

Xpand Fast Backup and Restore

See [also FTP Server Setup for Fast Backup and Restore](#).

- [Parallel Backup and Restore](#)
- [Backup](#)
 - [SFTP](#)
 - [Compression](#)
 - [Tuning BACKUP Performance](#)
 - [Example BACKUP Commands](#)
 - [Getting Information on Backups](#)
 - [Displaying all Backups on a Source](#)
 - [Displaying the Status of Backup/Restore](#)
 - [Displaying all Backed Up Tables on a Source](#)
- [Restore](#)
 - [REPLICAS = N](#)
 - [Tuning RESTORE Performance](#)
 - [Example RESTORE Commands](#)
- [Global Variables](#)
- [Using BACKUP/RESTORE to Resolve Replication Issues](#)
 - [Scenario I : Xpand slave does not exist or needs to be reset with master.](#)
 - [Scenario II: Xpand Master needs to be restored, Xpand Slave looks good.](#)
 - [Scenario III: Xpand Master and Slave both need to be restored and reset.](#)
- [Mysqldump](#)
- [Additional Resources](#)

Parallel Backup and Restore

Xpand implements Fast Backup and Restore as a binary backup mechanism that works at the row level. Each Xpand node sends its data directly to the backup target in parallel, eliminating bottlenecks and allowing backup to scale with cluster size. Restore is handled similarly; the initiating node coordinates with other participating nodes in parallel to read from the dump file and restore replicas.

The structure of the dumped data consists of text files describing the schema and metadata about the backup. The dumped data also contains binary files of compacted row data and data consistency information. Modifying these files may make it impossible to restore from the backup.

Xpand uses a Cyclic Redundancy Check (CRC) during Backup and Restore. Calculating the CRC allows the Restore function to validate the data written during the Backup function. CDC adherence also requires that any failed backup or restore be restarted from the beginning.

Xpand's Fast Backup and Restore:

- Is implemented as SQL commands that are executed from a MySQL client.
- Supports passive (versus active) mode FTP or SFTP
- Supports password-based authentication for SFTP
- Cannot be executed via stored procedures, functions or triggers.

Xpand recommends running BACKUP and RESTORE in a [screen session](#) or using [tmux](#) to avoid connection failures.

 Xpand also recommends running both BACKUP and RESTORE during off-peak hours.

Backup

Backup runs as a single transaction and will pin [BigC](#) (Xpand's garbage collection) for the duration of the operation. This may require operational consideration if your cluster does not have enough free space for the undo log created from this long running transaction.

Backup SQL Syntax

```
sql> BACKUP identifier [, identifier] [EXCLUDING excluded_identifier] TO 'target' [COMPRESSED]
```

Both the identifier and excluded_identifier should refer to valid system objects. This can be a fully qualified name, or include wildcards (*) for the database, object, or both. For example:

- `db_name.obj_name` - references a specific object in the given database (i.e. table name, view name).
- `db_name.*` - expands to all accessible objects in that database.
- `*.*` - expands to every accessible object in the system. (Meaningless for excluded objects.)

Additionally the excluded_identifier can specify a given table. For example:

- `*.tbl_name` - references a given table name in any database that is to be excluded from the backup.

Xpand uses username and password authentication. The server path for FTP represents the relative path from the FTP root directory whereas the path for the SFTP server should be the absolute path from the Linux root directory. For example: "ftp://user:password@servername/relative_path/from_ftp_root" or "sftp://user:password@servername:absolute_path/from_Linux_root".

Care should be taken when excluding objects from a backup. This can cause errors during RESTORE if another object depends on the excluded object.

 For example, if a trigger or view refers to an excluded object that does not already exist in the restored environment, an error will occur during the RESTORE.

Triggers, views, and foreign keys that reference an excluded object will be backed up, unless they are specifically excluded.

SFTP

To use an SFTP for Backup and Restore, specify the absolute path of the file location (versus relative path for FTP) and provide a username and password for authentication. For example: "sftp://user:password@servername:absolute_path/from_Linux_root"

Compression

To create a compressed backup, specify the COMPRESSED option. This functionality reduces bandwidth requirements for backup to remote locations and reduces the amount of storage utilized, but adds additional latency to complete a backup.

Tuning BACKUP Performance

You may want to tune BACKUP to be slower or faster.

```
sql> set global backup_backup_concurrency =
desired_value;
```

The global variable backup_backup_concurrency determines the number of tables that can be backed up simultaneously. With the default value of 1, the backup process is limited to one table at a time.

