for the next several backup jobs. Periodically, you then do another full backup to begin the cycle again.

This section outlines some of the considerations for making this most basic kind of backup. Because a full backup can take longer and produce larger backup files than other kinds of backups, your decisions about speed, capacity, and convenience are especially important for this part of the backup strategy.

For examples showing the commands to make a full backup, see [Section 3.2.1, “Backing Up an Entire MySQL Instance”](#).

Options on Command Line or in Configuration File?

For clarity, the examples in this manual typically show command-line options to demonstrate connection parameters and other information that might be the same for each backup job. For convenience and consistency, you can include these options in the `[mysqlbackup]` section of the MySQL configuration file that you pass to the `mysqlbackup` command; `mysqlbackup` also picks them up from the `[mysqld]` section if they are present. For example, relying on the port information in the configuration file avoids the need to edit your backup scripts if the database instance switches to a different port.

Output in Single Directory or Timestamped Subdirectories?

For convenience, the `--with-timestamp` option creates uniquely named subdirectories under the backup directory to hold the output from each backup job. The timestamped subdirectories make it simpler to establish retention periods, for example by removing or archiving backup data past a certain age. This option is NOT set on default.

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

With the `--incremental-base` option, as part of each incremental backup command, you specify the directory containing the previous backup. To make the directory names predictable, you might prefer to leave out the `--with-timestamp` option and instead generate a sequence of directory names as part of your backup script.

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 run a full backup each time. Typically, you can save time and storage space by running periodic full backups, and in between running several incremental backups, as described in [Section 3.3.2, “Making an Incremental Backup”](#).

Use Compression or Not?

Creating a compressed backup can save you considerable storage space and reduce I/O significantly. And with the LZ4 compression method (introduced since release 3.10), 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.

3.3.2 Making an Incremental Backup

An [incremental backup](#) only backs up data that changed since the previous backup. This technique provides additional flexibility in designing a backup strategy and reduces required storage for backups.

Incremental backups are typically smaller and take less time than a full backup, making them a good choice for frequent backup jobs. Taking frequent incremental backups ensures you can always restore the database to the same state as a few hours or days in the past, without as much load or storage overhead on the database server as taking frequent full backups.



Note

Because an incremental backup always adds to an existing set of backup files, make at least one [full backup](#) before doing any incremental backups.

Incremental backup is enabled through an option to the `mysqlbackup` command. For straightforward incremental backups, specify the `--incremental` option. An alternative method uses the `--incremental-with-redo-log-only` option, requiring additional planning on your part.

You also indicate 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, you can specify an explicit LSN value using the `--start-lsn` option, using the ending LSN from a previous full or incremental backup.

To prepare the backup data to be restored, you combine each incremental backup 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.

When running the “apply log” step for an incremental backup, you specify the option sequence `--incremental apply-log`, and the paths to 2 MySQL configuration files, first the `.cnf` file pointing to the full backup that you are updating, then the `.cnf` file pointing to the incremental backup data files. If you have taken several incremental backups since the last full backup, you might run several such “apply log” steps, one after the other, to bring the full backup entirely up to date.

Space 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. For non-InnoDB files, the entire file is included in an incremental backup if that file changed since the previous backup.

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

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.

Examples of Incremental Backups

This example uses the `mysqlbackup` command 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 specify the 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 use hardcoded names or generate a sequence of names in your own backup script, rather than using the `--with-timestamp` option.

```
$ mysqlbackup --defaults-file=/home/pekka/.my.cnf --incremental \
--incremental-base=dir:/incr-backup/wednesday \
--incremental-backup-dir=/incr-backup/thursday \
backup
...many lines of output...
mysqlbackup: Backup created in directory '/incr-backup/thursday'
mysqlbackup: start_lsn: 2654255717
mysqlbackup: incremental_base_lsn: 2666733462
mysqlbackup: end_lsn: 2666736714

101208 17:14:58 mysqlbackup: mysqlbackup completed OK!
```

With the `--start-lsn` option, you do have to record the LSN of the previous backup, but then the location of the previous backup is less significant, so you can use `--with-timestamp` to create named subdirectories automatically.

```
$ mysqlbackup --defaults-file=/home/pekka/.my.cnf --incremental \
--start-lsn=2654255716 \
--with-timestamp \
--incremental-backup-dir=/incr-backup \
backup
...many lines of output...
mysqlbackup: Backup created in directory '/incr-backup/2010-12-08_17-14-48'
mysqlbackup: start_lsn: 2654255717
mysqlbackup: incremental_base_lsn: 2666733462
mysqlbackup: end_lsn: 2666736714

101208 17:14:58 mysqlbackup: mysqlbackup completed OK!
```

Wherever you use the `--incremental` option, you can use the `--incremental-with-redo-log-only` option instead. Because `--incremental-with-redo-log-only` is more dependent on the precise LSN than the `--incremental` option is, use the `--incremental-base` option rather than the `--start-lsn` option with this kind of incremental backup.

For this alternative kind of incremental backup to work, the volume of changed information must be low enough, and the [redo log](#) files must be large enough, that all the changes since the previous incremental backup must be present in the redo log and not overwritten. See the `--incremental-with-redo-log-only` option description to learn how to verify those requirements.

```
$ mysqlbackup --defaults-file=/home/pekka/.my.cnf --incremental \
--incremental-base=dir:/incr-backup/wednesday \
--incremental-backup-dir=/incr-backup/thursday \
backup
...many lines of output...
mysqlbackup: Backup created in directory '/incr-backup/thursday'
mysqlbackup: start_lsn: 2654255717
```

```
mysqlbackup: incremental_base_lsn: 2666733462
mysqlbackup: end_lsn: 2666736714

101208 17:14:58 mysqlbackup: mysqlbackup completed OK!
```

See [Section C.3, “Sample Directory Structure for Incremental Backup”](#) for a listing of files from a typical incremental backup.

Once again, we 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:

```
$ 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.

This example shows an incremental backup. The last full backup we ran reported that the highest [LSN](#) was 2638548215:

```
mysqlbackup: Was able to parse the log up to lsn 2638548215
```

We specify that number again in the command here; the incremental backup includes all changes that came *after* the specified LSN.

```
$ mysqlbackup --defaults-file=/home/pekka/.my.cnf --incremental \
--start-lsn=2638548215 \
--incremental-backup-dir=/incr-backup/2010-12-08_17-14-48 \
--backup-dir=/full-backup/2010-12-08_17-14-11 \
backup
..many lines of output...
mysqlbackup: Scanned log up to lsn 2654252454.
mysqlbackup: Was able to parse the log up to lsn 2654252454.
mysqlbackup: Maximum page number for a log record 0
mysqlbackup: Backup contains changes from lsn 2638548216 to lsn 2654252454
101208 17:12:24 mysqlbackup: Incremental backup completed!
```

Next steps:

- Make a note of the LSN value in the message at the end of the backup, for example, `mysqlbackup: Was able to parse the log up to lsn LSN_number`. You specify this value when performing incremental backups of changes that occur after this incremental backup.
- [Apply the incremental backup](#) to the backup files, so that the backup is ready to be restored at any time. You can move the backup data to a different server first, to avoid the CPU and I/O overhead of this operation on the database server itself.
- On a regular schedule, determined by date or amount of database activity, take further [take incremental 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.

3.3.3 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 on 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 rather than the usual `.ibd` extension. To avoid wasting CPU cycles without saving additional disk space, `--compress` does not attempt to compress already-compressed tables that use the Barracuda file format; such tablespace files keep the usual `.ibd` extension.



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.

You can only use the `--compress` option for **full backups**, but 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 5.1.7, “Compression Options”](#) for details.

This is a sample command for making a compressed backup:

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

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
```

Next steps:

- Make a note of the LSN value in the message at the end of both the full and incremental backups (for example, in the line `mysqlbackup: Was able to parse the log up to lsn LSN_number`). You specify this value when performing incremental backups of changes that occur after this full backup.
- [Apply the log](#) to the compressed backup files, so that the full backup is ready to be restored at any time. You can move the backup data to a different server first, to avoid the CPU and I/O overhead of performing this operation on the database server.
- After applying the log, periodically [take incremental backups](#), which are smaller and can be made faster than a full backup.

3.3.4 Making a Partial Backup



Note

To facilitate the creation of partial backups, MySQL Enterprise Backup 3.10 introduces two new options for partial backup: `--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 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 Henry Spencer's 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 5.1.9, “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 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_BACKUPS_DIR/backupdir \ --include-tables='\.emp' \
backup
```

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_BACKUPS_DIR/backupdir \  
--exclude-tables='^(mysql|performance_schema)\.' \  
backup
```

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=$MEB_BACKUPS_DIR/backupdir \  
--include-tables='^sales\.' --exclude-tables='^sales\.hardware$' \  
backup
```

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

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

You can also make [compressed](#), [single-image](#), and other kinds of selective backups by adding the appropriate options.

Next Steps:

- Make a note of the LSN value in the message at the end of both full and incremental backups, for example, `mysqlbackup: Was able to parse the log up to lsn LSN_number`. You specify this value when performing incremental backups of changes that occur after this full backup.
- [Apply the log](#) to the uncompressed backup files, so that the full backup is ready to be restored at any time. You can move the backup data to a different server first, to avoid the CPU and I/O overhead of performing this operation on the database server.
- After applying the log, periodically [take incremental backups](#), which are much faster and smaller than a full backup like these ones.

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 Henry Spencer's regular expression library.

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

IMPORTANT: Although the `mysqlbackup` command 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 3.1 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 3.2 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 the `mysqlbackup` command 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.

3.3.5 Making a Single-File Backup

To avoid a large number of backup files to track and keep safe, and to simplify moving backup data around, the `mysqlbackup` command can create a backup in a single-file format, pack an existing backup into a single file, unpack the single file back to the original backup directory structure, list the contents of a single-file backup, verify the contents of a single-file backup against embedded checksums, or extract a single file or directory tree. For the syntax of the relevant `mysqlbackup` options, see [Section 5.1.10, “Single-File Backup Options”](#).

Because the single-file backup can be streamed or piped to another process, such as a tape backup or a command such as `scp`, you can use this technique to put the backup on another storage device or server without significant storage overhead on the original database server. (During preparation of the single-file backup, some small work files are prepared temporarily inside the specified backup directory.)

To create a single-file backup, use the `mysqlbackup` subcommand `backup-to-image`.

Example 3.3 Single-File Backup to Absolute Path

This command creates a single backup image in 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 3.4 Single-File Backup to Relative Path

This command specifies `--backup-image` with a relative path underneath the backup directory. The resulting single-file backup is created as `/backups/sales.mbi`.

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

Example 3.5 Single-File Backup to Standard Output

The following command dumps the backup output to standard output. Again, the `--backup-dir` directory specified in `my.cnf` is used as a temporary directory.

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

Example 3.6 Convert Existing Backup Directory to Single Image

The `backup-dir` directory specified in `my.cnf` is bundled into the `/backup/my.mbi` file. The directory can contain anything, not necessarily a backup produced by MySQL Enterprise Backup.

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

Example 3.7 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 3.8 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 3.9 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 3.10 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 `--backup-dir` directory.

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

Example 3.11 Selective Extract of Single File

The following command extracts the single file `meta/comments.txt` 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 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` (inside a single-file backup) to standard output.

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

Example 3.12 Selective Extract of Single Directory

The following command extracts a single directory `meta` into a local file system path `./meta`. Extracting a directory extracts all its contents, including any subdirectories.

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

The following command extracts all `meta` directory contents (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 3.13 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.

3.3.5.1 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 `image-to-backup-dir` option on the other end to reproduce the directory structure of a regular backup.

Example 3.14 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 3.15 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 3.16 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'
```

3.3.5.2 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](#).

