How Transparent Data Encryption is built in MySQL and Percona Server ?

Robert Golebiowski, Senior Software Engineer at Percona

PERCONA

Transparent
Data Encryption

KEYRINGS

FOSDEM 2020

KEYRINGS

- What is keyring ?
- Plugin installation
  - always successful
  - keyring variables may need correction:
    - keyring_vault_config
    - keyring_file_data

KEYRINGS

Keyring file

KEY ID	KEY TYPE	KEY OWNER	KEY LENGTH	KEY
MK 1	AES		32	001010101 …
Key 1	AES	Robert	16	100111010 …

KEYRINGS

Keyring vault

KEY ID	KEY TYPE	KEY OWNER	KEY LENGTH	KEY
MK 1		NULL		
Key 1		Robert		

FOSDEM 2020

KEYRINGS

PERCONA
Server for MySQL

- Writes to keyring_file
  - backup file keyring.backup (whole content is rewritten)
- Writes to keyring_vault
  - connection lags (only one key is sent)

KEYRINGS

PERCONA Server for MySQL

Each keyring should store keys in a separate place.

- why needed ?
- natural for keyring_file
- work needed for keyring_vault

KEYRINGS

PERCONA
Server for MySQL

- separate mount point per MySQL/PS:

```
curl -L -H "X-Vault-Token: TOKEN" –cacert VAULT_CA
--data '{"type":"generic"}' --request POST
VAULT_URL/v1/sys/mounts/SECRET_MOUNT_POINT
```

- separate *directory* inside mount point per each server:

```
config for server1:
secret_mount_point= <mount_point>/server1
config for server2:
secret_mount_point=<mount_point>/server2
```

FOSDEM 2020

KEYRINGS

PERCONA Server for MySQL

keyring_vault's configuration file

- vault_url
- secret_mount_point
- token
- vault_ca

OPTIONAL

KEYRINGS

Keys inside Vault server are base64 encoded

```
echo NDhfSU5OT0RCS2V5LTc2NGQzODJhLTczMjQtMTFlOS1hZDhmLTljYjZkMGQ1ZGM5OS0xMF8= | base64 -d
```

```
48_INNODBKey-764d382a-7324-11e9-ad8f-9cb6d0d5dc99-10_
```

KEYRINGS

PERCONA Server for MySQL

keyring_udf

Used for storing user’s secret inside keyrings.

Set of UDFs:

- keyring_key_generate
- keyring_key_fetch
- keyring_key_length_fetch
- keyring_key_type_fetch
- keyring_key_store
- keyring_key_remove

Transparent
Data Encryption

InnoDB Encryption

InnoDB Encryption

Reminder: Tablespace consists of pages

What is Master Key encryption ?

