



# SOLR Power FTW

## 10 EZ Steps to SOLR Domination

Alex Pinkin

@apinkin

# Alex Pinkin

- Director of Engineering,  
Data and Infrastructure,  
Bazaarvoice



@apinkin

# Bazaarvoice

- SaaS company  
powering user generated content  
on thousands of web sites
- 5+ billion page views per month
- 300+ billion impressions




# SOLR Stats

<b>Documents</b>	<b>350 MM</b>
<b>Index size</b>	<b>330 GB</b>
<b>AVG QPS</b>	<b>3,400</b>
<b>MAX QPS</b>	<b>12,500</b>
<b>AVG Response time</b>	<b>7 ms</b>
<b>95P Response time</b>	<b>14 ms</b>

# Initial State - back in 2007...

- SoR: MySQL
- Search: MySQL
- Numerous indexes
- Numerous aggregates



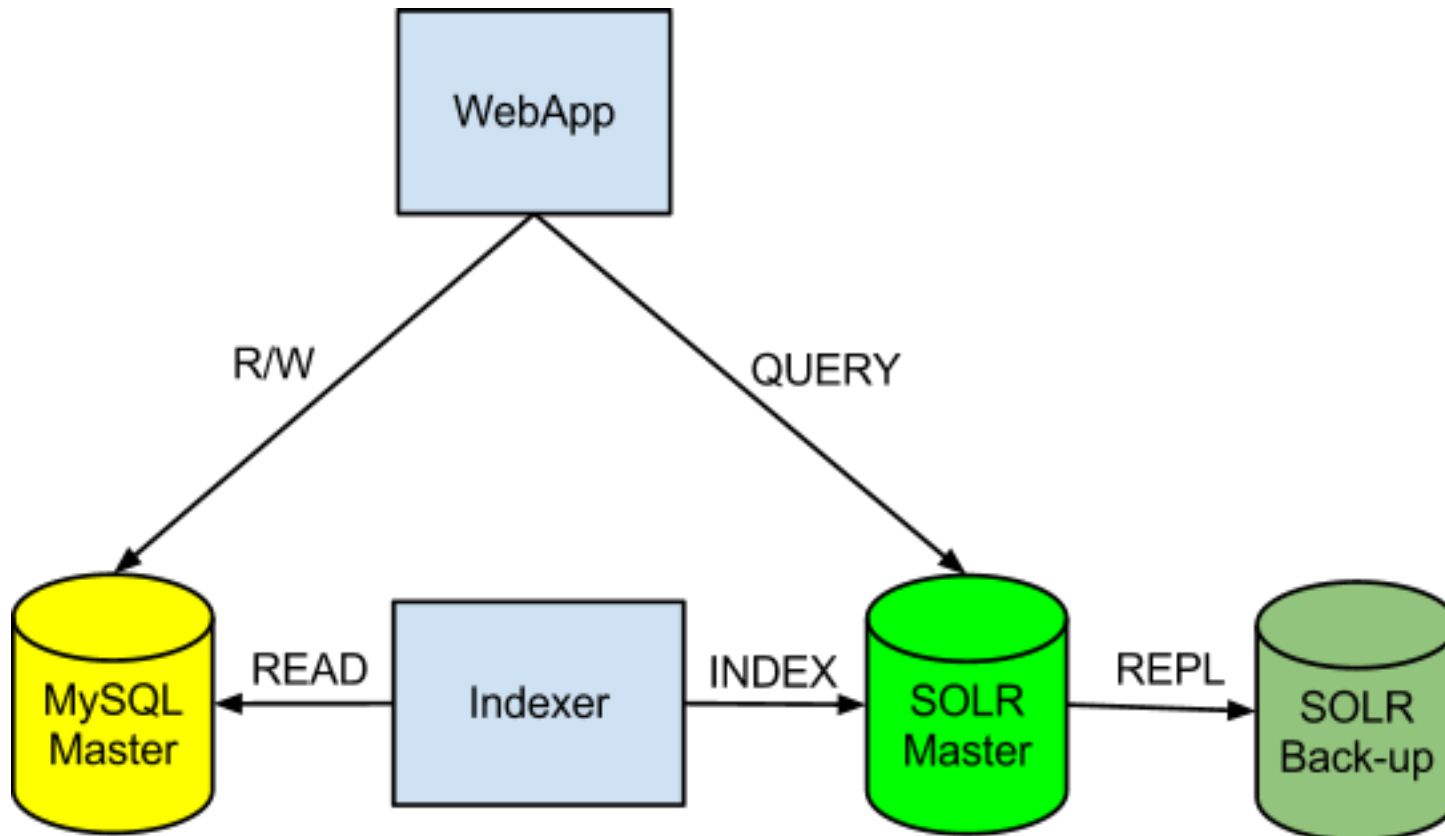
# Problem 1

FULL TEXT SEARCH



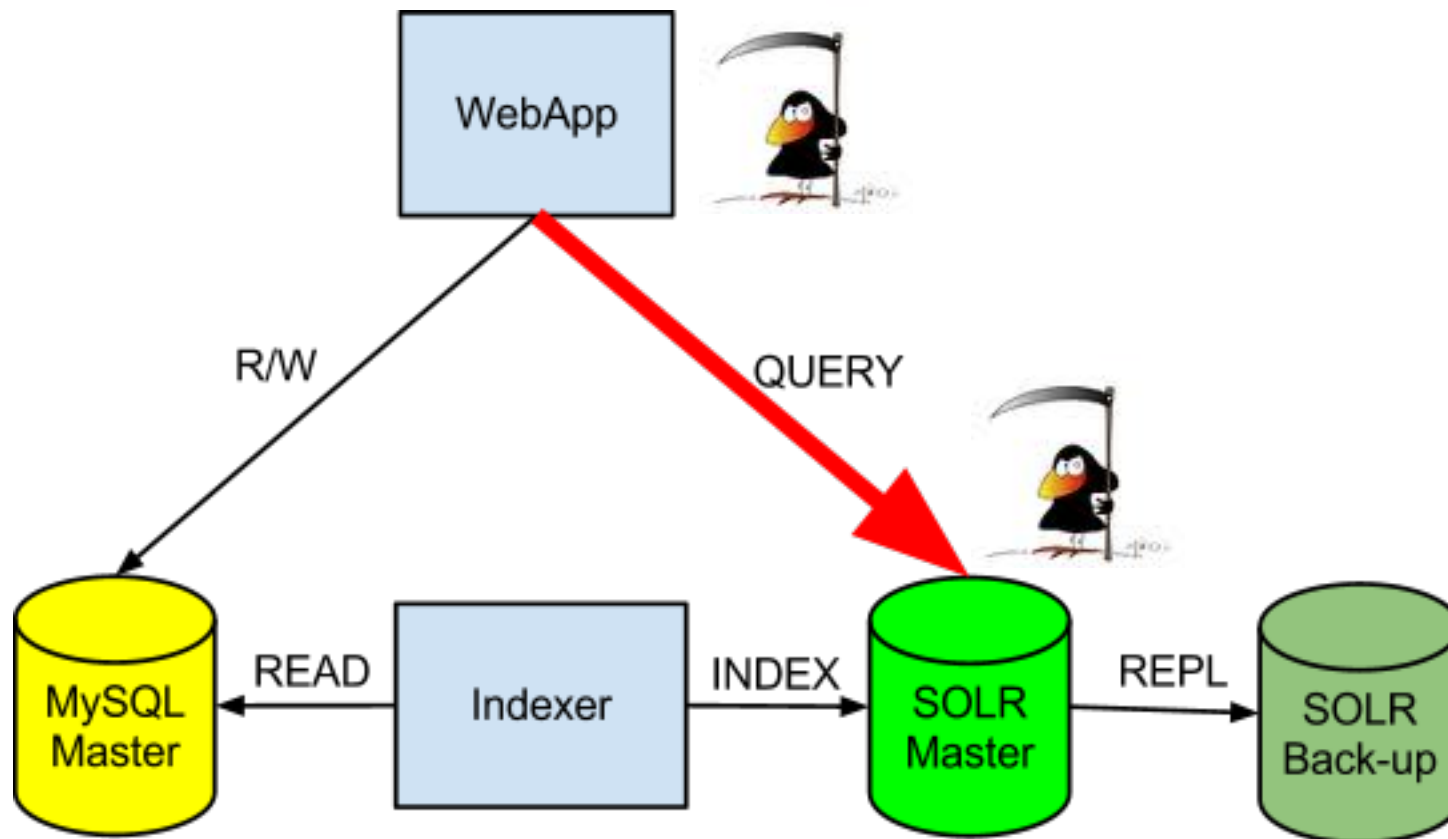
# Solution 1: FT Search

- SOLR Master + Hot Backup
- Single index



# Problem 2

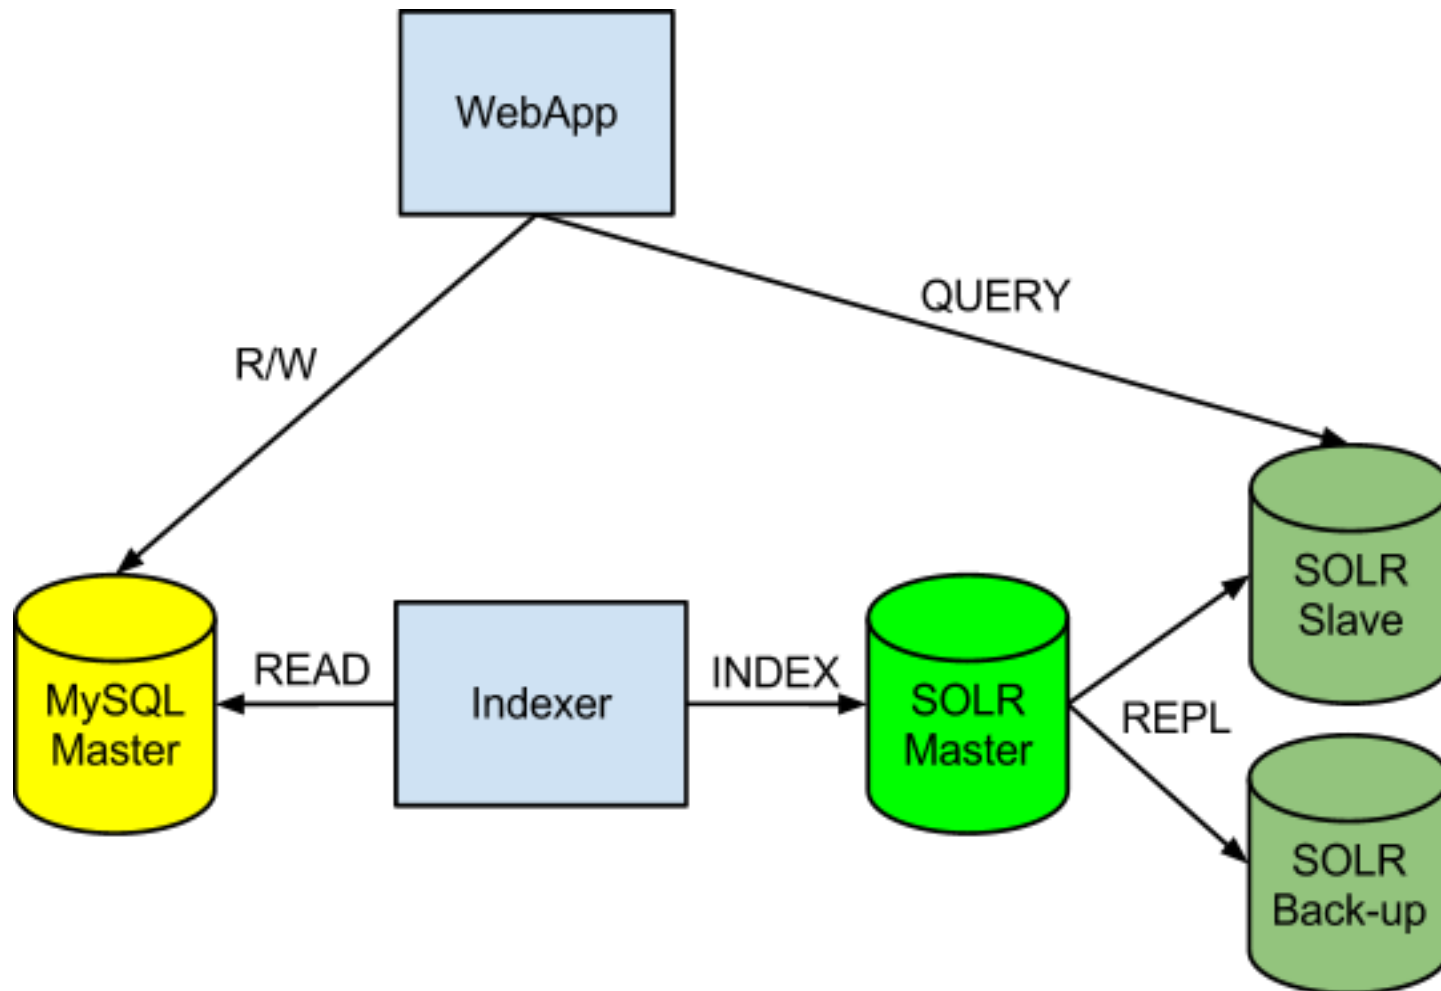
- Reads overwhelming master
- No R/W isolation