3.3.5.3 Backing Up to Cloud Storage

MySQL Enterprise Backup supports cloud backup since version 3.10.2.

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 backup or restoration.

Currently, Amazon S3 is the only cloud service supported by MySQL Enterprise Backup.

A cloud backup is created using the cloud options for `mysqlbackup`, which are described in details in [Section 5.1.15, “Cloud Storage Options”](#). This is a sample command for creating a cloud backup:

Example 3.17 Creating a Cloud Backup

```
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
```

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 using cloud storage.

Example 3.18 Extract an Existing Image from Cloud Storage to Backup Directory

Extract a backup image from cloud storage, 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 [Example 4.9, “Restoring a Single-file Backup from Cloud Storage to a MySQL Server”](#) on how to restore a backup image from a cloud storage.

3.3.6 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 the other relevant files and copied or processed with the MySQL instance locked. Also, if it turns out that some of the inactive tables have actually been modified after they were backed up in the optimistic phase, they are copied again in the normal phase. The redo logs, undo logs, and the system table space 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 5.1.11, “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 3.19 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
```

Example 3.20 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
```

Example 3.21 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 3.22 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
```

3.3.7 Making a Back Up of In-Memory Database Data

The `--exec-when-locked` option of the `mysqlbackup` command lets you specify a command and 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 parameter value within single quotes.

3.3.8 Making Scheduled Backups

Maintaining a regular backup schedule is an important measure for preventing data loss for you 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:

```
@daily /path-to-mysqlbackup/mysqlbackup -uroot --backup-dir=/path-to-backup-folder/cronbackups --with
```

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-dir=/path-to-backup-folder/cronbackups --with
```

Check your platform's documentation for 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 your Windows platform for instructions.

Chapter 4 Recovering or Restoring a Database

Table of Contents

| | |
|---|----|
| 4.1 Preparing the Backup to be Restored | 45 |
| 4.2 Performing a Restore Operation | 46 |
| 4.3 Point-in-Time Recovery from a Hot Backup | 49 |
| 4.4 Backing Up and Restoring a Single .ibd File | 49 |
| 4.5 Restoring a Backup with a Database Upgrade or Downgrade | 51 |

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 various 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 properly after being recovered.

4.1 Preparing the Backup to be Restored

Immediately after the backup job completes, the backup files might not be in a consistent state, because data could be inserted, updated, or deleted while the backup is running. These initial backup files are known as the [raw backup](#).

You must update the backup files so that they reflect the state of the database corresponding to a specific InnoDB [log sequence number](#). (The same kind of operation as [crash recovery](#).) When this step is complete, these final files are known as the [prepared backup](#).

During the backup, [mysqlbackup](#) copies the accumulated InnoDB log to a file called [ibbackup_logfile](#). This log file is used to “roll forward” the backed-up data files, so that every page in the data files corresponds to the same log sequence number of the InnoDB log. This phase also creates new [ib_logfiles](#) that correspond to the data files.

The [mysqlbackup](#) option for turning a raw backup into a prepared backup is [apply-log](#). 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.



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](#) option, which allows the data and log files created by the failed [apply-log](#) operation to be overwritten.

For simple backups (without compression or incremental backup), you can combine the initial backup and the [apply-log](#) step using the option [backup-and-apply-log](#).

You can also perform `apply-log` and `copy-back` (which restores the prepared backup) with a single `copy-back-and-apply-log` command.

Example 4.1 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 4.2 Applying the Log to a Compressed Backup

If the backup is compressed, as in [Section 3.3.3, “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
```

Example 4.3 Applying an Incremental Backup to a Full Backup

After you take an incremental backup, as in [Section 3.3.2, “Making an Incremental Backup”](#), the changes reflected in those backup files must be applied to a full backup to bring the full backup up-to-date, in the same way that you apply changes from the binary log.

To bring the data files from the full backup up to date, first run the apply log step so that the data files include any changes that occurred while the full backup was running. Then apply the changes from the incremental backup to the data files produced by the full backup:

```
mysqlbackup --backup-dir=/export/backups/full apply-log
mysqlbackup --backup-dir=/export/backups/full \
  --incremental-backup-dir=/export/backups/incremental \
  apply-incremental-backup
```

Now the data files in the `full-backup` directory are fully up-to-date, as of the time of the incremental backup.

4.2 Performing a Restore Operation

The `mysqlbackup` options to perform a restore operation are `copy-back` and `copy-back-and-apply-log`. The restoration process requires the database server to be already shut down (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 4.4 Shutting Down and Restoring a Database

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



Note

The restored data includes the `backup_history` table, where MySQL Enterprise Backup records details of each backup. Restoring this table to its

earlier state removes information about any subsequent backups that you did. This is the correct starting point for future incremental backups, particularly those using the `--incremental-base` 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. The same cleanup is not required for restoring backups created with the `--use-tts` option (in which case other requirements described in [Restoring Backups Created with the --use-tts Option](#) apply though), and usually not necessary for restoring a partial backup.

You can combine the `apply-log` and the `copy-back` operations (as well as a number of other operations, depending on the kind of backup you are restoring) into a single step by using the `copy-back-and-apply-log` option instead:

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

Example 4.5 Restoring a Compressed Backup

Restore a compressed backup at `<backupDir>` to `<restoreDir>` on the server using `copy-back-and-apply-log`:

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

Do the same for 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
```

See [Section 3.3.3, “Making a Compressed Backup”](#) and [Section 5.1.7, “Compression Options”](#) for more details on compressed backups.

Example 4.6 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>`:

```
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 5.1.14, “Encryption Options”](#) for more details on backup encryption and decryption.

Example 4.7 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.

Alternatively 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:

```
$ 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.

See [Section 3.3.2, “Making an Incremental Backup”](#), and [Section 5.1.8, “Incremental Backup Options”](#), for more details on incremental backups.

Example 4.8 Restoring Selected Tables from a TTS Backup

Selected tables can be restored from a backup created with [transportable tablespaces \(TTS\)](#) (that is, a backup created with the `--use-tts` option) 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”:

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

See [Restoring Backups Created with the --use-tts Option](#) for the special requirements that apply when you restore selected tables from a TTS backup.

Example 4.9 Restoring a Single-file Backup from Cloud Storage to a MySQL Server

Restore a backup image from cloud storage to `datadir` on the server, using the `--backup-dir` option to specify the temporary directory into which the image will be extracted (see [Section 5.1.15, “Cloud Storage Options”](#) for information on the cloud service options):

```
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
```

Restoring Backups Created with the --use-tts Option

There are some special requirements for restoring backups 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` 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.

4.3 Point-in-Time Recovery from a Hot Backup

Using MySQL Enterprise Backup and the [binary log](#) 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](#) file[s] included in the incremental backup that covers t_R . Use the `mysqlbinlog` utility to extract from the [binary log](#) file[s] all the SQL activities that happened in between the t_{LB} (which you noted in step 3) 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](#).

4.4 Backing Up and Restoring a Single .ibd File

A table with a table-specific tablespace (stored in an `.ibd` file) can be restored individually without taking down the MySQL server. This technique is applicable if you delete or update the table data by mistake, without actually losing the table itself through a `DROP TABLE`, `TRUNCATE TABLE`, or `DROP DATABASE` statement.

If you have a clean backup of an `.ibd` file, you can restore it to the MySQL installation from which it originated as follows:

1. For MySQL 5.5 and earlier, the table must already exist and not have been dropped or truncated since taking the backup. When an InnoDB table is truncated, or dropped and recreated, it gets a new table ID. Any ID mismatch between the table in the database and the backed-up table can prevent it from being restored. The requirement for matching table IDs is also the reason why you must restore to the same MySQL server from which the backup data came, not another server with a similar set of databases and tables. This restriction does not apply to MySQL 5.6 and later, as long as the restoration is made from one Generally Available (GA) version to another in the same series of MySQL servers.
2. Prevent write operations for the table to be restored. This prevents users from modifying the table while the restore is in progress.

```
LOCK TABLES tbl_name WRITE;
```

3. Issue this `ALTER TABLE` statement:

```
ALTER TABLE tbl_name DISCARD TABLESPACE;
```

Caution: This deletes the current `.ibd` file.

4. Copy the backup `.ibd` file back to the appropriate database directory.
5. Issue this `ALTER TABLE` statement:

```
ALTER TABLE tbl_name IMPORT TABLESPACE;
```

6. Release the write lock to complete the restore procedure:

```
UNLOCK TABLES;
```

In this context, a clean `.ibd` file backup means:

- There are no uncommitted modifications by transactions in the `.ibd` file.
- There are no unmerged insert buffer entries in the `.ibd` file.
- Purge has removed all delete-marked index records from the `.ibd` file.
- `mysqld` has flushed all modified pages of the `.ibd` file from the buffer pool to the file.

You can make such a clean backup `.ibd` file with the following method:

1. Stop all activity from the `mysqld` server and commit all transactions.
2. Wait until `SHOW INNODB STATUS` shows that there are no active transactions in the database, and the main thread status of InnoDB is `Waiting for server activity`. Then you can make a copy of the `.ibd` file.

Another method for making a clean copy of an `.ibd` file is to use `mysqlbackup`:

1. Use `mysqlbackup` with the `--only-innodb` or `--only-innodb-with-frm` option to back up the InnoDB installation.
2. Run `mysqlbackup ... apply-log` to create a consistent version of the backup database.
3. Start a second (dummy) `mysqld` server on the backup and let it clean up the `.ibd` files in the backup. Wait for the cleanup to end.

4. Shut down the dummy `mysqld` server.
5. Take a clean `.ibd` file from the backup.

4.5 Restoring a Backup with a Database Upgrade or Downgrade



Important

Due to the changes to the InnoDB storage engine going from MySQL 5.5 to 5.6, restoring a backup of a MySQL 5.5 database to a MySQL 5.6 server requires some extra steps beyond the general restore and upgrade procedures, the skipping of which will crash the target server. For such a restoration, follow the [steps](#) described below.

You can back up a server running one MySQL version and restore on a server running a different MySQL version. After the restore, perform the appropriate upgrade steps as if you are running the new MySQL version for the first time. (Or, if you installed on a server running an older MySQL, perform the appropriate downgrade steps.) For information about upgrading and downgrading, see [Upgrading MySQL](#) and [Downgrading MySQL](#).