FOSDEM 2020

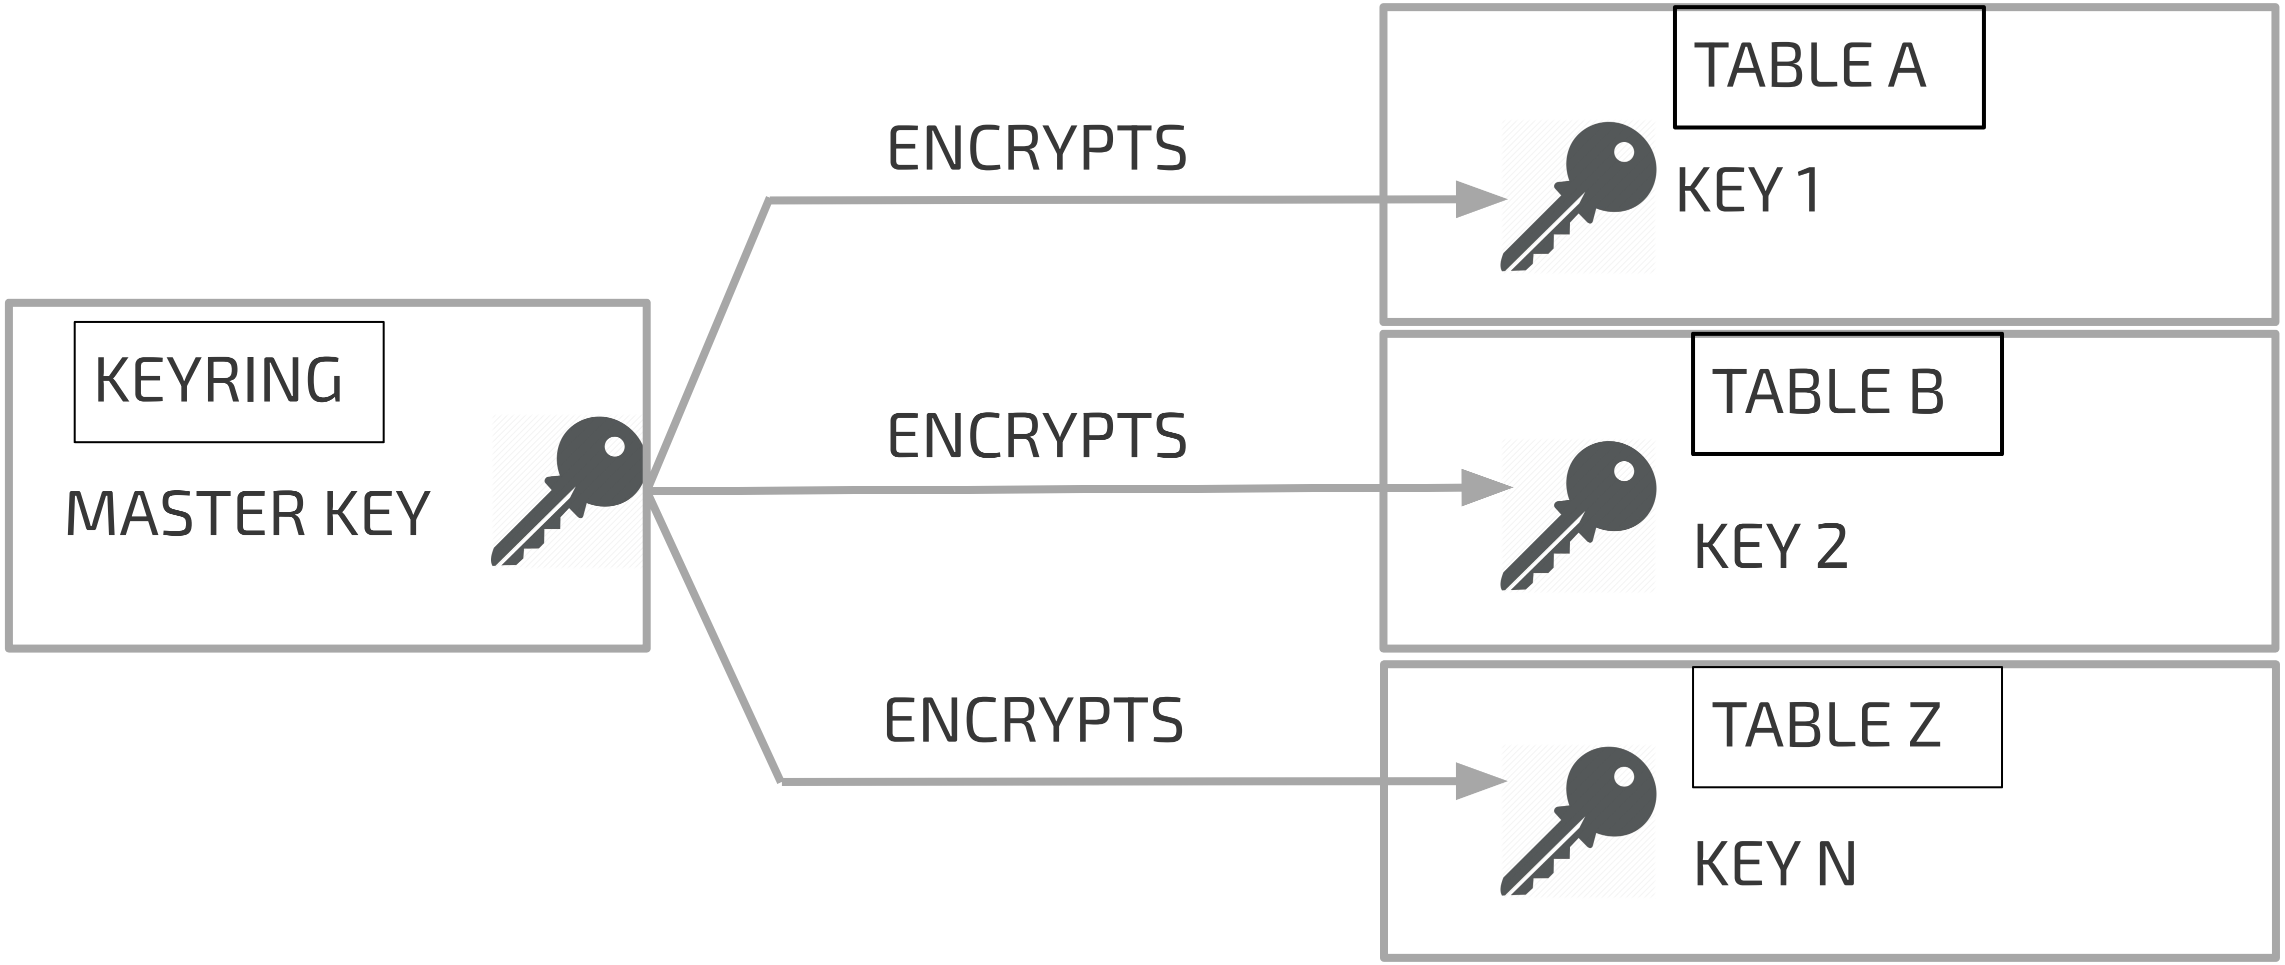
KEYRING

MASTER KEY

ENCRYPTS

ENCRYPTS

ENCRYPTS

TABLE A

KEY 1

TABLE B

KEY 2

TABLE Z

KEY N

FOSDEM 2020

InnoDB Encryption

Tablespace’s encryption header.

Resides in page 0. Page 0 is never encrypted.

ENCRYPTION_KEY_MAGIC (_V1,_V2,_V3)
KEY ID
UUID
ENCRYPTED (TABLESPACE KEY, IV)
CRC32 OF (TABLESPACE KEY,IV)

INNODBKey-UUID-KEY_ID

InnoDB Encryption

PERCONA Server for MySQL

- How do we know which Master Key we should fetch from keyring to decrypt a table? (question from client)
- How do we know if the key used is the correct one?
- How do we make sure that we are able to decrypt table when we need it?

FOSDEM 2020

InnoDB Encryption

PERCONA
Server for MySQL

Encrypted tables validation

- Read page 0
- Read encryption information from page 0
- Get master key from keyring
- Decrypt tablespace key and iv with master key
- Make sure crc32 is correct

If any failed: Mark tablespace as missing

InnoDB Encryption

PERCONA
Server for MySQL

What crypto are used ?

AES 256 ECB for tablespace key and iv encryption (hardcoded)

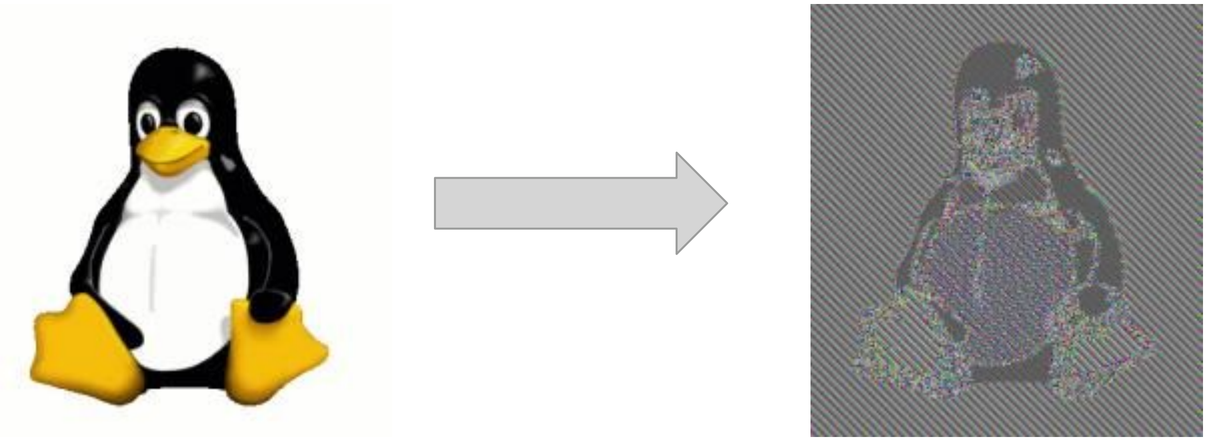

InnoDB Encryption

What crypto are used ?

