



PERCONA  
TRAINING

# Introduction to MongoDB

<http://www.percona.com/training/>

# Why MongoDB

---

## What is MongoDB?

"MongoDB is an open-source document database that provides high performance, high availability, and automatic scaling."

# Fun Facts

---

- MongoDB was initially created in 2008 as part of a hosted application stack
- The company was originally called 10gen
- As part of their overarching plan to create the 10gen platform, the company built a database
- Suddenly everybody said: "I like that! Give me that database!"

# Why Mongo

---

## Where does it fit?

- You need automated HA without using 3rd party tools
- You need to quickly prototype and make structural changes to your data layout
- You need rich data storage to be fast and flexible
- Source data is JSON
- You want out-of-the-box sharding support

# Why Mongo

---

## What is amazing?

- Each document is atomic so very few mutex ranges needed -> potentially faster performance
- Nesting sub-documents unlike N tables relationship in SQL
- High availability and sharding out-of-the-box
- Capped collections
- TTL indexes

# Why Mongo

---

## What is concerning?

- No enforcing of data structure
- No good rollback (multi-document) or even transactions until 4.0
- Too many indexes (even 10-20) can kill MongoDB performance
- Loss of relational integrity (FK)
- Custom query language

# Original goals for MongoDB

---

- IoT, Cloud, and Distributed DB
- Scaling out/horizontally vs. up/vertically
- Ease application development and release cycle
- Complex/extended data types to aid rapid development

# MongoDB flavors

---

- MongoDB Community
- MongoDB Enterprise
  - adds LDAP, Kerberos, encryption and auditing
- Percona Server for MongoDB
  - same features as Enterprise but no license fee



# Percona Server for MongoDB

---

- Optional Percona Memory Engine storage engine
- Data at rest encryption
- External LDAP authentication & authorization
- Audit logging
  - Track and query database interactions of users
- Hot Backup for WiredTiger
- Profiling Rate Limit
  - Decrease the impact of the profiler on performance



PERCONA  
TRAINING

# Questions