Note

After upgrading between certain combinations of MySQL versions, you might see error messages about missing or mismatching definitions for system tables. Use the [mysql_upgrade](#) command as directed in the upgrade instructions to correct such issues. See [mysql_upgrade — Check and Upgrade MySQL Tables](#) for instructions on this command.



Warning

Restoring a database to an older MySQL version (i.e., server [downgrading](#)) is only supported when the original and the final versions are in the same release series (e.g. going from 5.5.30 to 5.5.29). [Downgrading](#) to a lower series (e.g. from 5.6.33 to 5.5.33) might cause server crashes or data corruption.

Steps to Back Up on MySQL 5.5 and Restore on MySQL 5.6

- Back up the data on the MySQL 5.5 server.
- Restore the data to the directory you plan to use as the MySQL 5.6 server's data directory by running an [apply-log](#) and then a [copy-back](#) operation on the backup.
- Restart the MySQL 5.5 server, using the intended data directory for the MySQL 5.6 server as its own data directory.



Note

This is an extra step beyond the normal restore and upgrade procedures, special to the restoration of MySQL 5.5 data to MySQL 5.6 server; with it, the MySQL 5.5 server prepares the data for an upgrade to MySQL 5.6 by performing some clean-up steps on the data, similar to what the server would do during a [crash recovery](#).

- Stop the MySQL 5.5 server.
- Install the MySQL 5.6 server.
- Start the MySQL 5.6 server.
- Run [upgrade steps](#) as documented in the MySQL reference manual on the restored data. Make sure the [mysql_upgrade](#) that comes with MySQL 5.6 is applied.

- Check data.

Steps to Back Up on MySQL 5.1 and Restore on MySQL 5.5

- Back up on MySQL 5.1.
- Install MySQL 5.5.
- Restore on MySQL 5.5.
- Run [upgrade steps](#) as documented in the MySQL reference manual.
- Check data.

Chapter 5 mysqlbackup Command Reference

Table of Contents

| | |
|---|----|
| 5.1 mysqlbackup Command-Line Options | 54 |
| 5.1.1 Subcommands | 59 |
| 5.1.2 Standard Options | 66 |
| 5.1.3 Connection Options | 67 |
| 5.1.4 Server Repository Options | 69 |
| 5.1.5 Backup Repository Options | 71 |
| 5.1.6 Metadata Options | 73 |
| 5.1.7 Compression Options | 74 |
| 5.1.8 Incremental Backup Options | 75 |
| 5.1.9 Partial Backup and Restore Options | 78 |
| 5.1.10 Single-File Backup Options | 84 |
| 5.1.11 Performance / Scalability / Capacity Options | 86 |
| 5.1.12 Message Logging Options | 92 |
| 5.1.13 Progress Report Options | 93 |
| 5.1.14 Encryption Options | 96 |
| 5.1.15 Cloud Storage Options | 96 |
| 5.1.16 Options for Special Backup Types | 97 |
| 5.2 Configuration Files and Parameters | 99 |

The `mysqlbackup` command 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 under 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.

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.

Sample command line arguments to start `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 [mysqlbackup], 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 3.1.2, “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 [Section 5.1, “mysqlbackup Command-Line Options”](#). For information about configuration parameters, see [Section 5.2, “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 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 the `mysqlbackup` command 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:

- Do not run long `SELECT` queries or other SQL statements at the time of the backup run.
- 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 `mysqlbackup` has completed successfully, either by verifying that the `mysqlbackup` command returned exit code 0, or by observing that `mysqlbackup` has printed the text “mysqlbackup completed OK!”.
- The `mysqlbackup` command 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 A.1, “Limitations of MySQL Enterprise Backup”](#) for restrictions on backups at the same time as DDL operations.

5.1 mysqlbackup Command-Line Options

The following sections describe the different modes of operation for `mysqlbackup`, and explain in details the applicable options for each of the modes of operation and for different situations.

The table below list all the command options for `mysqlbackup`. Use the hotlinks 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 [Section 5.2, “Configuration Files and Parameters”](#). The `mysqlbackup` command follows 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 5.1 List of All Options

| Format | Description | Introduced |
|--|--|------------|
| --backup-dir | The directory to store the backup data. | |
| --backup-image | Specifies the path name of the backup image. | |
| --backup_innodb_checksum_algorithm | The name of the checksum algorithm used for validating InnoDB tablespaces. | |
| --backup_innodb_data_file_path | Specifies InnoDB system tablespace files' path and size in backup. | |
| --backup_innodb_data_home_dir | Backup base directory for all InnoDB data files in the system tablespace. | |
| --backup_innodb_log_file_size | The size in bytes of each InnoDB backup log file. | |
| --backup_innodb_log_files_in_group | Number of InnoDB log files in backup. | |
| --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. | |
| --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. | |
| --backup_innodb_undo_tablespaces | The number of tablespace files that the undo logs are divided between when a non-zero <code>innodb_undo_logs</code> 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 <code>mysqlbackup</code> access for S3. | |
| --cloud-bucket | The storage bucket on Amazon S3 for the backup image. | |
| --cloud-object-key | The Amazon S3 object key for the backup image. | |
| --cloud-proxy | Proxy address and port number for overriding the environment's default proxy settings for accessing cloud service. | |
| --cloud-secret-access-key | AWS secret access key. | |
| --cloud-service | Cloud service for data backup or restoration. | |

| Format | Description | Introduced |
|---------------------------------------|---|------------|
| <code>--cloud-trace</code> | Print trace information for cloud operations. | |
| <code>--comments</code> | Specifies comments string. | |
| <code>--comments-file</code> | Specifies path to comments file. | |
| <code>--compress</code> | Create backup in compressed format. | |
| <code>--compress-level</code> | Specifies the level of compression. | |
| <code>--compress-method</code> | Specifies the compression algorithm. | |
| <code>--connect-if-online</code> | Use connection only if available. | |
| <code>--connect_timeout</code> | Connection timeout in seconds. | |
| <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 a backup (or in a restore of a TTS 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</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 a backup (or in a restore of a TTS 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. | |

| Format | Description | Introduced |
|---|--|------------|
| <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. | |
| <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). | 3.11.0 |
| <code>--login-path</code> | Read options from the named login path in the .mylogin.cnf 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. | 3.11.0 |
| <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 .frm 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. | 3.11.0 |
| <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. | 3.11.0 |

| Format | Description | Introduced |
|--|--|------------|
| <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. | |
| <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). | 3.11.0 |
| <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). | 3.11.0 |
| <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. | 3.11.0 |
| <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. | 3.11.0 |
| <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. | |

| Format | Description | Introduced |
|---------------------------------------|---|------------|
| <code>--src-entry</code> | Identifies a file or directory to extract from a single-file backup. | |
| <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</code> | Trace level of messages by <code>mysqlbackup</code> . | |
| <code>--uncompress</code> | Uncompress the compressed backup before an <code>apply-log</code> , <code>copy-back</code> , or <code>copy-back-and-apply-log</code> operation. | |
| <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. | |

5.1.1 Subcommands

These options represent the major operations or modes for the `mysqlbackup` command. Only one can be specified for each `mysqlbackup` invocation, and the name is not preceded by any dashes.

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

The major groups of subcommands are:

- Backup operations: `backup`, `backup-and-apply-log`, `backup-to-image`
- Apply 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`

5.1.1.1 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` command 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]
            backup | backup-and-apply-log

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
```

- `backup`

Performs the initial phase of a backup. The second phase is performed later by running `mysqlbackup` again with the `apply-log` option.

- `backup-and-apply-log`

A combination of `backup` and `apply-log`. Not compatible with incremental backups. Also incompatible with the `--compress` option.

- `backup-to-image`

Produces a single-file backup rather than a directory structure holding the backup files. 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.

Example 5.1 Simple Backup with Connection Parameters from Default Configuration File

The following example shows a minimal backup with the `mysqlbackup` command, with any necessary connection parameters for the database in the `[mysqlbackup]` section of the default MySQL configuration file:

```
mysqlbackup --backup-dir=/export/backups/latest backup
```

Example 5.2 Basic Incremental Backup

```
mysqlbackup --incremental --start-lsn=12345 --incremental-backup-dir=/path/to/incbackup backup
```

There is a separate directory dedicated to incremental backup. Both this directory and the one for full backups can be specified in the `my.cnf` file, and the appropriate directory is used depending on the type of backup. Both the incremental backup data and an earlier full backup are needed to do a successful restore operation.

5.1.1.2 Apply-Log Operations for Existing Backup Data

These operations bring the backup files up-to-date with any changes to InnoDB tables that happened while the backup was in progress. Although for convenience you can combine this operation with the initial backup using the `backup-and-apply-log` option, you must run the steps separately when performing incremental or compressed backups.

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

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

- `apply-log`

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

- `apply-incremental-backup`

Brings the backup up-to-date using the data from an incremental backup.

Example 5.3 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.

Because the `apply-log` operation does not apply to incremental backups, no `incremental-backup-dir` is needed for this operation.

You can also perform `apply-log` and `copy-back` (which restores the prepared backup) together with a single `copy-back-and-apply-log` command.

5.1.1.3 Restore an Existing Backup

Restores the data files from a backup to their original locations within the database server. The MySQL instance must be shut down first before a restore operation. 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 and examples, see [Chapter 4, Recovering or Restoring a Database](#).

```
mysqlbackup [STD-OPTIONS]
            [SERVER-REPOSITORY-OPTIONS]
```

```
    [--backup-dir=PATH]
    [MESSAGE-LOGGING-OPTIONS]
    [PARTIAL-BACKUP-RESTORE-OPTIONS]
    [PROGRESS-REPORT-OPTIONS]
    [CLOUD-STORAGE-OPTIONS]
    copy-back

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
```

- `copy-back`

Restores files from a 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 used to replace all data of an existing MySQL server). See [Section 4.2, “Performing a Restore Operation” \[47\]](#) for details.

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

- `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 option 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 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 single-file incremental backup, besides specifying the location of the incremental backup image with the `--backup-image` option, also supply 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.
- To restore a backup created with the `--use-tts` option:

- See the general requirements described in [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` subcommand, see [Section 4.2, "Performing a Restore Operation"](#).



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.

5.1.1.4 Validating a Backup

To ensure the integrity of the backup data, MySQL Enterprise Backup provides a `validate` subcommand 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:

- 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.
- If any `.ibd` or `.frm` files are missing from a backup directory during a backup or have been deleted from a backup directory after the backup was made, the `validate` operation will not be able to detect the problem.
- If a backup directory 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.

Here is a sample command for validating a backup directory:

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

Here is a sample command for validating a backup image:

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

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
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 : 
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 : 
.....
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.
```

5.1.1.5 Work with Single-File Backups

To simplify transfer and management of backup data, you can keep each backup in a single file (the backup image). The `backup-to-image` option performs a backup directly to a single file, or the options here can pack an existing backup into a single file or unpack a single-file backup to a full backup directory structure. There are other options for working with single-file backups, which are explained below. For usage and examples, see [Section 3.3.5, “Making a Single-File Backup”](#).

```
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

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
```

- `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 in which to unpack. For usage and examples, see [Section 3.3.5, “Making a Single-File Backup”](#).