AES 256 ECB for tablespace key and iv encryption (hardcoded)

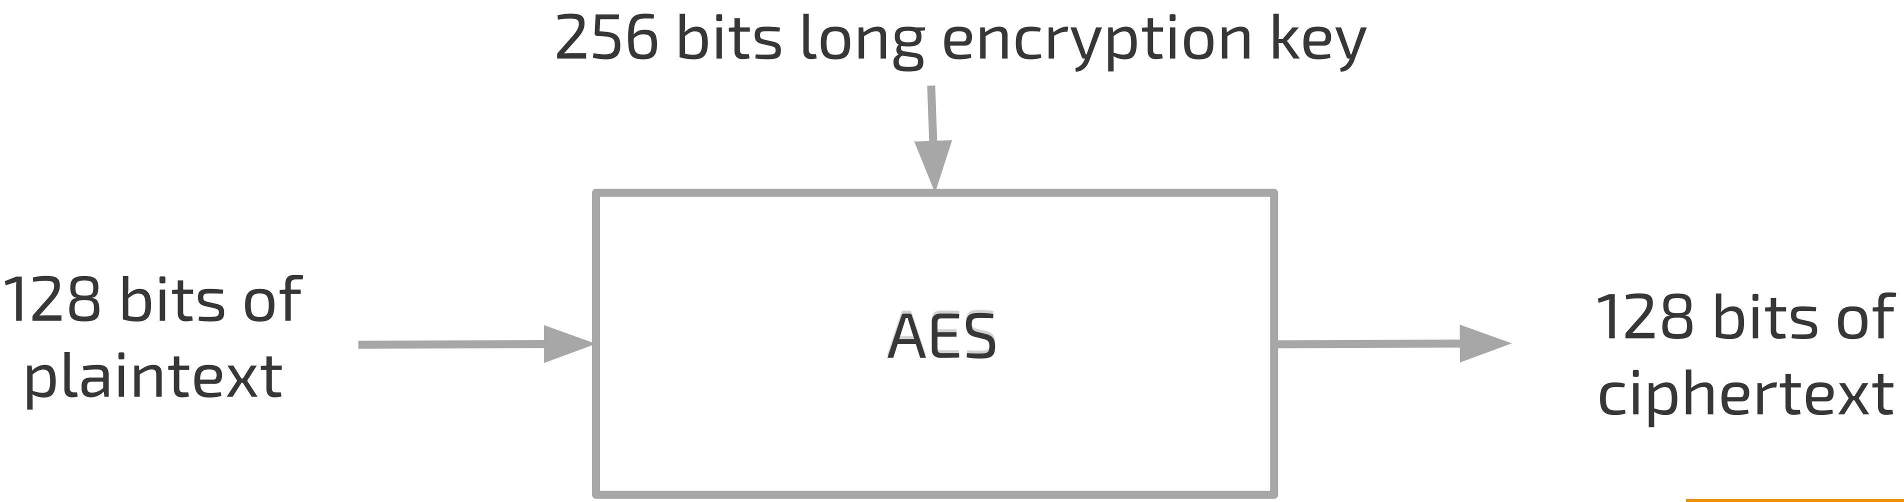

FOSDEM 2020

InnoDB Encryption

What crypto are used ?

AES 256 CBC for page encryption (hardcoded)

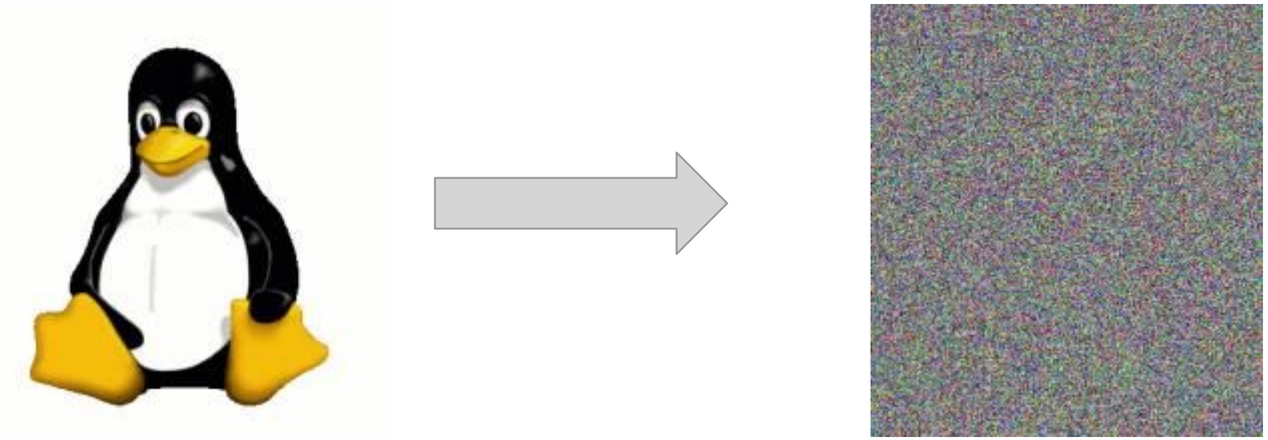

InnoDB Encryption

What crypto are used ?