Example BACKUP Commands

```
sql> BACKUP db_name1.*, db_name2.tbl_name1, db_name2.tbl_name2 TO 'ftp://storage01/uploads/johndoe+kibbles.
jun01';
sql> BACKUP *.* EXCLUDING db_name2.tbl_name1 TO 'ftp://storage01/uploads/johndoe.kibbles.jun01';
sql> BACKUP test TO 'sftp://plato:fy569gk@piquet.colorado.sproutsys.com:22/tmp/tc5326/tc5326.bkup';
```

Getting Information on Backups

The following queries show the status of all currently running backups and restores.

Displaying all Backups on a Source

```
sql> SELECT * from system.backups where source="ftp://storage01
/uploads"\G
***** 1. row *****
source: ftp://storage01/uploads
backup: 1_52T_large_dogfood_dbx2
cluster_name: c15d97786267d0d4cb
version: 5.0.45-clustrix-7.5.2
status: ERROR
start_time: 2016-06-15 01:43:36
completed_time: 2016-06-15 03:19:31
size: 0
```

Displaying the Status of Backup/Restore

Run the following query to view information for a backup in progress

```

sql> select * from system.backup_status\G
***** 1. row *****
nodeid: 2
id: 5756997494884148226
type: BACKUP
url: ftp://storage01/uploads/1_52T_large_dogfood_dbx2_2
objs: (("test" . "" )("clustrix_ui" . "" )("bugstest_01" . "" )("jttunes_01" . "" )("longstats_01" . "" )("npktest_01" . "" )("perfstats_01" . "" )("run_resources_01" . "" )("statd_01" . "" )("bugstest_02" . "" )("jttunes_02" . "" )("longstats_02" . "" )("npktest_02" . "" )("perfstats_02" . "" )("run_resources_02" . "" )("statd_02" . "" ))
start_time: 2016-06-22 22:50:32
expected_bytes: 793977192448
replicas: 1070
rows: 49589227
bytes: 23991322310
1 row in set (0.01 sec)

```

Displaying all Backed Up Tables on a Source

```

sql> SELECT * from system.backup_tables where source="ftp://storage01/uploads" limit 2;
+-----+-----+-----+-----+-----+
| source                | backup                | db                | table                | size |
+-----+-----+-----+-----+-----+
| ftp://storage01/uploads | 1_52T_large_dogfood_dbx2 | bugstest_01      | attach_data         | 0    |
+-----+-----+-----+-----+-----+
| ftp://storage01/uploads | 1_52T_large_dogfood_dbx2 | bugstest_01      | attachments         | 0    |
+-----+-----+-----+-----+-----+
2 rows in set (12.16 sec)

```

Restore

Use the following syntax to restore objects previously backed up with Fast Backup.

Restore SQL Syntax

```

sql> RESTORE identifier [AS identifier] [, identifier [AS identifier]] [EXCLUDING excluded_identifier] FROM
target [REPLICAS=N]

```

Both the identifier and excluded_identifier should refer to valid system objects. This can be a fully qualified name, or include wildcards (*) for the database, object, or both. For example:

- db_name.obj_name - references a specific object in the given database (i.e. table name, view name).
- db_name.* - expands to all accessible objects in that database.
- *.* - expands to every accessible object in the system. (Meaningless for excluded objects.)

Additionally, the excluded_identifier can specify a given table. For example:

- *.tbl_name - references a given table name in any database that is to be excluded from the restore.

Tables may be renamed during the restore process by providing a table alias [AS identifier].

The server path for FTP represents the relative path from the FTP root directory whereas the path for the SFTP server should be the absolute path from the Linux root directory. For example: "ftp://user:password@servername/relative_path/from_ftp_root" or "sftp://user:password@servername:absolute_path/from_Linux_root".

Care should be taken when restoring from a backup that used the EXCLUDING option. This can cause errors during restoration if another object depends on the excluded object(s).

For example, if a trigger or view refers to an object that does not already exist in the restored environment, an error will occur.

Triggers, views, and foreign keys that reference an object excluded from the backup will be successfully restored only if the object exists in the restored environment.

REPLICAS = N

The number of replicas for a table or index is recorded within the backups. Normally when RESTORE is run, tables and indexes are restored with the same number of replicas as the original. This option allows you to explicitly specify the number of replicas that should be created as part of the RESTORE operation.

For example, if a table had three replicas when backed up, the following example would create the table with only two replicas.

```
sql> RESTORE * FROM 'ftp://storage01/uploads/johndoe.kibbles.jun01' REPLICAS = 2;
```

See more on specifying the number of replicas in a table in the section on [Managing Data Distribution](#).