- `backup-dir-to-image`

Packs an existing backup into a single file. 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. The `--backup-image` parameter is either `-` or an absolute path outside the `backup-dir` directory. For usage and examples, see [Section 3.3.5, “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 and examples, see [Section 3.3.5, “Making a Single-File Backup”](#).

**Note**

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

- `extract`

Unpacks an individual file or directory from a single-file backup. For troubleshooting or restoration operations that do not require the full set of backup data. The resulting file or directory goes in the current directory, or in `backup-dir` if specified. All files and directory contents in the image with absolute path names are extracted into the same absolute path names on the local system. For usage and examples, see [Section 3.3.5, “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 path names even if `--backup-dir` option is specified. The `--dst-entry` option must be used to relocate an absolute pathname.

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

See description for `copy-back-and-apply-log` in [Section 5.1.1.3, “Restore an Existing Backup”](#).

5.1.2 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 op
--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 operations such as overwrite files, create backup directory.
--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 overwriting of InnoDB data and log files during the `apply-log` and `apply-incremental-backup` operations and allows the replacing of an image file during an `backup-to-image` or `backup-dir-to-image` option. 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)

5.1.3 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. The `mysqlbackup` command 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 the `mysqlbackup` command 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.

5.1.4 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 live 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 live. 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, that 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) live 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 filename 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 live 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, that path is located 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.

5.1.5 Backup Repository Options

These options specify various parameters related to the layout of the backup directory.

The `--backup-dir` option is the only one from this group that you typically specify. Values for all the `backup_innodb_*` options can be derived automatically from the corresponding configuration options without the `backup_` prefix (see descriptions for individual options below). Those options need to be specified only if, during a backup, the user wants to specify particular values for those options that are to be used during restorations of the backup. Values specified by `backup_innodb_*` options are stored in the `backup_my.cnf` file, to be used during future restorations.

When a database connection is available during a backup, those `backup_innodb_*` options are ignored, with their values overridden by the corresponding values retrieved from the database connection. .

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`.

The following parameters describe the layout of files in the backup directory:

- `--backup_dir=PATH`

Same as `--backup-dir`. The directory under which to store the backup data. This is a crucial parameter required for most kinds of backup operations. It cannot be a subdirectory of the directory specified by `--datadir`. An additional level of subdirectory is created when the `--with-timestamp` option is also specified.

Also, when restoring a single-file backup with `copy-back-and-apply-log`, `--backup-dir` is used for supplying the location of a folder that will be used for storing temporary files produced during the restoration process.

- `backup_innodb_data_home_dir=PATH`

Specifies the directory where backup InnoDB data files live. Usually same as `backup-dir`, but can be different.

This parameter together with `backup_innodb_data_file_path` determines where the InnoDB data files (such as `ibdata1`, `ibdata2`, ...) are stored inside the backup directory structure.

This parameter is applicable only for backup operations, not restore.

For the backup operations (such as `backup`, `backup-and-apply-log`, `backup-to-image`), the value of the backup destination directory is derived as follows:

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

To make it easy to relocate the backup directory and avoid editing the `backup-my.cnf` file, the backup operation writes this value into `backup-my.cnf` only if it is different than the `backup-dir` value, and using a relative path if possible.

For `backup-to-image` operations, the final value of the `backup_innodb_data_home_dir` option must be a relative path, so that the single-file backup is machine-independent.

- `backup_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
```

This parameter together with `backup_innodb_data_home_dir` determines where the InnoDB data files (such as `ibdata1`, `ibdata2`, ...) live in the backup repository.

Within the backup directory, any data files specified with relative paths are located relative to the `backup_dir` path. Any data files specified with absolute paths are placed inside the `backup_innodb_data_home` directory.

When the parameter is not specified, it inherits the value from the value of the `innodb_data_file_path` option. If both source and destination attempt to use an absolute path that resolve to the same files, the backup is cancelled.

To specify absolute paths for InnoDB datafiles in backup, you must also set the `backup_innodb_data_home` option to `" "`.

- `backup_innodb_log_group_home_dir=PATH`

Specifies where backup InnoDB logs live. Usually the same as `backup-dir`, but can be different.

The names of the log files are fixed and not reconfigurable.

This parameter is applicable only for backup operations (not restore).

The backup operation uses this value and writes it as `innodb_log_group_home_dir=value` in `backup-my.cnf`.

For `copy-back` and `apply-log` operations, `innodb_log_group_home_dir` in `backup-my.cnf` is treated in a way that is compatible with how it was created.

- `backup_innodb_log_files_in_group=N`

Specifies the number of InnoDB log files in backup before being rotated. Example: 5.

Usually same as `innodb_log_files_in_group`, but can be different.

The value for this parameter is derived as:

- Specified `backup_innodb_log_files_in_group` value from command line or configuration file.
- Else `innodb_log_files_in_group` value from the database connection, if available.
- Else the `innodb_log_files_in_group` value from the command line or configuration file.
- `backup_innodb_log_file_size=SIZE`

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

Usually the same as `innodb_log_file_size`, but can be different.

The value for this parameter is derived as:

- Specified `backup_innodb_log_file_size` value from command line or configuration file.
- Else `innodb_log_file_size` value from database connection, if available.
- Else specified `innodb_log_file_size` value from command line or configuration file.
- `backup_innodb_page_size=SIZE`

The page size for all InnoDB tablespaces in a MySQL instance. The page size must be the same value as in MySQL instance. By default, it assumes the innodb page size in server.

- `backup_innodb_checksum_algorithm=NAME`

The name of the checksum algorithm used for validating InnoDB tablespaces. Default is 'innodb'.

- `backup_innodb_undo_directory=PATH`

The relative or absolute directory path where InnoDB creates separate tablespaces for the undo logs. Typically used to place those logs on a different storage device.

- `backup_innodb_undo_logs=NUMBER`

Number of rollback segments in the system tablespace that InnoDB uses within a transaction. This setting is appropriate for tuning performance if you observe mutex contention related to the undo logs.

- `backup_innodb_undo_tablespaces=NUMBER`

The number of tablespace files that the undo logs are divided between, when you use a non-zero `innodb_undo_logs` setting. 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`.

- `--with-timestamp`

Creates a subdirectory underneath the backup directory, with a name formed from the timestamp of the backup operation. Useful to maintain a single backup directory containing many backup snapshots.

Default: no timestamped subdirectory is created. To reuse the same backup directory for a new backup, either remove the previous backup files manually or specify the `--force` option to overwrite them.

5.1.6 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 11.3, “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`

| | | |
|----------------------------|--------------------------------|---------------------|
| Command-Line Format | <code>--comments=STRING</code> | |
| Permitted Values | Type | <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`

| | | |
|----------------------------|-----------------------------------|------------------------|
| Command-Line Format | <code>--comments-file=PATH</code> | |
| Permitted Values | Type | <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.

5.1.7 Compression Options

For an overview on backup compression, see [Section 3.3.3, “Making a Compressed Backup”](#).

- `--compress`

Create backup in compressed format. For a regular backup, only the InnoDB data files are compressed, and they bear the `.ibz` extension after the compression. Similarly, for a single-image backup, only the InnoDB data files inside the backup image are compressed.

Default: compression is disabled.

- `--compress-method=ALGORITHM`

| | | |
|----------------------------|--|--------------------------|
| Command-Line Format | <code>--compress-method=ALGORITHM</code> | |
| Permitted Values | Type | <code>enumeration</code> |
| | Default | <code>lz4</code> |
| | Valid Values | <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.

5.1.8 Incremental Backup Options

For an overview of incremental backups and usage information about these options, see [Section 3.3.2, “Making an Incremental Backup”](#).

To take an incremental backup, specify the `--incremental` or `--incremental-with-redo-log-only`, along with the `--incremental-backup-dir`. All InnoDB data modified after the specified LSN is copied in the incremental backup. Depending on the choice of `--incremental` or `--incremental-with-redo-log-only` 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.

The incremental aspect applies only to InnoDB tables. By default, all non-InnoDB and `.frm` files are also included in incremental backup. To exclude non-InnoDB data in an incremental backup, use the `--only-innodb` option.

- `--incremental-with-redo-log-only`

Specifies an alternative form of [incremental](#) backup for a `backup` or `backup-to-image` operation. Also requires either the `--incremental-base` option, or the combination of the `--start-lsn` and `--incremental-backup-dir` options.

The incremental backup performed by this option has different performance characteristics and operational limitations than with the `--incremental` option:

- 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 that depends on 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 this case, the `mysqlbackup` command quickly determines it cannot proceed and returns an error. Your backup script can catch the error and do 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 `innodb_log_files_in_group`. To compute redo log size at the physical level, look in the `datadir` directory of the MySQL instance and sum 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 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 described 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.

- This type of incremental backup is not so forgiving of too-low `--start-lsn` values as the standard `--incremental` option. 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 using this option, Oracle recommends always using 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 see 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, to avoid the likelihood of the backup failing due to historical data 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 if a sudden flurry of updates produced more redo 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, amount of DML activity, and size of your redo log files; do your testing on a server with a realistic data volume and running 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.

As with the `--incremental` option, the incremental aspect applies only to InnoDB tables. By default, all non-InnoDB and `.frm` files are also included in incremental backup. To exclude non-InnoDB data in an incremental backup, use the `--only-innodb` option.

- `--incremental-base=mode:argument`

| | | |
|----------------------------|---|---------------------|
| Command-Line Format | <code>--incremental-base=mode:argument</code> | |
| Permitted Values | Type | <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 root 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 in a subsequent incremental backup using this mode of the `--incremental-base` option.

- `--start-lsn=LSN`

| | | |
|----------------------------|------------------------------|----------------------|
| Command-Line Format | <code>--start-lsn=LSN</code> | |
| Permitted Values | Type | <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`

| | | |
|----------------------------|--|-----------------------------|
| Command-Line Format | <code>--incremental-backup-dir=PATH</code> | |
| Permitted Values | Type | <code>directory name</code> |

Specifies the location under which to store data from an incremental backup. This is the same location you specify with `--incremental-base` if you use that option for a subsequent incremental backup.

Example 5.4 Incremental Backup

These examples show typical combinations of options used for incremental backups.