AES 256 ECB for tablespace key and iv encryption (hardcoded)

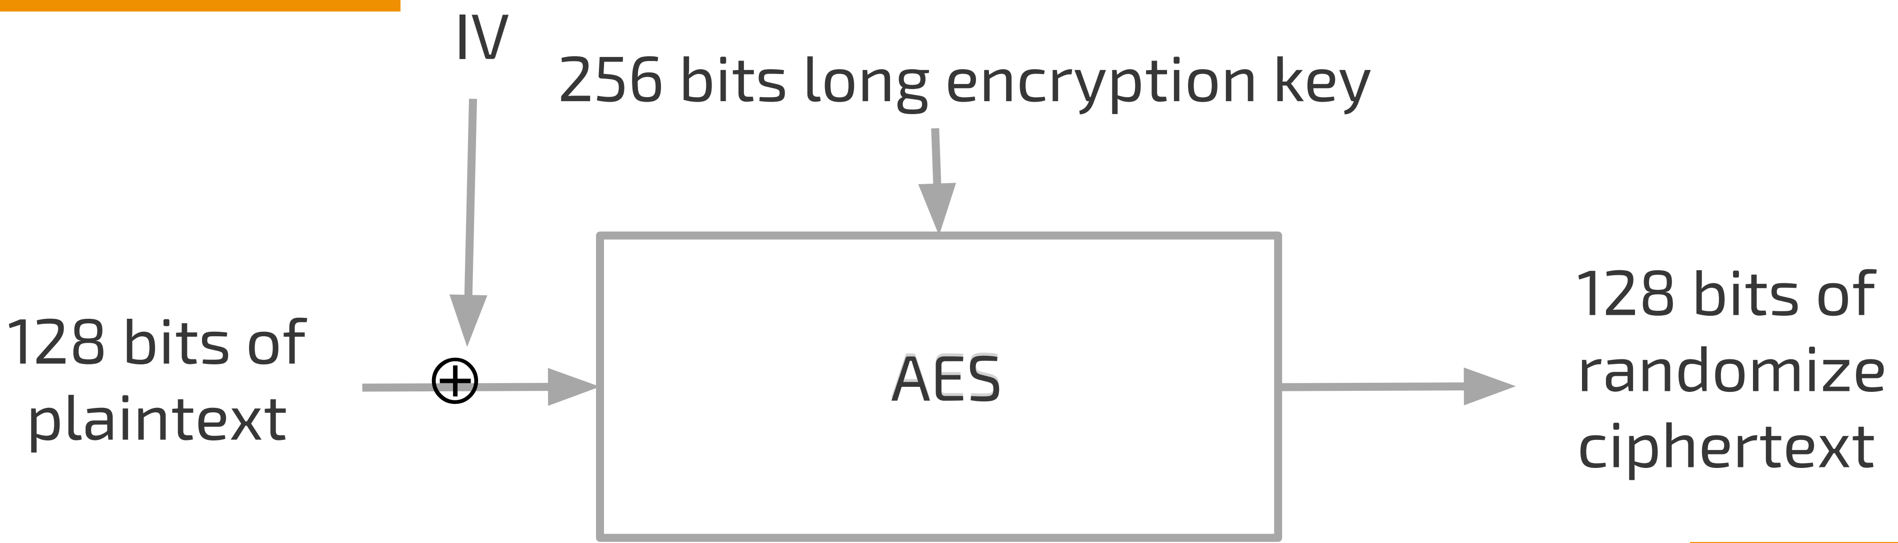
IV
256 bits long encryption key
128 bits of plaintext
⊕
AES
128 bits of randomize ciphertext

FOSDEM 2020

InnoDB Encryption

PERCONA
Server for MySQL

Master key rotation:

- Generate new Master Key
- Go over all encrypted tables. For each table:
    - Re-encrypt tablespace key and iv with new Master Key
    - Update the encryption information in tablespace header (page 0)

FOSDEM 2020

InnoDB Encryption

ENCRYPTION_KEY_MAGIC (_V1,_V2,_V3)
~~KEY ID~~ NEW KEY ID
~~UUID~~ NEW UUID
~~ENCRYPTED (TABLESPACE KEY, IV)~~ RE-ENCRYPTED
CRC32 OF (TABLESPACE KEY,IV)

InnoDB Encryption

Master Key Rotation. Why needed?

- Improves safety
- Speeds up the innodb startup in case we have restored tables from different backups (for keyring_vault without per server separation of keys in Vault server)

InnoDB Encryption

PERCONA
Server for MySQL

Drawbacks of Master Key encryption.

FOSDEM 2020

Transparent
Data Encryption

Binlog encryption

Binlog encryption

Binlog encryption, 5.7

- --encrypt_binlog
- --master_verify_checksum

Binlog encryption

Binlog encryption, 5.7

PERCONA
Server for MySQL

- new event: Start_encryption_event

After Start_encryption_event rest of the binlog is encrypted.

This event is never send over the network.

The events between master and slave are not encrypted (use TLS)

mysqlbinlog cannot decrypt, however there is --read-from-remote-server

Binlog encryption

Binlog encryption, 8.0

Upstream implementation. Follows Master Key encryption rules.