Tuning RESTORE Performance

You may want to tune RESTORE to be slower or faster.

```
sql> set global backup_restore_concurrency =  
desired_value;
```

The global variable `backup_restore_concurrency` determines the number of slices that can be restored simultaneously to each node.

Example RESTORE Commands

```
sql> RESTORE db_name.tbl_name AS db_name_backup.tbl_name FROM 'ftp://username@server.com/backups/backupfolders  
/backup_file';  
sql> RESTORE tbl_name, tbl_name2, tbl_name3 FROM 'ftp://username@ftp://server.com/backups/backupfolders  
/backup_file';  
sql> RESTORE ehms FROM 'sftp://root@kartma6127/root/ehmsftp';
```

Global Variables

The following global variables impact Fast Backup and Restore. The defaults provided are generally acceptable. These variables are not available by session.

Name	Description	Default Value
<code>backup_backup_concurrency</code>	The number of tables that can be backed up simultaneously.	1
<code>backup_restore_concurrency</code>	The maximum number of slices restored concurrently on each node.	16
<code>backup_write_compression_level</code>	Compression level from 1 (fastest) to 9 (best compression)	6

Using BACKUP/RESTORE to Resolve Replication Issues

Fast Backup and Restore can also be used to resolve the following issues in replicated environments.

Scenario I : Xpand slave does not exist or needs to be reset with master.

1. Take a fresh backup using BACKUP sql command on master. This will minimize the time it will take for the slave to catch up if older backups were to be used.
2. Stop the slave if one already exists. If none exists as yet, please refer to [Using Xpand as a Replication Slave](#) and [Starting and Stopping a Slave Process](#) for steps to create a new Xpand slave and leave it in stop mode after successful creation.
3. Restore backup taken in step 1 on slave using RESTORE command.
4. Master's binlog position when backup initiates is stored under `your_backup_location/metadata/binlogs` file. This file contains relevant information in the format `logfile.seq#:position`, for example, `foo1.002843:99744547`.
5. Reset the position recorded in the binlogs file by issuing the command:

```
sql> CHANGE MASTER TO MASTER_LOG_FILE='foo1.002843', MASTER_LOG_POS=99744547, MASTER_HOST='foo1',  
MASTER_USER='root', MASTER_PORT=3306;
```

6. Start the slave.

Scenario II: Xpand Master needs to be restored, Xpand Slave looks good.

1. Promote Xpand slave to act as master. Please refer to [Promoting a Xpand Slave to Master](#) for the necessary steps to promote a passive Xpand slave to an active Master mode.
2. Follow steps outlined in Scenario I to restore and reset the old master which is now in passive slave mode.
3. If necessary, repeat the steps to switch over the current passive slave to active master once both master and slave are fully caught up.

Scenario III: Xpand Master and Slave both need to be restored and reset.

1. Stop the slave.
2. Restore the master and slave using RESTORE command.
3. Reset the slave to the current position of the master (i.e. from SHOW MASTER STATUS on the master), after the restore, before it begins taking writes from clients. Note that using this technique, you are losing incremental changes since the backup. An advanced method to capture the incremental changes since backup would be to replay the binlogs if they are still available, and change the serverid of the master, in order to have it reapply the events in the binlog. Please contact [Xpand Support](#) for assistance with this advanced procedure.
4. Start the slave.

If you receive the following timeout error during the backup process, "*Socket timeout while waiting for server response: read_timer*", the global variable `ftp_read_timeout_secs` can be increased to a higher value, up to 3000, to prevent this error. This global applies to backups and restores of both FTP and SFTP servers.

Mysqldump

To create a dump of the Xpand, use the `mysqldump` utility, which is included on each node in your cluster. If you intend to import the dump into a Slave, create the binary log for the replication stream before you create the dump.

To create the dump, issue the following command:

```
shell> mysqldump -u user -h clustrix host --single-transaction --master-data=2 --all-databases > mydumpfile.dump
```

- The `--single-transaction` flag ensures a consistent snapshot of the database by querying for all of the data in a single transaction and permits continued access to the Xpand while the dump is being created.
- The `--master-data=2` flag inserts a `CHANGE MASTER` command in a comment near the top of the dump file, indicating the location in the binary log where a slave must start to be consistent with the dump. If no binlog is being created, omit the `--master-data` flag.

Additional Resources

[FTP Server Setup for Fast Backup and Restore](#)

[List of Errors for Backup and Restore](#)