```
mysqlbackup --incremental \
--incremental-backup-dir=/var/mysql/backup/latest \
--incremental-base=dir:/var/mysql/backup/previous \
... backup

mysqlbackup --incremental-with-redo-log-only \
--incremental-backup-dir=/var/mysql/backup/latest \
--incremental-base=dir:/var/mysql/backup/previous \
... backup

mysqlbackup --incremental --start-lsn=12345 \
--incremental-backup-dir=/var/mysql/backup/inc \
... backup

mysqlbackup --incremental-with-redo-log-only --start-lsn=12345 \
--incremental-backup-dir=/var/mysql/backup/inc \
... backup
```

5.1.9 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 information on the following options for partial backup, see [Section 3.3.4, “Making a Partial Backup”](#).

- `--include-tables=REGEXP`

| | | |
|----------------------------|--------------------------------------|---------------------|
| Command-Line Format | <code>--include-tables=REGEXP</code> | |
| Permitted Values | Type | <code>string</code> |

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. Some limitations apply when using the option to select database names or file names that contain special characters (spaces, dashes, periods, etc.); see [this description](#) in [Appendix A, MySQL Enterprise Backup Limitations](#) for details. `mysqlbackup` throws an error when the option is used without a regular expression being supplied with it.

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 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`

| | | |
|----------------------------|--------------------------------------|---------------------|
| Command-Line Format | <code>--exclude-tables=REGEXP</code> | |
| Permitted Values | Type | <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. Some limitations apply when using the option to select database names or file names that contain special characters (spaces, dashes, periods, etc.); see [this description](#) in [Appendix A, MySQL Enterprise Backup Limitations](#) for details. `mysqlbackup` throws an error when the option is used without a regular expression being supplied with it.

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 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 [Section 1.4, “Files that Are Backed Up”](#) 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.

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 to select 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](#))
- 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 be used for incremental backups
- Does not include the binary log or the relay log in the backup

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: back up 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 3.1.2, “Grant MySQL Privileges to Backup Administrator”](#) for details.

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

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 information about 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](#) 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 whitespaces surrounding the database or table names, as the whitespaces 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 the `mysqlbackup` command 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 4.2, “Performing a Restore Operation”](#) for details.

5.1.10 Single-File Backup Options

These options are associated with single-file backups. You use them in combination with the `mysqlbackup` subcommands `backup-to-image`, `image-to-backup-dir`, `backup-dir-to-image`, `list-image`, and `extract` that pack or unpack single-image backups. For usage information, see [Section 3.3.5, “Making a Single-File Backup”](#).

- `--backup-image=IMAGE`

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

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

You can optionally prefix the image name with `file:` to signify file I/O (the default). For tape backups, prefix the image name with `sbt:`. See [Section 3.3.5.2, “Backing Up to Tape”](#) for details about tape backups.

- `--src-entry=PATH`

| | | |
|----------------------------|-------------------------------|------------------------|
| Command-Line Format | <code>--src-entry=PATH</code> | |
| Permitted Values | Type | <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`

| | | |
|----------------------------|-------------------------------|------------------------|
| Command-Line Format | <code>--dst-entry=PATH</code> | |
| Permitted Values | Type | <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 selected entry in backup image corresponding to entry specified by `--src-entry=PATH` option. 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`

| | | |
|----------------------------|---------------------------------------|---------------------|
| Command-Line Format | <code>--sbt-database-name=NAME</code> | |
| Permitted Values | Type | <code>string</code> |
| | Default | <code>MySQL</code> |

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 3.3.5.2, “Backing Up to Tape”](#) for usage details.

- `--sbt-lib-path=PATH`

| | | |
|----------------------------|----------------------------------|------------------------|
| Command-Line Format | <code>--sbt-lib-path=PATH</code> | |
| Permitted Values | Type | <code>file name</code> |

Path name of the SBT library used by 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 3.3.5.2, “Backing Up to Tape”](#) for usage details.

- `--sbt-environment=VAR=value,...`

| | | |
|----------------------------|--|---------------------|
| Command-Line Format | <code>--sbt-environment=VAR1=value1[,VAR2=value2[,...]] SBT API provider)</code> | |
| Permitted Values | Type | <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.

5.1.11 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`

| Command-Line Format | <code>--number-of-buffers=NUMBER</code> | |
|---------------------|---|---------|
| Permitted Values | Type | numeric |
| | Default | 14 |
| | Min Value | 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`

| | | |
|----------------------------|------------------------------------|---------|
| Command-Line Format | <code>--read-threads=NUMBER</code> | |
| Permitted Values | Type | numeric |
| | Default | 1 |
| | Min Value | 1 |
| | Max Value | 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`

| | | |
|----------------------------|---------------------------------------|---------|
| Command-Line Format | <code>--process-threads=NUMBER</code> | |
| Permitted Values | Type | numeric |
| | Default | 6 |
| | Min Value | 1 |
| | Max Value | 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`

| | |
|----------------------------|-------------------------------------|
| Command-Line Format | <code>--write-threads=NUMBER</code> |
|----------------------------|-------------------------------------|

| | | |
|-------------------------|------------------|---------|
| Permitted Values | Type | numeric |
| | Default | 1 |
| | Min Value | 1 |
| | Max Value | 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`

| | | |
|----------------------------|--------------------------------|--|
| Command-Line Format | <code>--limit-memory=MB</code> | |
| Permitted Values | Type | numeric |
| | Default | 100 for <code>apply-log</code> (without uncompression), 300 for other operations |
| | Min Value | 0 |
| | Max Value | 999999 |

Specify maximum memory in megabytes that can be used by the `mysqlbackup` command. 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`

| | | |
|----------------------------|-------------------------|---------|
| Command-Line Format | <code>--sleep=MS</code> | |
| Permitted Values | Type | numeric |
| | Default | 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`

| | | |
|----------------------------|------------------------------------|----------------------|
| Command-Line Format | <code>--page-reread-time=MS</code> | |
| Permitted Values | Type | <code>numeric</code> |
| | Default | <code>100</code> |

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`

| | | |
|----------------------------|---|----------------------|
| Command-Line Format | <code>--page-reread-count=number</code> | |
| Permitted Values | Type | <code>numeric</code> |
| | Default | <code>500</code> |

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}`

| | | |
|----------------------------|------------------------------------|-------------------------------|
| Command-Line Format | <code>--on-disk-full=option</code> | |
| Permitted Values | Type | <code>enumeration</code> |
| | Default | <code>abort</code> |
| | Valid Values | <code>abort</code> |
| | | <code>warn</code> |
| | | <code>abort_and_remove</code> |

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`

Do not include binary log files in a backup. Binary log files are included by default for all kinds of online backups (full, incremental, compressed, partial, single-file, etc.). See [Section 1.4, “Files that Are Backed Up”](#), for details. Use this option to skip backing up binary logs if resource, performance, or other issues arise.



Note

Due to some known issues, users should always use the `--skip-binlog` option when creating offline backups, or when creating an incremental backup that is based on a full backup created with the `--no-locking` option. See [Appendix A, MySQL Enterprise Backup Limitations](#) for details.

- `--skip-relaylog`

Do not include relay log files in a backup. Relay log files are included by default for all kinds of online backups (full, incremental, compressed, partial, single-file, etc.) of a slave server. See [Section 1.4, “Files that Are Backed Up”](#), 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.

- `--log-bin-index[=PATH]`

| | | |
|----------------------------|---------------------------------------|---|
| Command-Line Format | <code>--log-bin-index=FILENAME</code> | |
| Permitted Values | Type | <code>file name</code> |
| | Default | <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 filename 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]`

| | | |
|----------------------------|---|---|
| Command-Line Format | <code>--relay-log-index=FILENAME</code> | |
| Permitted Values | Type | <code>file name</code> |
| | Default | <code>data_dir/host_name-relay-bin.index</code> |

For offline backups of slave servers only: specify the absolute path (including filename 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]`

| | | |
|----------------------------|--|-----------------------------------|
| Command-Line Format | <code>--master-info-file=FILENAME</code> | |
| Permitted Values | Type | <code>file name</code> |
| | Default | <code>data_dir/master.info</code> |

For offline backups of slave servers only: specify the absolute path (including filename 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]`

| | | |
|----------------------------|--|--------------------------------------|
| Command-Line Format | <code>--relaylog-info-file=FILENAME</code> | |
| Permitted Values | Type | <code>file name</code> |
| | Default | <code>data_dir/relay-log.info</code> |

For offline backups of slave servers only: specify the absolute path (including filename 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]`

| | | |
|----------------------------|--|---------------------|
| Command-Line Format | <code>--optimistic-time=DATE-TIME</code> | |
| Permitted Values | Type | <code>string</code> |
| | Default | <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 3.3.6, “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`

| | | |
|----------------------------|--|---------------------|
| Command-Line Format | <code>--optimistic-busy-tables=REGEXP</code> | |
| Permitted Values | Type | <code>string</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, 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 3.3.6, “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`.

5.1.12 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 | <code>directory name</code> |
| | 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.



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
```

5.1.13 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 | <code>stderr</code> |
| | | <code>stdout</code> |
| | | <code>file:FILENAME</code> |
| | | <code>fifo:FIFONAME</code> |
| | | <code>table</code> |
| | | <code>variable</code> |