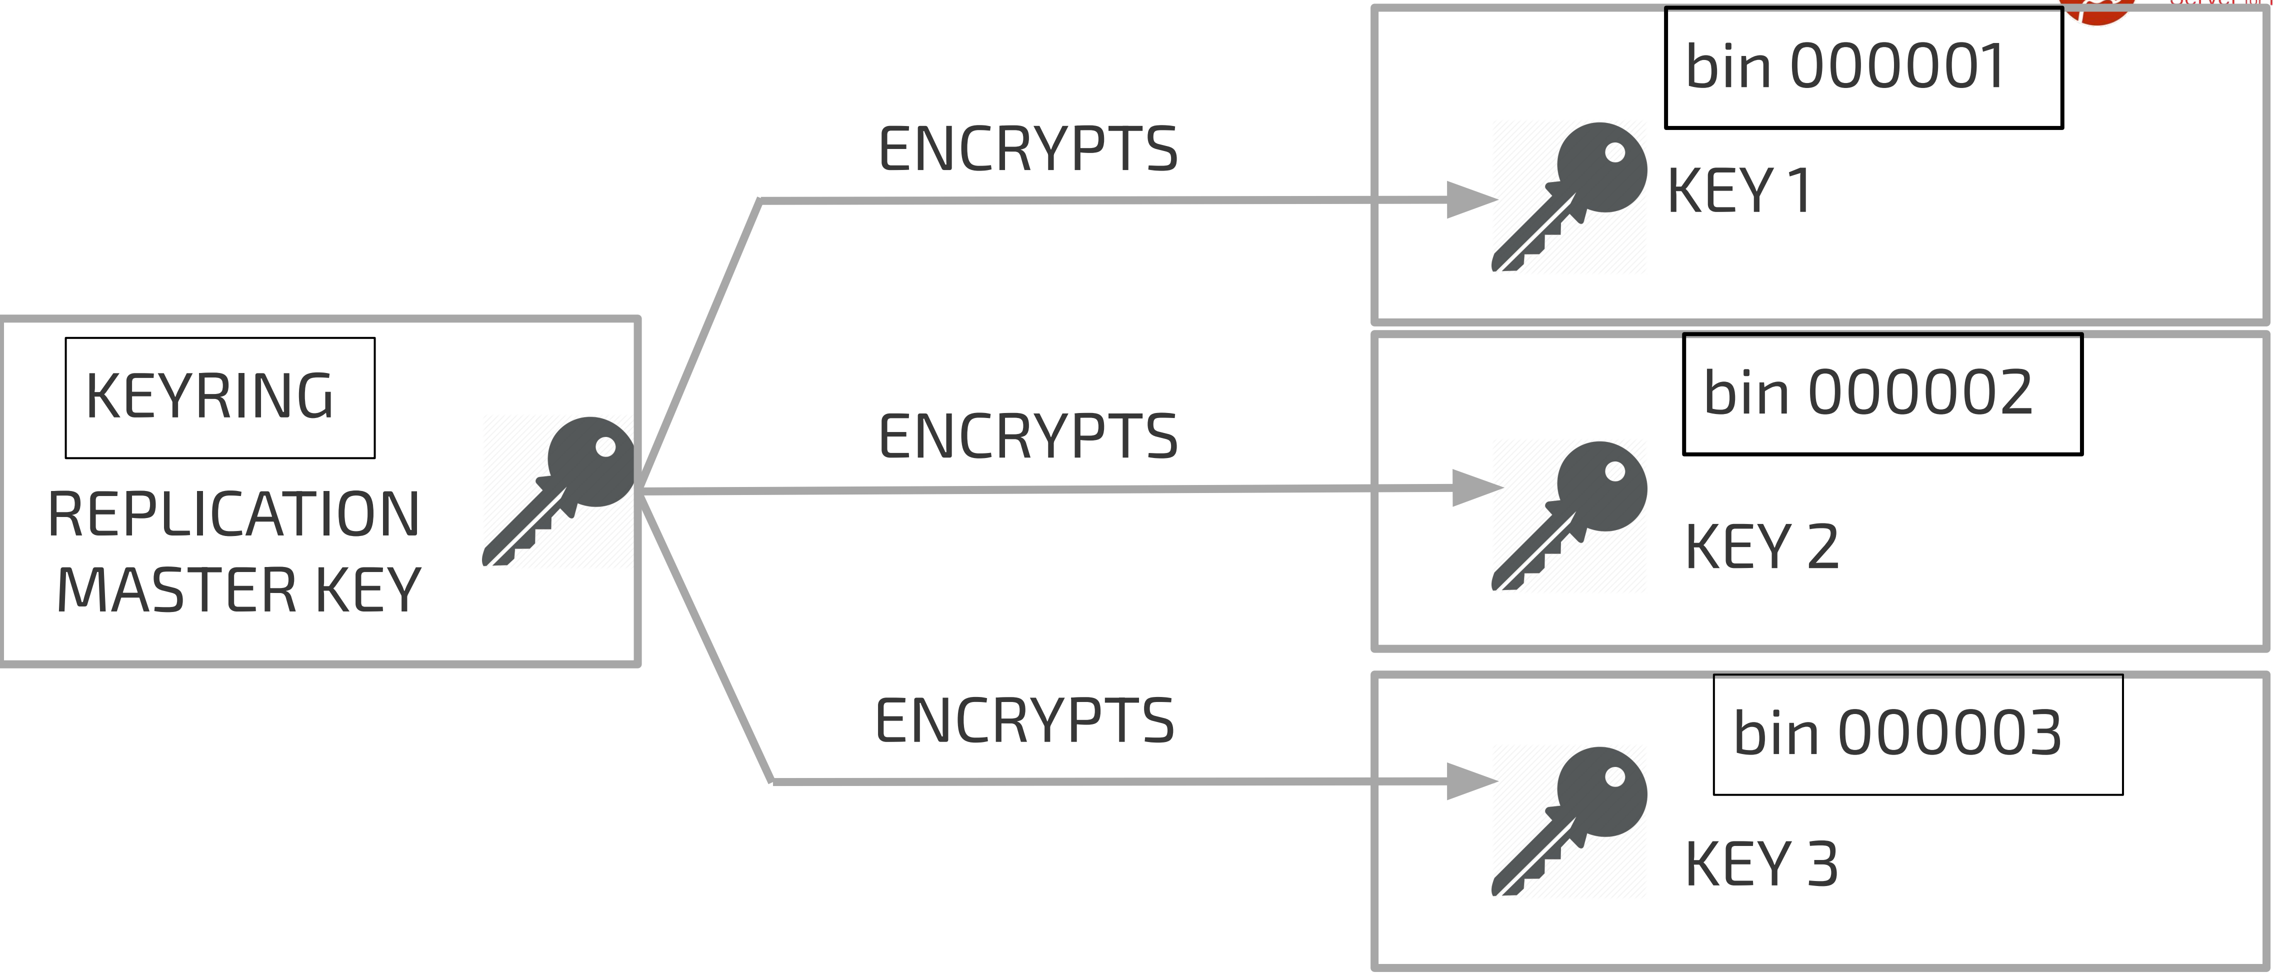
KEYRING
REPLICATION MASTER KEY
ENCRYPTS
ENCRYPTS
ENCRYPTS
bin 000001
KEY 1
bin 000002
KEY 2
bin 000003
KEY 3

FOSDEM 2020

Undo and Redo Log Encryption

Undo tablespace encryption:

- for MK pages are encrypted/decrypted as innodb_undo_log_encrypt is ON/OFF

Redo log encryption almost the same as binary log encryption.

