

MongoDB Replica Sets, EC2 Hosting, and Random Tomfoolery

Cody Powell
Cofounder, Famigo
@codypo

Intro

- MongoDB's database replication is done via replica sets
 - New feature, only added in v1.6 (we're at 1.8.1 now)
- Previously had master/slave db replication, like MySQL
- What are replica sets?
 - Asynchronous master/slave replication on steroids

How do Replica Sets work?

- A replica set has 1 primary node and ≥ 1 secondary node
- Each node has a special oplog collection, an ordered list of all writes performed on the data
- All writes occur on the primary node, and quickly appear in the primary's oplog
- Secondary nodes listen to primary's oplog, and apply all new operations in the right order

What makes Replica Sets great?

- Automatic failover
 - When primary node goes down, new primary node is elected within seconds
- Automatic recovery of member nodes
 - When secondary node falls behind, it catches up by iterating through the primary's oplog
- Simply scale your reads
 - All writes occur on primary, reads can occur on any node

What does a Replica Set look like?

- A replica set should have at least 3 nodes
- Why 3?
 - Primary node is determined by node voting, you want an odd number to prevent weirdness that occurs from tied votes.
- 3 node replica sets can be:
 - 1 primary node, 2 secondary nodes
 - 1 primary, 1 secondary, 1 arbiter
- An arbiter is a lightweight node that only participates in voting

MongoDB and EC2

- Amazon's EC2 service is a simple, cheap way to host MongoDB
- Full MongoDB instances (replica set or no) should be 64-bit
 - MongoDB uses memory-mapped files for performance
 - Thus, 32 bit instances are limited to storing 2.5 GB of data
- Arbiter instances can be 32-bit, though

Replica Set Deployment Model for EC2

- 2 Large Nodes - 1 Primary, 1 Secondary
 - \$0.34 per hour, each
 - Instances don't have to be dedicated to MongoDB, can host lots of other stuff too
- 1 Micro Node - Arbiter
 - Freeeeeeeeeeeeeeeeeeee!

Demo!

3 Nodes:

- mongo1, small EC2 instance
- mongo2, small EC2 instance
- mongo-arb, micro EC2 instance

These are all 32 bit micro Amazon Linux instances.

EC2 Hints:

- Default hostnames are hellaciously long, use shortcuts via /etc/hosts entries, pointing to your internal IPs
- Your replica set configs should reference these hostnames
- Ensure appropriate ports are open