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 `meb_output`:

```
mysqlbackup --show-progress --show-progress=file:meb_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 | <code>numeric</code> |
| | Default | <code>2</code> |
| | Min Value | <code>1</code> |
| | Max Value | <code>100000</code> |

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.

5.1.14 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 information on 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` subcommand. It cannot be used with the `backup` or `backup-and-apply-log` subcommand.

- `--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` subcommand, the encrypted backup image has to be unpacked and decrypted first using the `image-to-backup-dir` or `extract` subcommand, together with the `--decrypt` option.

- `--key=STRING`

| | | |
|----------------------------|------------------------|---------------------|
| Command-Line Format | <code>--key=KEY</code> | |
| Permitted Values | Type | <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`

| | | |
|----------------------------|------------------------------|------------------------|
| Command-Line Format | <code>--key-file=FILE</code> | |
| Permitted Values | Type | <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.

5.1.15 Cloud Storage Options

These options are for creating or restoring a single-file backup to or from cloud storage. See [Section 3.3.5.3, “Backing Up to Cloud Storage”](#) for more information and instructions on using cloud storage with MySQL Enterprise Backup.

- `--cloud-service`

Cloud service for data backup or restoration. Currently, only the Amazon S3 service is supported, and “s3” is the only value `mysqlbackup` accepts for this option.

- `--cloud-bucket`

The storage bucket on Amazon S3 for the backup image. The option only has meaning if Amazon S3 is used for cloud backup.

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`

The Amazon S3 object key for the backup image. The option only has meaning if Amazon S3 is used for cloud backup.

- `--cloud-access-key-id`

AWS access key ID for logging onto Amazon S3. The option only has meaning if Amazon S3 is used for cloud backup.

- `--cloud-secret-access-key`

AWS secret access key that goes with the AWS access key id used in `--cloud-access-key-id`. The option only has meaning if Amazon S3 is used for cloud backup.

- `--cloud-aws-region`

Region for Amazon Web Services that `mysqlbackup` accesses for S3. The option only has meaning if Amazon S3 is used for cloud backup.

- `--cloud-trace`

Print trace information for cloud operations. It works independently of `--trace`, 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 Amazon S3.

**Note**

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

5.1.16 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 InnoDB tables.

- `--exec-when-locked="utility arg1 arg2 ..."`

| Command-Line Format | <code>--exec-when-locked="utility arg1 arg2 ..."</code> | |
|---------------------|---|---------------------|
| Permitted Values | Type | <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.

5.2 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 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 parameter 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 documentation typically lists the names with underscores, to match the output of the `SHOW VARIABLES` statement.

Options Files

The `mysqlbackup` command reads the location of the MySQL data to back up from (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.

Because `mysqlbackup` **does not overwrite any files** during the initial backup step, the backup directory must not contain any old backup files. `mysqlbackup` stops when asked to create a file that already exists, to avoid harming an existing backup. For convenience, specify the `--with-timestamp` option, which always creates a unique timestamped subdirectory for each backup job underneath the main backup directory.

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 5.5 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 minimal required parameters are stored in `backup-my.cnf`, to allow you to restore the backup to a different location without extensive changes to that file. For example, although the `innodb_data_home_dir` and `innodb_log_group_home_dir` options can go into `backup-my.cnf`, they are omitted when those values are the same as the `backup-dir` value.

Chapter 6 Using MySQL Enterprise Backup with Replication

Table of Contents

| | |
|---|-----|
| 6.1 Setting Up a New Replication Slave | 101 |
| 6.2 Backing up and Restoring a Slave Database | 102 |
| 6.3 Restoring a Master Database | 102 |

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 to deal with entire system images to set up new slave servers, or to restore a master server in an efficient way that avoids unnecessary work for the slave servers. 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 to a point in time, even a time that is later than the last backup.

GTID support with MySQL Server 5.6 and above

MySQL Enterprise Backup supports the [GTID feature](#) of MySQL 5.6:

- The GTID feature is compatible with the CSV tables that the `mysqlbackup` command uses to log job progress and history.
- When a server using the GTID feature is backed up, `mysqlbackup` produces a file `gtid_executed.sql` as part of the backup data. This file contains a SQL statement that sets the `GTID_PURGED` configuration option. Execute this file using the `mysql` command line after restoring the backup data on a slave server. Optionally, you can uncomment the `CHANGE MASTER` command in this file and add any needed authentication parameters to it, before running it through `mysql`.
- For servers not using GTIDs, you can use the `--slave-info` option as before, then edit and execute the `ibbackup_slave_info` file afterward.

6.1 Setting Up a New Replication Slave

If you use MySQL replication, MySQL Enterprise Backup allows you to set up a slave database without stopping the master, by backing up the master and restoring that backup on a new slave server.

1. Take the backup, transfer it to the slave server, use `mysqlbackup` with the `apply-log` option to prepare it, and put the restored backup and the log files in the right directories for the new slave.
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` 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.

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. Using the backup of the master database, do the `apply-log` operation, shut down the database, and do the `copy-back` operation.
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 | 105 |
| 7.2 Optimizing Restore Performance | 108 |

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 5.1.8, “Incremental Backup Options”](#) for information about these options. For usage instructions for the backup and apply stages of incremental backups, see [Section 3.3.2, “Making an Incremental Backup”](#) and [Example 4.3, “Applying an Incremental Backup to a Full 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 3.3.3, “Making a Compressed Backup”](#) and [Section 5.1.7, “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 5.1.1.5, “Work with Single-File Backups”](#) for the options related to single-file backups, and [Section 3.3.5, “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

The `mysqlbackup` command can take advantage of modern multicore CPUs and operating system threads to perform backup operations in parallel. See [Section 5.1.11, “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 the `mysqlbackup` command 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 the `mysqlbackup` command returned exit code 0, or by observing that `mysqlbackup` has printed the text “mysqlbackup completed OK!”.
- The `mysqlbackup` command 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 A.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 5.1.9, “Partial Backup and Restore Options”](#) for information about the relevant options, and [Section 3.3.4, “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 3.3.4, “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.

The Apply-Log Phase

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. [Section 5.1.1.2, “Apply-Log Operations for Existing Backup Data”](#)

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 [Section 5.1.1.2, “Apply-Log Operations for Existing Backup Data”](#).

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 5.1.1.3, “Restore an Existing Backup”](#).

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 5.1.9, “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. 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. See [Section 5.1.7, “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 5.1.1.5, “Work with Single-File Backups”](#) for information about making single-file backups.

The Apply-Log Phase

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. [Section 5.1.1.2, “Apply-Log Operations for Existing Backup Data”](#)

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 [Section 5.1.1.2, “Apply-Log Operations for Existing Backup Data”](#).

[Section 5.1.1.2, “Apply-Log Operations for Existing Backup Data”](#)

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

The `mysqlbackup` command can take advantage of modern multicore CPUs and operating system threads to perform backup operations in parallel. See [Section 5.1.11, “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, since version 3.10, 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
6ald325e6ef0577f3400b7cd624ae574f5186d0da2eeb946895de418297ed75b
```

The encryption function uses MySQL Enterprise Backup's own encryption format—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 example, 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 5.1.14, “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 | 113 |
|--|-----|

This section describes how you can use MySQL Enterprise Backup in combination with media management software (MMS) products. Such products are typically used for managing large volumes of backup data, often with high-capacity backup devices such as tape drives.

9.1 Backing Up to Tape with Oracle Secure Backup

Tape drives are affordable, high-capacity storage devices for backup data. The MySQL Enterprise Backup product 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.

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 the `mysqlbackup` command to control various aspects of the SBT processing. The `--sbt-*` options are listed in [Section 5.1.10, “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 the `mysqlbackup` command 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 the `mysqlbackup` command 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.

To restore MySQL data from tape:

- Specify the `--backup-image=sbt:name` parameter of the `mysqlbackup` command as part of the restore operation. Use the same `name` value as during the original backup. This single parameter retrieves the appropriate data from the appropriate tape device.

- Optionally use the `--sbt-lib-path` option, using the same values 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 in standard places):
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 Monitoring Backups with MySQL Enterprise Monitor

The MySQL Enterprise Monitor is a companion product to the MySQL Server that enables monitoring of MySQL instances and their hosts, notification of potential issues and problems, and advice on how to correct issues. Among its other functions, it can be used to monitor the progress and history of backup jobs. Check the [MySQL Enterprise Monitor User's Guide](#) for detail.

Chapter 11 Troubleshooting for MySQL Enterprise Backup

Table of Contents

| | |
|---|-----|
| 11.1 Error codes of MySQL Enterprise Backup | 117 |
| 11.2 Working Around Corruption Problems | 117 |
| 11.3 Using the MySQL Enterprise Backup Logs | 118 |
| 11.4 Using the MySQL Enterprise Backup Manifest | 119 |

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 A, MySQL Enterprise Backup Limitations](#).
- If the `mysqlbackup` command encounters problems during operating system calls, it returns the corresponding OS error codes. You might need to consult your operating system documentation for the meaning and solution of these error codes.
- The output from the `mysqlbackup` command 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 5.1.12, “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 the `mysqlbackup` command 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 these tables to monitor ongoing jobs, see how much time was needed for various stages, and check if any errors occurred.

11.1 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.

11.2 Working Around Corruption Problems

Sometimes the operating system or the hardware can corrupt a data file page, in 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` disables 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.

11.3 Using the MySQL Enterprise Backup Logs

The `mysql.backup_progress` table lets you monitor backup jobs as they run. 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 the `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 specify the end LSN from that backup as the starting LSN for the next incremental backup.

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.

11.4 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 substantial 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. The contents of this file might be expanded in the future. A management tool might use this information to confirm which tables are part of a full backup, or a partial backup performed with the `--databases` option. (The information is not present for partial backups taken with the `--include`, `--incremental`, `--incremental-with-redo-log-only`, `--only-innodb`, or `--only-innodb-with-frm` options.) 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` options.) 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` options, or to remap the paths of extracted files with the `--src-entry` and `--dst-entry` options.

Chapter 12 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

- **12.1:** Does MySQL Enterprise Backup work with MySQL Server version `x.y.z`?
- **12.2:** What is the big `ibdata` file that is in all the backups?
- **12.3:** Can I back up non-InnoDB data with MySQL Enterprise Backup?
- **12.4:** What happens if “apply” step is interrupted?
- **12.5:** Why is the option `--defaults-file` not recognized?
- **12.6:** Can I back up a database on one OS platform and restore it on another one using MySQL Enterprise Backup?

Questions and Answers

12.1: Does MySQL Enterprise Backup work with MySQL Server version `x.y.z`?

See [Section B.2, “Compatibility with MySQL Versions”](#) for details of compatibility between different releases of MySQL Enterprise Backup and MySQL Server.

12.2: 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, 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, which are more flexible in terms of freeing disk space when dropped or truncated, and can be individually included or excluded from backups.
- 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 dump the entire instance, then turn on `innodb_file_per_table` before re-creating it, so that all the table data is kept outside the system tablespace.

12.3: 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.

12.4: What happens if “apply” step is interrupted?

If the `mysqlbackup` command 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.

12.5: Why is the option `--defaults-file` not recognized?

When you specify the `--defaults-file` option, it must be the first option after the name of the command. Otherwise, the error message makes it look as if the option name is not recognized.

12.6: Can I back up a database on one OS platform and restore it on another one using MySQL Enterprise Backup?

See [Section B.1, “Cross-Platform Compatibility”](#) for details.

Part III Appendixes

Table of Contents

| | |
|--|-----|
| A MySQL Enterprise Backup Limitations | 127 |
| A.1 Limitations of MySQL Enterprise Backup | 127 |
| B Compatibility Information for MySQL Enterprise Backup | 131 |
| B.1 Cross-Platform Compatibility | 131 |
| B.2 Compatibility with MySQL Versions | 131 |
| B.3 Compatibility with Older Versions of MySQL Enterprise Backup | 131 |
| B.4 Compatibility Notes for Specific MySQL Versions | 131 |
| C Extended Examples | 133 |
| C.1 Sample Directory Structure for Full Backup | 133 |
| C.2 Sample Directory Structure for Compressed Backup | 137 |
| C.3 Sample Directory Structure for Incremental Backup | 137 |
| D MySQL Enterprise Backup Release Notes | 139 |
| MySQL Enterprise Backup Glossary | 141 |

Appendix A MySQL Enterprise Backup Limitations

Table of Contents

| | |
|--|-----|
| A.1 Limitations of MySQL Enterprise Backup | 127 |
|--|-----|

Please refer to the MySQL Enterprise Backup version history in [Appendix D, MySQL Enterprise Backup Release Notes](#) for a list of fixed `mysqlbackup` bugs.

A.1 Limitations of MySQL Enterprise Backup

- The group commit feature in MySQL 5.6 and higher changes the frequency of flush operations for the `InnoDB` redo log, which could affect the point in time associated with the backup data from `InnoDB` tables. See [Section B.4, “Compatibility Notes for Specific MySQL Versions”](#) for details.
- For MySQL 5.5 and earlier, when restoring an individual `InnoDB` table, as described in [Section 4.4, “Backing Up and Restoring a Single .ibd File”](#), the table must not have been dropped or truncated in the MySQL server after the backup. Dropping or truncating an `InnoDB` table changes its internal table ID, and when the table is re-created the ID will not match the table ID from the backup data. This restriction does not apply to MySQL 5.6 and later, as long as the restoration is made from one Generally Available (GA) version to another in the same series of MySQL servers.
- In Linux, Unix, and OS X systems, the `mysqlbackup` command 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` 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, such as by 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 [Chapter 5, mysqlbackup Command Reference](#) for guidelines about backup timing and performance.
- Do not run the DDL operations `ALTER TABLE`, `TRUNCATE TABLE`, `OPTIMIZE TABLE`, `REPAIR TABLE`, or `RESTORE TABLE` while a backup operation is going on. The resulting backup might be 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 you use those temporary tables to update or insert into normal tables, 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.
- When saving any tables into the data directory, the MySQL server performs conversions for any special characters in any database or table names to come up with the file names for storage. Here are two examples of MySQL commands including special characters (“.” and “-”) in their database and file names:

```
mysql> CREATE TABLE `db.2`.`t.2` (c1 INT) ENGINE=InnoDB;
mysql> CREATE TABLE `db-3`.`t-3` (c1 INT) ENGINE=InnoDB;
```

The tables created by these commands are stored in the file system as these files:

```
db@002e2/t@002e2.ibd
db@002d3/t@002d3.ibd
```

Currently, `mysqlbackup` does not perform the same conversion with database and table names supplied using the `--include-tables` and `--exclude-tables` options. Using, for example, `--include-tables="db\2\2|db-3\t-3"` will not select the tables described above.