FOSDEM 2020

Transparent
Data Encryption

System Tablespace and Double Write Buffers Encryption

FOSDEM 2020

System tablespace and double write buffer encryption

System tablespace encryption in PS (possible at bootstrap):

- --innodb_sys_tablespace_encrypt (5.7 and 8.0)
- mysql.ibd by enabling --default-table-encryption = on (bootstrap) ALTER TABLESPACE mysql ENCRYPTION='Y' (mysql and PS)
- double write buffer encrypted (part of system tablespace

Parallel double write buffer encryption:

- --innodb_parallel_dblwr_encrypt

Transparent Data Encryption

PERCONA Server for MySQL

Thank you !

FOSDEM 2020

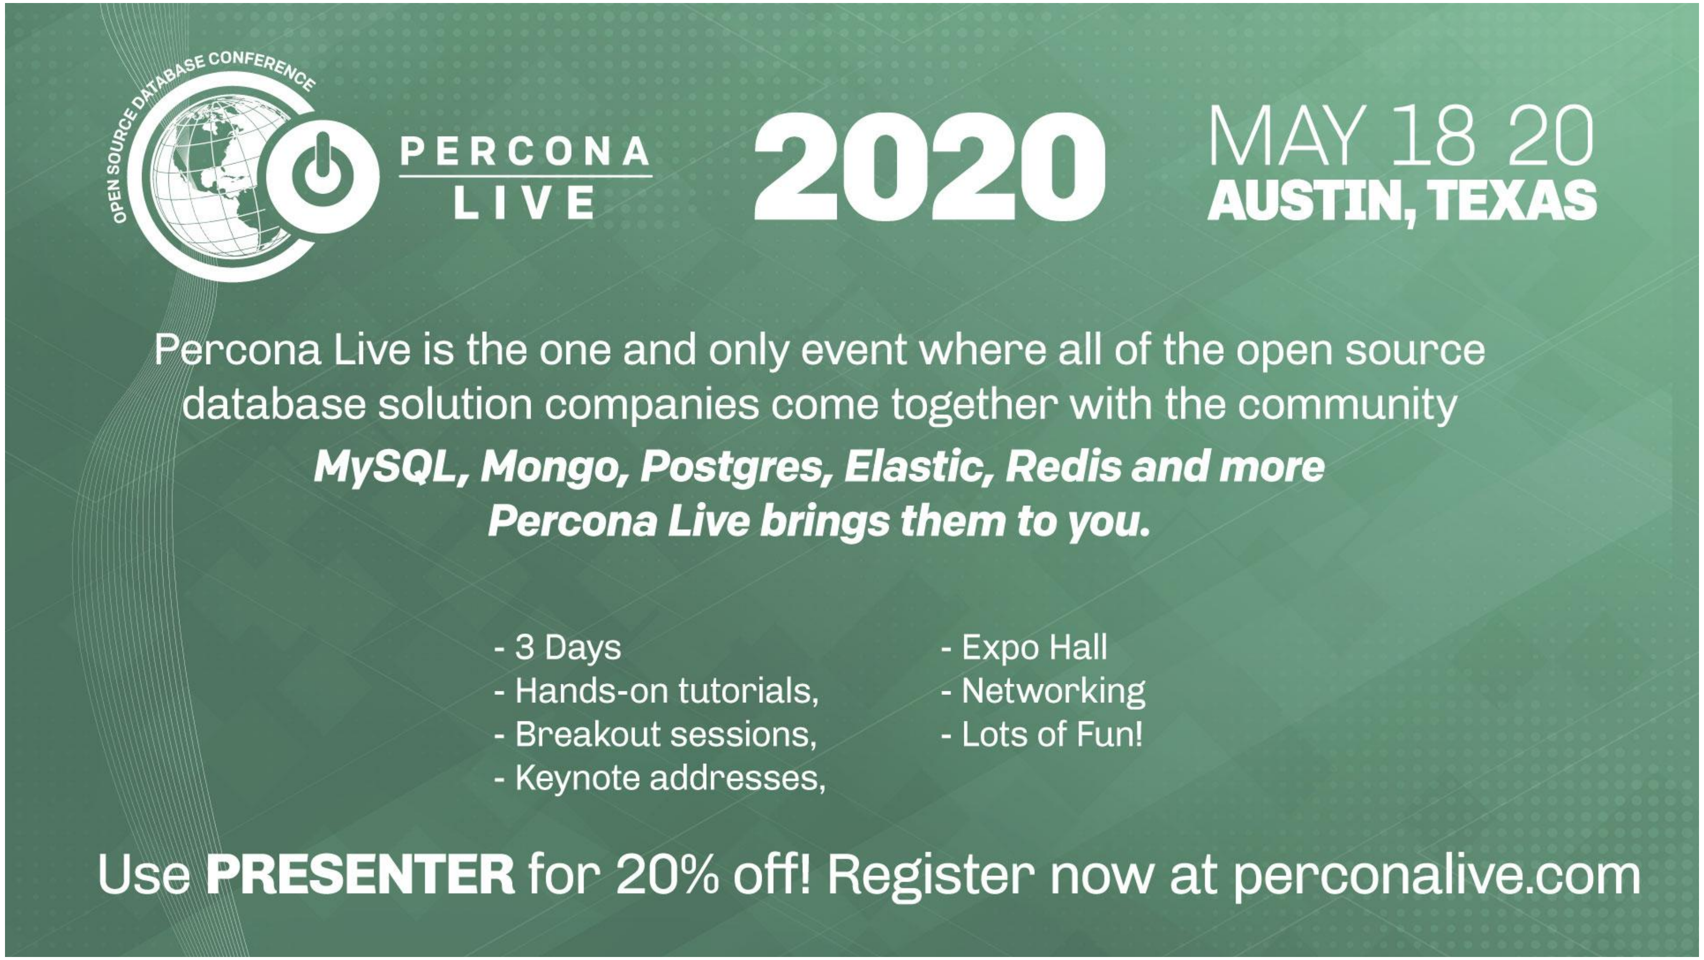

OPEN SOURCE DATABASE CONFERENCE

PERCONA LIVE

2020

MAY 18 20
AUSTIN, TEXAS

Percona Live is the one and only event where all of the open source database solution companies come together with the community

MySQL, Mongo, Postgres, Elastic, Redis and more
Percona Live brings them to you.

