

Percona XtraBackup Current and Future State

What's the future of the open-source industry standard for MySQL hot backup?

Tyler Duzan

Product Manager for MySQL & Cloud

FOSDEM 2020

February 1st, 2020



Who Am I?

- **My name is Tyler Duzan**
- **Product Manager for MySQL Software and Cloud at Percona**
- **Previously was an Ops/Security Engineer for 13 years**
- **Active in the Open Source community since 2003**



What is XtraBackup?

- **Enterprise-Grade**
- **Cloud-Integrated**
- **Supports MyRocks**
- **100% Free and Open Source Software**
- **Online Hot Backup for MySQL**



PERCONA
XtraBackup

What is XtraBackup

- **xtrabackup** –
 - Main binary utility, performs the backup and preparation stages
- **xbstream** –
 - Both a file format and a utility. This allows simultaneous compression and streaming of backup data
- **xbcrypt** –
 - Supports encryption and decryption of backups, integrates with xbstream
- **xbcloud** -
 - Provides a method to perform streaming backups to supported cloud object storage. Currently supports the S3 API and the Swift API

Current State

Split Binaries for PXB 8.0 and PXB 2.4

- **MySQL 8.0 introduced a new data dictionary and redo log format**
- **Data dictionaries in MySQL 8.0 give many great advantages, but are fundamentally incompatible with the necessary code paths for supporting 5.7**
- **Oracle also split their MySQL Enterprise Backup binaries for the same reasons**
- **PXB 2.4 and 8.0 are independent release series, both still actively developed.**

MyRocks Consistent Backups

- **New data dictionary format relies on InnoDB, which means even in a purely MyRocks database the system tables would still use InnoDB**
- **Added checkpointing into MyRocks both in Percona Server and upstream in FB-MySQL**
- **Checkpointing allows us to gather an LSN at the point in time of the backup, in order to appropriately process the MyRocks WAL**
- **Utilizing the log_status table allows us to match consistency to InnoDB binary log while backing up all InnoDB databases including the system tables**

Support for Streaming Cloud Object Storage

- **Added into xcloud as a necessary feature for enabling the use of MySQL w/ XtraBackup in a Kubernetes environment**
- **Expanded capabilities of xcloud and added support for the S3 API**
- **S3 API works with AWS S3, Minio, Google Cloud Storage, and others**
- **Allows a backup to be executed where nothing is written locally to disk, it is all immediately streamed offsite.**

Roadmap Overview

Security Improvements

- **Support transition encryption keys to ensure that PXB based SST can succeed when every PXC member in a cluster uses a different unique encryption key**
- **Tighter integration with Hashicorp Vault, including the use of Vault to generate a unique encryption key for each backup before it transits the wire.**
- **Support for the PS 8.0 Encryption Threads and Key Rotation feature when it becomes GA**

Performance Improvements

- **New flag that will auto-tune the value typically set by --use-memory by analyzing the MySQL redo log**
- **Support adjustable chunk sizing in xstream and xcloud**
- **Make --parallel work on a chunk level in xtrabackup rather than a file level for speeding up backups of large single-file databases**
- **Add exponential back-off into our xcloud API wrappers to reduce failures during retries**

Improved Integrations

- **Add Native Google Cloud Storage support to xbccloud**
- **Add Native Azure Blob support to xbccloud**
- **MyRocks Incremental Backups**
- **Support restoring/retrieving from tarballs with xbccloud**
- **Add support for IAM Roles in our S3 integration in xbccloud**

Questions?

Please Reach Out if You Have Ideas/Concerns

- tyler.duzan@percona.com
- **Zendesk Support Tickets**
- **File a New Feature in JIRA**



**Champions of Unbiased
Open Source Database Solutions**