As a workaround, for a backup that does not use [transportable tablespace \(TTS\)](#) (that is, the `--use-tts` option is not used during the creation of the backup), users can look up the file names for the desired tables in the data directories of the databases and, in the arguments for the `--include-tables` and `--exclude-tables` options, use the converted names for the database directories as database names and the converted file names (without the file extension) as table names. For example, for the tables mentioned above, select them for inclusion in a backup with the option argument `--include-tables="db@002e2\t@002e2|db@002d3\t@002d3"`. Using the same argument with the option `--exclude-tables` excludes the two tables in the backup.

However, the suggested workaround does not work for TTS backups; therefore, do not use TTS if you want to do a partial backup that includes tables with special characters in their names.

- 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 3.3.6, “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 bug with the InnoDB storage engine (see Bug# 72851 on the MySQL Bug System).
- The attempt to copy the binary log files from a server into a backup (which is the default behavior of MySQL Enterprise Backup since 3.11.0) will cause an incremental backup to fail if the full backup it is based on was created with the `--no-locking` option. As a workaround, users should use the `--skip-binlog` option for an incremental backup performed in that situation.
- The attempt to copy the binary log files from a server into a backup (which is the default behavior of MySQL Enterprise Backup since 3.11.0) will cause an offline backup to fail. As a workaround, always use the `--skip-binlog` option for offline backups.

- 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 (Bug #20583946).
- Tables created on the MySQL server with the `ANSI_QUOTES` SQL mode cannot be backed up using transportable tablespace (TTS).

Appendix B Compatibility Information for MySQL Enterprise Backup

Table of Contents

| | |
|--|-----|
| B.1 Cross-Platform Compatibility | 131 |
| B.2 Compatibility with MySQL Versions | 131 |
| B.3 Compatibility with Older Versions of MySQL Enterprise Backup | 131 |
| B.4 Compatibility Notes for Specific MySQL Versions | 131 |

This section describes information related to compatibility issues for MySQL Enterprise Backup releases.

B.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](#).

B.2 Compatibility with MySQL Versions

MySQL Enterprise Backup 3.11 supports all the MySQL server versions that are currently in the General Availability (GA) status (that is, MySQL 5.5 and 5.6).

B.3 Compatibility with Older Versions of MySQL Enterprise Backup

MySQL Enterprise Backup 3.11 can be used to restore the following kinds of backups created by earlier versions of the product:

- Backups created for MySQL 5.5 by MySQL Enterprise Backup 3.5 to 3.10.
- Backups created for MySQL 5.6 by MySQL Enterprise Backup 3.8 to 3.10.



Note

- For any restore that involves a server upgrade or downgrade, see the important discussion in [Section 4.5, “Restoring a Backup with a Database Upgrade or Downgrade”](#)

B.4 Compatibility Notes for Specific MySQL Versions

This section lists any performance-related features and settings in specific MySQL Server versions that affect various aspects of the backup process.

MySQL 5.6

Some new MySQL 5.6 features introduce changes in directory layout and file contents for InnoDB tables. Backing up servers that use these features requires MySQL Enterprise Backup 3.8.1 or higher:

- `innodb_page_size` configuration option.

- `innodb_undo_directory`, `innodb_undo_logs`, and `innodb_undo_tablespaces` configuration options.
- `innodb_checksum_algorithm` configuration option.
- `DATA DIRECTORY` clause of the `CREATE TABLE` statement, which produces a `.isl` file in the database directory and stores the `.ibd` file in a user-specified location.
- [Online DDL](#).

See [MySQL Enterprise Backup 3.9 Release Notes](#) for details on the fixes and enhancements related to these MySQL 5.6 features.

Appendix C Extended Examples

Table of Contents

| | |
|---|-----|
| C.1 Sample Directory Structure for Full Backup | 133 |
| C.2 Sample Directory Structure for Compressed Backup | 137 |
| C.3 Sample Directory Structure for Incremental Backup | 137 |

This section illustrates the commands and associated output for various backup and restore operations.

C.1 Sample Directory Structure for Full Backup

Here is an example of the subdirectories and files underneath a typical backup directory. The `--with-timestamp` option creates a new subdirectory for each backup, named according to the timestamp of the job. The backups contain the files for the InnoDB system tablespace, `.ibd`, `.frm`, `.MYD`, `.MYI`, `.CSV`, and `.CSM` files representing table and index data from various storage engines, and `.par` and `#P#` files representing partitioned tables.

```
$ find ~/backups
/Users/cirrus/backups
/Users/cirrus/backups/2011-06-16_10-33-47
/Users/cirrus/backups/2011-06-16_10-33-47/backup-my.cnf
/Users/cirrus/backups/2011-06-16_10-33-47/datadir
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/ib_logfile0
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/ib_logfile1
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/ibbackup_logfile
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/ibdata1
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/db.opt
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_double#P#p0.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_double#P#p1.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_double#P#p2.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_double#P#p3.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_double#P#p4.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_double#P#p5.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_double#P#p6.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_double#P#p7.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_double#P#p8.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_double.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_double.par
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_long#P#p0.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_long#P#p1.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_long#P#p2.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_long#P#p3.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_long#P#p4.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_long#P#p5.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_long#P#p6.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_long#P#p7.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_long#P#p8.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_long.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_long.par
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_string#P#p0.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_string#P#p1.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_string#P#p2.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_string#P#p3.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_string#P#p4.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_string#P#p5.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_string#P#p6.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_string#P#p7.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_string#P#p8.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_string.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/dc_p_string.par
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/graph_dc_schedules.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/graph_dc_schedules.ibd
```

Sample Directory Structure for Full Backup

[illegible]

Sample Directory Structure for Full Backup

```
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/mos_service_requests.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/resource_bundle.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/resource_bundle.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/resource_bundle_map.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/resource_bundle_map.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/rule_alarms.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/rule_alarms.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/rule_dc_schedules.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/rule_dc_schedules.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/rule_eval_result_vars.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/rule_eval_result_vars.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/rule_eval_results.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/rule_eval_results.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/rule_schedule_email_targets.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/rule_schedule_email_targets.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/rule_schedules.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/rule_schedules.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/rule_tags.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/rule_tags.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/rule_thresholds.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/rule_thresholds.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/rule_variables.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/rule_variables.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/rules.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/rules.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/schema_version_v2.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/schema_version_v2.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/statement_data.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/statement_data.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/statement_examples.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/statement_examples.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/statement_explain_data.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/statement_explain_data.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/statement_summaries.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/statement_summaries.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/statement_summary_data.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/statement_summary_data.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/system_maps.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/system_maps.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/tags.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/tags.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/target_email.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/target_email.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/user_form_defaults.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/user_form_defaults.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/user_preferences.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/user_preferences.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/user_tags.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/user_tags.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/users.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/users.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/whats_new_entries.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mem/whats_new_entries.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/backup_history.CSM
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/backup_history.CSV
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/backup_history.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/backup_progress.CSM
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/backup_progress.CSV
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/backup_progress.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/columns_priv.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/columns_priv.MYD
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/columns_priv.MYI
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/db.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/db.MYD
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/db.MYI
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/event.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/event.MYD
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/event.MYI
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/func.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/func.MYD
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/func.MYI
```

Sample Directory Structure for Full Backup

```
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/general_log.CSM
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/general_log.CSV
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/general_log.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/help_category.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/help_category.MYD
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/help_category.MYI
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/help_keyword.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/help_keyword.MYD
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/help_keyword.MYI
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/help_relation.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/help_relation.MYD
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/help_relation.MYI
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/help_topic.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/help_topic.MYD
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/help_topic.MYI
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/host.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/host.MYD
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/host.MYI
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/ibbackup_binlog_marker.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/ibbackup_binlog_marker.ibd
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/inventory.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/inventory.MYD
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/inventory.MYI
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/ndb_binlog_index.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/ndb_binlog_index.MYD
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/ndb_binlog_index.MYI
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/plugin.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/plugin.MYD
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/plugin.MYI
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/proc.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/proc.MYD
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/proc.MYI
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/procs_priv.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/procs_priv.MYD
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/procs_priv.MYI
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/servers.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/servers.MYD
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/servers.MYI
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/slow_log.CSM
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/slow_log.CSV
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/slow_log.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/tables_priv.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/tables_priv.MYD
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/tables_priv.MYI
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/time_zone.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/time_zone.MYD
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/time_zone.MYI
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/time_zone_leap_second.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/time_zone_leap_second.MYD
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/time_zone_leap_second.MYI
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/time_zone_name.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/time_zone_name.MYD
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/time_zone_name.MYI
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/time_zone_transition.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/time_zone_transition.MYD
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/time_zone_transition.MYI
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/time_zone_transition_type.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/time_zone_transition_type.MYD
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/time_zone_transition_type.MYI
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/user.frm
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/user.MYD
/Users/cirrus/backups/2011-06-16_10-33-47/datadir/mysql/user.MYI
/Users/cirrus/backups/2011-06-16_10-33-47/meta
/Users/cirrus/backups/2011-06-16_10-33-47/meta/backup_content.xml
/Users/cirrus/backups/2011-06-16_10-33-47/meta/backup_create.xml
/Users/cirrus/backups/2011-06-16_10-33-47/meta/backup_variables.txt
/Users/cirrus/backups/2011-06-16_10-34-12
/Users/cirrus/backups/2011-06-16_10-34-12/backup-my.cnf
/Users/cirrus/backups/2011-06-16_10-34-12/datadir
/Users/cirrus/backups/2011-06-16_10-34-12/datadir/ib_logfile0
/Users/cirrus/backups/2011-06-16_10-34-12/datadir/ib_logfile1
/Users/cirrus/backups/2011-06-16_10-34-12/datadir/ibbackup_logfile
```

```
/Users/cirrus/backups/2011-06-16_10-34-12/datadir/ibdata1
/Users/cirrus/backups/2011-06-16_10-34-12/datadir/mem
/Users/cirrus/backups/2011-06-16_10-34-12/datadir/mem/db.opt
/Users/cirrus/backups/2011-06-16_10-34-12/datadir/mem/dc_p_double#P#p0.ibd
...same database and table files as the previous backup...
/Users/cirrus/backups/2011-06-16_10-34-12/meta
/Users/cirrus/backups/2011-06-16_10-34-12/meta/backup_content.xml
/Users/cirrus/backups/2011-06-16_10-34-12/meta/backup_create.xml
/Users/cirrus/backups/2011-06-16_10-34-12/meta/backup_variables.txt
```