- 3 Days
- Hands-on tutorials,
- Breakout sessions,
- Keynote addresses,

- Expo Hall
- Networking
- Lots of Fun!

Use **PRESENTER** for 20% off! Register now at perconalive.com

Undo and Redo Log Encryption

Transparent Data Encryption

PERCONA
Server for MySQL

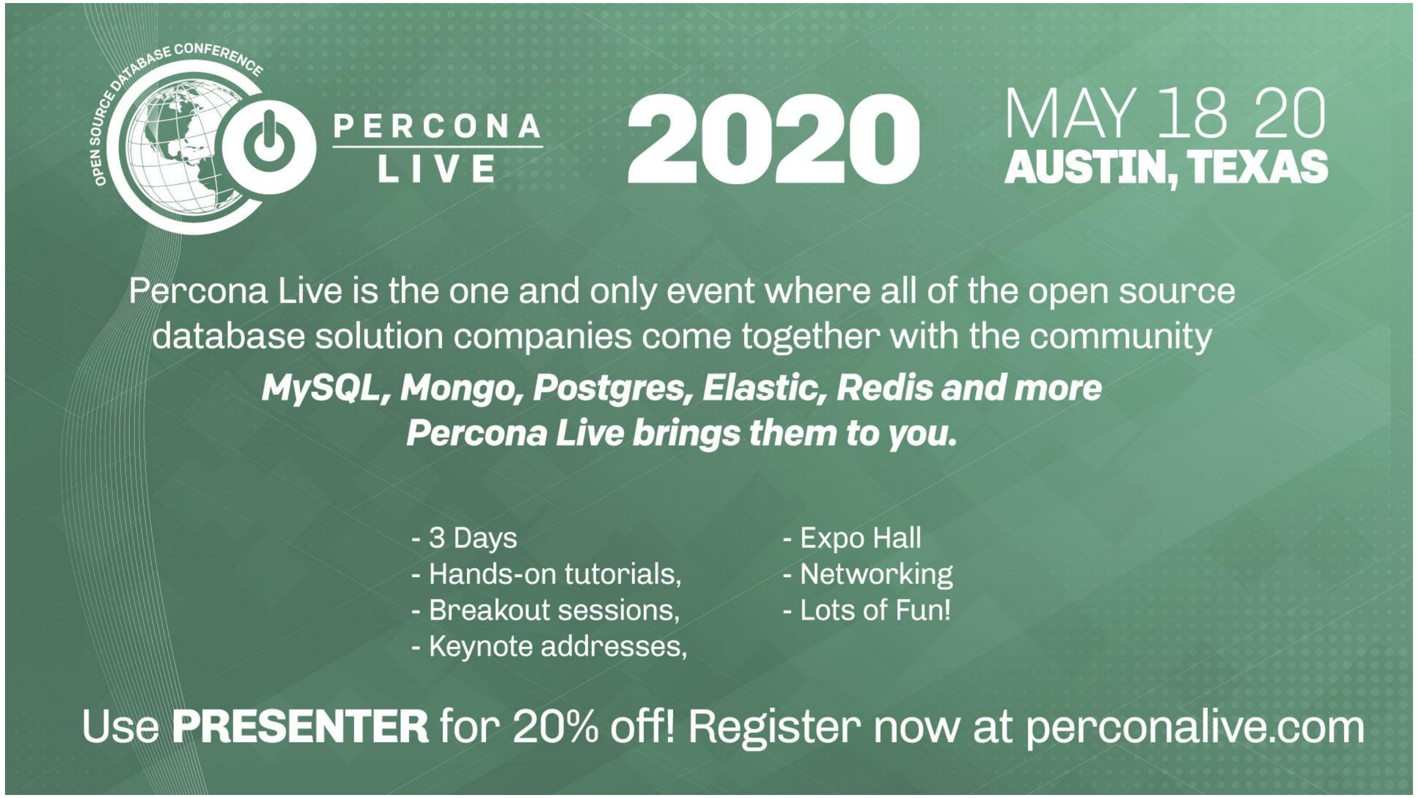

OPEN SOURCE DATABASE CONFERENCE

PERCONA LIVE **2020**

MAY 18 20
AUSTIN, TEXAS

Percona Live is the one and only event where all of the open source database solution companies come together with the community

MySQL, Mongo, Postgres, Elastic, Redis and more
Percona Live brings them to you.

- 3 Days
- Hands-on tutorials,
- Breakout sessions,
- Keynote addresses,
- Expo Hall
- Networking
- Lots of Fun!

Use **PRESENTER** for 20% off! Register now at perconalive.com

PERCONA

All the benefits of Percona Server for MySQL, with the MyRocks storage engine

Percona Server for MySQL

- Based on RocksDB key-value store
- Requires less storage space
- Provides more storage endurance
- Ensures better IO capacity
- Available for most popular 64-bit Linux distributions

MyRocks

> “The efficiency improvements in MyRocks make it a gre
> InnoDB. Including it in Percona Server for MySQL makes it possible for the MySQL community to use it. I am thrilled that we worked with Percona to make this possible.”
>
> — Mark Callaghan, MTS, Facebook

PERCONA

InnoDB Encryption

ENCRYPTION_KEY_MAGIC (_V1,_V2,_V3)

~~KEY ID~~ NEW KEY ID

~~UUID~~ NEW UUID

~~ENCRYPTED (TABLESPACE KEY, IV)~~
RE-ENCRYPTED

~~CRC32 OF (TABLESPACE KEY,IV)~~ RE-CALCULATED

Percona Software

Percona XtraDB Cluster

PERCONA XtraDB Cluster

100% open source, free to download and use:

- Cost-effective HA and scalability solution for MySQL

Works on-premises, in the cloud, or hybrid scenario:

- Enterprise ready
- Highly secure
- Provides deep visibility into database performance

Percona Software

Percona XtraDB Cluster

High availability for MySQL

Combines Percona Server 5.7 and Codership Galera Replicator 3.17

ProxySQL load balancer built-in

- Support thousands of concurrent connections, multiplexed to hundreds of backend servers

Improved security

- Percona XtraDB Cluster strict-mode
- Data at rest encryption

Deliver higher performance

- Increased read and write scalability
- Multi-AZ deployment support
- Automatic node provisioning to ease scaling requirements

Percona Software

Percona XtraBackup

PERCONA XtraBackup

Seamless integration into your existing workflow

- Uninterrupted transaction processing during backups
- Fast and reliable database backups with minimal impact

Save on disk space and network bandwidth

- Advanced compression

Validate the integrity of your backup

- Automatic backup verification

Restore your data to any desired time

- Point-in-time recovery

Percona Software

Percona XtraBackup

PERCONA XtraBackup

100% open source, free, database backup solution

- MySQL, Percona Server for MySQL, MariaDB

Works on-premises, in the cloud, or a hybrid

- Enterprise ready
- Simplifies operations by speeding up the addition on new slaves
- Delivers non-blocking backups to minimize impact
- Backup automation enables regular backups and verification

Percona Software

Percona Server for MongoDB

PERCONA
Server for MongoDB

Full drop-in replacement for MongoDB Community Edition

- Fully compatible with MongoDB Community Edition
- 100% open source, free to download and use
- Works on-premises, in the cloud, or a hybrid

Provides advanced security and compliance

- Provides deep visibility into database performance
- Improved efficiency with server consolidation to reduce OPEX
- Improved ROI through lower hosting fees and power usage

Percona Software

Percona Server for MongoDB

PERCONA
Server for MongoDB

Enhanced security with binary log and data-at-rest encryption

Full support for transactions

Enterprise ready, with free enterprise features

- Plug-in authentication and auditing functionality
- WiredTiger, MMAPv1 and Percona Memory Engine storage engines
  - Percona Memory Engine for in-memory computing workloads is equivalent to proprietary MongoDB Enterprise in-memory engine
- Integrated open source hot backup system for WiredTiger

Percona Software

Percona Server for MongoDB

PERCONA Server for MongoDB

The only MongoDB variant with storage solutions for all workloads

- Traditional OLTP workloads with WiredTiger
- In-memory computing workloads with Percona Memory Engine

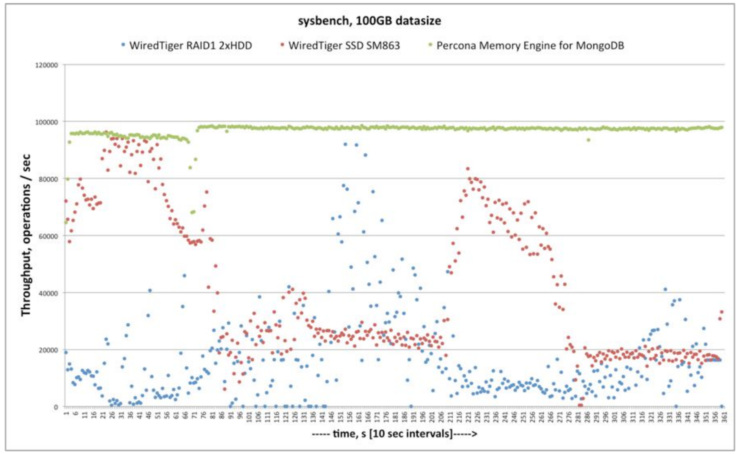

Percona Software

Percona Monitoring and Management

Percona Monitoring and Management (PMM) is a single pane of glass to help manage complex database environments in public, private or on-premises environments.

Designed to help DBAs and developers **gain deep insight into their applications and databases**, PMM is used by thousands of organizations around the globe to manage complex database environments.

PMM is an **award-winning database monitoring tool built** by Percona, the database performance and scalability experts, using best-of-breed tools.

Percona Software

Percona Monitoring and Management

Customer Story

PERCONA
Monitoring and Management

Keep your revenue engine up and running. With PMM, you can keep your databases running smoothly and continuously, with consistent end-user experience for applications. Easily find, fix, and prevent scaling issues, bottlenecks, and potential outages.

Spend less time managing complex environments. Enable developers and DBAs to be able to view and monitor complex environments with multi-databases, multiple technologies, and multiple providers.

Percona Software

Percona Monitoring and Management

Customer Story

PERCONA
Monitoring and Management

Speed up development. PMM creates a common language between DBA’s, developers, and sysadmins to help speed development and release cycles. With PMM, high-quality releases won’t negatively impact performance, scale, or security.

Percona Monitoring and Management helps improve the quality of your releases and applications by identifying bottlenecks and issues and helps you deal with problems easily and efficiently.

Percona Software

Percona Kubernetes Operators

PERCONA
Kubernetes Operator
FOR XTRADB CLUSTER

Percona Kubernetes Operator for Percona XtraDB Cluster

- Deploy easily
- Scale your Percona XtraDB Cluster
- Automate Your Backups
- Integrate with Percona Monitoring and Management (PMM)
- Rely on ProxySQL to Remove Single Point of Failure
- Automate node recovery
- Provide data encryption
- Support private data registries

Percona Software

Percona Kubernetes Operators

Percona Kubernetes Operator for Percona Server for MongoDB

- Deploy easily
- Scale Your Replica Set
- Add Monitoring
- Manage your Backups
- Set Node as Arbiter
- Automate node recovery
- Provide data encryption
- Support private data registries

Percona Software

Percona Toolkit

PERCONA Toolkit

Simplify operations – save time and resources

- Complex tasks are scripted
- Locate potential issues before they impact your environment

Alter your environment with little to no user impact

- On-line schema change tool

Perform complex tasks with ease and reliable repetition

- Archiver tool

Percona Software

Percona Toolkit

PERCONA Toolkit

100% open source, free command-line tools

- Developed and used by Percona experts

Works on-premises, in the cloud, or a hybrid

- Enterprise ready
- Full customization allows you to alter the tools to meet your specific needs
- Supports Percona Server for MySQL, MariaDB, MySQL, Percona XtraDB Cluster, Percona Server for MongoDB, and MongoDB

Percona Software