C.2 Sample Directory Structure for Compressed Backup

Here is an excerpt from the file listing under `backup-dir/datadir/mem` for a backup from a MySQL Enterprise Monitor repository database. Notice how the `.ibd` files for InnoDB tables are now compressed to `.ibz` files, while other kinds of files are left unchanged.

```
inventory_types.frm
inventory_types.ibz
log_db_actions#P#p0.ibz
log_db_actions#P#p1.ibz
log_db_actions#P#p2.ibz
log_db_actions#P#p3.ibz
log_db_actions#P#p4.ibz
log_db_actions#P#p5.ibz
log_db_actions#P#p6.ibz
log_db_actions#P#p7.ibz
log_db_actions#P#p8.ibz
log_db_actions.frm
log_db_actions.par
loghistogram_data.frm
loghistogram_data.ibz
```

C.3 Sample Directory Structure for Incremental Backup

An incremental backup produces a directory structure containing a subset of the files from a full backup. All non-InnoDB files such as `*.frm` `*.MYD`, and so on are included. `*.ibd` files are included only if they changed since the full backup, that is, if their maximum `logical sequence number` is higher than the value specified by the `--start-lsn` option.

```
$ find /tmp/backups
/tmp/backups
/tmp/backups/backup-my.cnf
/tmp/backups/datadir
/tmp/backups/datadir/ibbackup_ibd_files
/tmp/backups/datadir/ibbackup_logfile
/tmp/backups/datadir/ibdata1
/tmp/backups/datadir/mem
/tmp/backups/datadir/mem/db.opt
/tmp/backups/datadir/mem/dc_p_double.frm
/tmp/backups/datadir/mem/dc_p_double.par
/tmp/backups/datadir/mem/dc_p_long.frm
/tmp/backups/datadir/mem/dc_p_long.par
/tmp/backups/datadir/mem/dc_p_string.frm
/tmp/backups/datadir/mem/dc_p_string.par
/tmp/backups/datadir/mem/graph_dc_schedules.frm
/tmp/backups/datadir/mem/graph_schedules.frm
... many more files...
```

Appendix D MySQL Enterprise Backup Release Notes

Release notes for MySQL Enterprise Backup are published separately. See [MySQL Enterprise Backup 3.11 Release Notes](#).

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

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 Log 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](#).

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, because that would save little or no space.

See Also [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 3.3.2, “Making an 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 3.3.6, “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 1.3, “Overview of Backup Performance and Capacity Considerations”](#) for an overview, [Section 5.1.11, “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](#) and [--exclude-tables](#).

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 [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

.frm file, 33

A

Antelope, 53, 141
apply, 141
apply-incremental-backup option, 46, 61
--apply-log option, 61
.ARM file, 141
.ARZ file, 141

B

backup, 141
backup directory, 141
backup option, 60
backup repository, 141
backup-and-apply-log option, 60
--backup-dir option, 71
backup-dir-to-image option, 66
backup-image option, 84
backup-my.cnf, 141
backup-my.cnf file, 7
backup-to-image option, 60, 84
backups
 cold, 5
 compressed, 6, 33, 37, 46, 74, 137
 controlling overhead, performance, and scalability, 86
 encrypted, 111
 full, 29, 133
 hot, 5
 incremental, 6, 30, 75, 137
 InnoDB tables only, 53
 logical, 6
 message logging, 92
 parallel, 6
 partial, 33, 78
 physical, 6
 prepared, 7, 45
 preparing to restore, 45
 progress report, 93
 raw, 7, 45
 scheduled, 43
 single-file, 6
 streaming, 6, 40
 to cloud, 41
 to tape, 40, 113
 troubleshooting, 117
 uncompressed, 6, 36
 verifying, 28
 warm, 5
backup_content.xml, 7
backup_content.xml file, 120
backup_create.xml, 7
backup_create.xml file, 120

BACKUP_HISTORY table, 118
BACKUP_PROGRESS table, 118
backup_variables.txt file, 7
Barracuda, 53, 142
benchmarking, 105
binary log, 49, 142
binlog, 142

C

changelog, 139
changes
 release notes, 139
cloud backups, 41
--cloud-access-key-id option, 97
--cloud-aws-region option, 97
--cloud-bucket option, 97
--cloud-object-key option, 97
--cloud-proxy option, 97
--cloud-secret-access-key option, 97
--cloud-service option, 97
--cloud-trace option, 97
cold backup, 5, 142
command-line tools, 6
--comments option, 74
--comments-file option, 74
comments.txt file, 7, 74
--compress option, 33, 74
--compress-level option, 33, 75
--compress-method option, 74
compressed backup, 137
compressed backups, 6, 33, 37, 46, 74
compression, 142
compression level, 143
configuration file, 143
configuration options, 99
connection, 143
connection options, 67
copy-back option, 16, 28, 45, 62
copy-back-and-apply-log option, 62
corruption problems, 117
crash recovery, 45, 143
cron jobs, 43
.CSM file, 7
.CSV file, 7

D

data dictionary, 143
database, 143
--databases option, 82
--databases-list-file option, 82
datadir directory, 7
--datadir option, 69
--data_home_dir option, 69
--decrypt option, 96
decryption, 111
--disable-manifest option, 86
disk storage for backup data, 6, 40

downtime, 143
--dst-entry option, 85

E

--encrypt option, 96
encrypted backups, 111
encryption, 111
error codes, 117
exclude, 144
--exclude-tables option, 79
--exec-when-locked option, 98
extract, 144
extract option, 66, 84

F

FAQ, 121
file format, 144
files backed up, 7
frequently asked questions, 121
.frm file, 7, 144
full backup, 29, 133, 144

G

GRANT statement, 25

H

history table, 144
hot backup, 5, 144

I

ibbackup_logfile file, 7
.ibd file, 7, 49, 145
ibdata file, 7, 145
ibreset command, 117
.ibz file, 7, 145
ib_logfile file, 7
image, 145
image-to-backup-dir option, 66, 84, 84
image_files.xml file, 7, 120
include, 145
--include option, 33, 82
--include-tables option, 79
incremental backup, 6, 75, 137, 145
--incremental option, 75
--incremental-backup-dir option, 78
--incremental-base option, 77
--incremental-with-redo-log-only option, 76
InnoDB, 145
InnoDB tables, 5, 7, 53, 53
 compressed backup feature, 33
 incremental backup feature, 30
installing MySQL Enterprise Backup, 17
instance, 146

K

--key option, 96
--key-file option, 96

L

--limit-memory option, 88
list-image option, 66, 84
locking, 146
log, 7, 61, 146
--log-bin-index, 90
logical backup, 6, 146
logs
 of backup operations, 118
LSN, 30, 75, 146

M

manifest, 7, 86, 119, 146
master, 102, 147
--master-info-file, 91
media management software, 147
media management software (MMS) products, 113
MEMORY tables, 43
message logging, 92
meta directory, 7
MMS products, 113
monitoring backup jobs, 115
.MRG file, 146
my.cnf, 147
my.ini, 147
.MYD file, 7
.MYD file, 146
.MYI file, 7
.MYI file, 146
MyISAM, 147
MyISAM tables, 53
MySQL Enterprise Backup, 147
mysqlbackup, 53, 147
 and media management software (MMS) products, 113
 configuration options, 99
 examples, 29
 files produced, 7
 modes of operation, 59
 options, 54
 overview, 6
 required privileges, 25
 using, 23
mysqlbinlog command, 49
mysqldump, 43, 147

N

--no-history-logging option, 74
--no-locking option, 89
non-TTS backup, 147
--number-of-buffers option, 86

O

offline, 147
--on-disk-full option, 89
online, 148
--only-innodb option, 80

- only-innodb option, 83
- only-known-file-types option, 80
- .opt file, 7, 147
- optimistic backup, 148
- optimistic-busy-tables, 92
- optimistic-time, 91
- options, mysqlbackup, 54
 - connection, 67
 - for cloud storage, 96
 - for compression, 74
 - for controlling backup overhead, performance, and scalability, 86
 - for controlling message logging, 92
 - for controlling progress reporting, 93
 - for encryption, 96
 - for generating metadata, 73
 - for incremental backups, 75
 - for partial backups, 78
 - for single-file backups, 84
 - for special types of backups, 97
 - in configuration files, 99
 - layout of backup files, 71
 - layout of database files, 69
 - modes of operation, 59
 - options in common with mysql, 66
 - standard options, 66
- Oracle Secure Backup, 148
- OSB, 148

P

- page-reread-count option, 89
- page-reread-time option, 89
- .par file, 7, 148
- parallel backup, 105, 108, 148
- parallel backups, 6
- partial backup, 33, 78, 148
- partial restore, 148
- performance
 - of backups, 105
 - of restores, 108
- performance of backup operations, 6
- physical backup, 6, 149
- point in time, 149
- point-in-time recovery, 49
- posix_fadvise() system call, 6
- prepared backup, 7, 45, 149
- privileges, 25
- process-threads option, 87
- progress indicator, 93
- progress table, 149
- progress-interval, 96

R

- RAID, 105, 108
- raw backup, 7, 45, 149
- read-threads option, 87
- redo log, 149

- regular expression, 149
- relay log, 149
- relay-log-index, 90
- relaylog-info-file, 91
- release notes, 139
- replication, 101, 102, 102, 149
- repository, 150
- restore, 150
- restoring a backup, 45
 - at original location, 28
 - backup created with the --use-tts option, 48
 - examples, 46
 - mysqlbackup options, 61
 - overview, 16
 - point-in-time recovery, 49
 - preparation, 45
 - single .ibd file, 49
- row format, 150

S

- SBT, 150
- sbt-database-name option, 85
- sbt-environment option, 86
- sbt-lib-path option, 85
- selective backup, 150
- selective restore, 150
- server, 150
- server repository, 150
- show-progress, 93
- single-file backup, 6, 65, 84, 150
- skip-binlog, 90
- skip-relaylog, 90
- skip-unused-pages, 89
- slave, 101, 102, 150
- slave-info option, 98
- sleep option, 88
- space for backup data, 6
- src-entry option, 84
- start-lsn option, 78
- streaming, 40, 151
- streaming backups, 6
- suspend, 151
- suspend-at-end option, 98
- system backup to tape, 151
- system tablespace, 7, 151

T

- table, 151
- tablespace, 151
- tape backups, 40, 113
- transportable tablespace, 151
- .TRG file, 7
- .TRG file, 151
- .TRN file, 7
- troubleshooting for backups, 117
- TTS, 152

U

--uncompress option, 75
uncompressed backups, 6, 36
--use-tts option, 80

V

validate option, 64
validating a backup, 63
verifying a backup, 28

W

warm backup, 5, 152
--with-timestamp option, 73
--write-threads option, 88