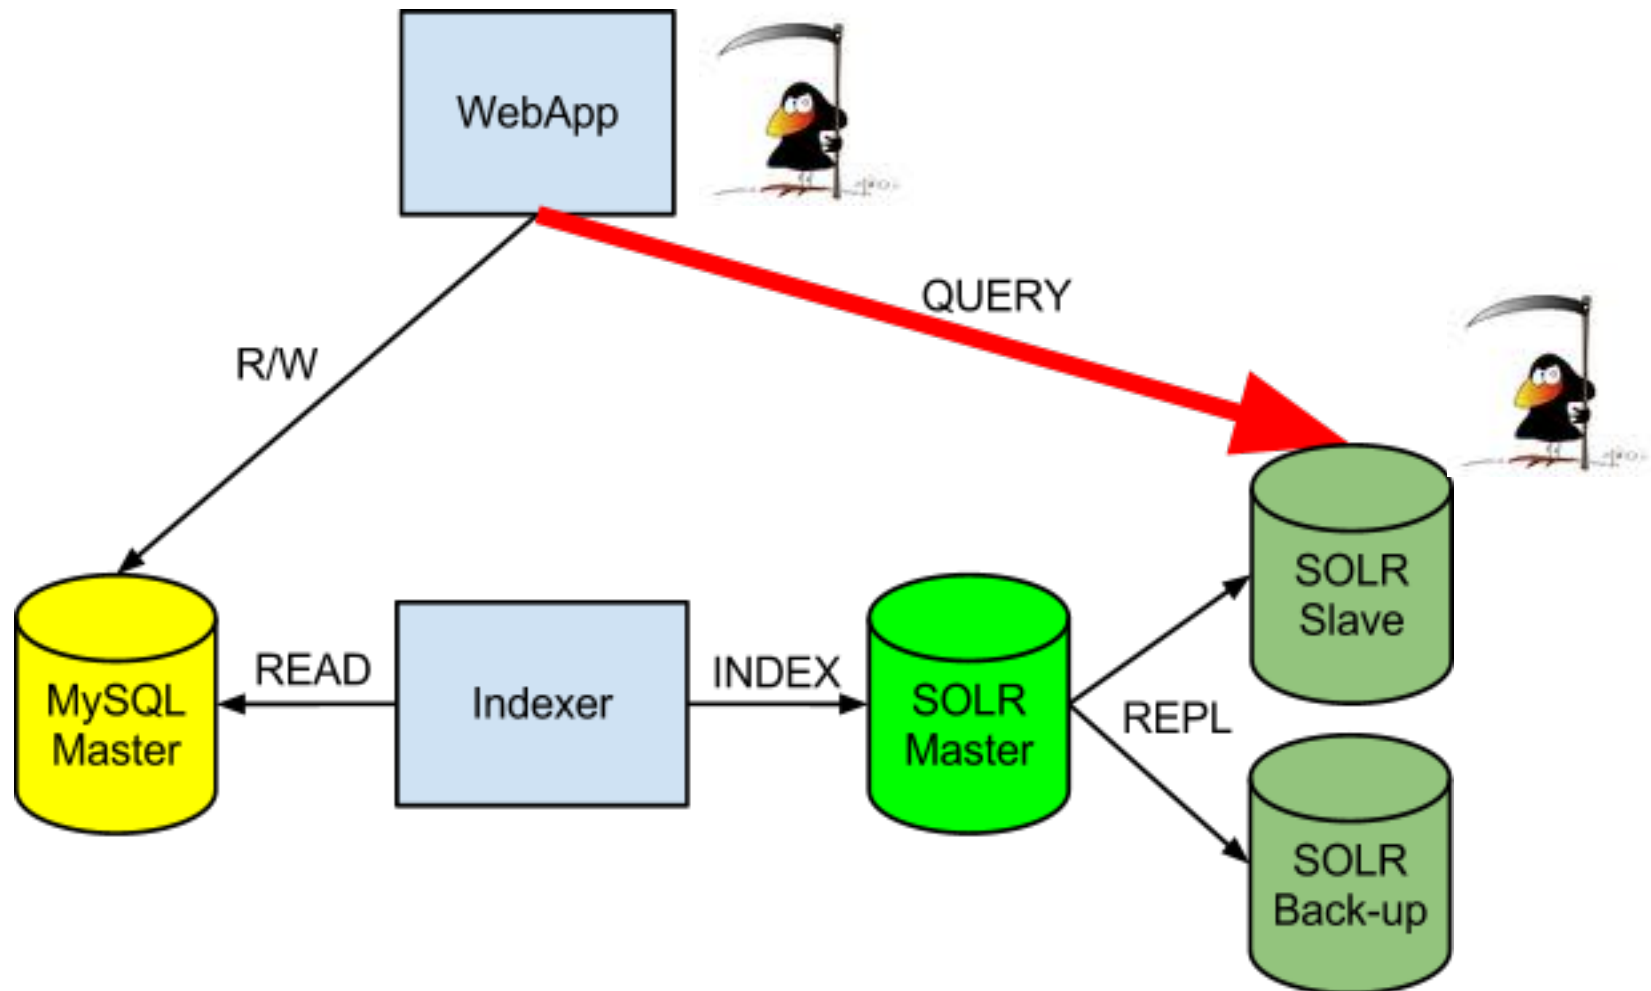
# Solution 2

- R/W separation



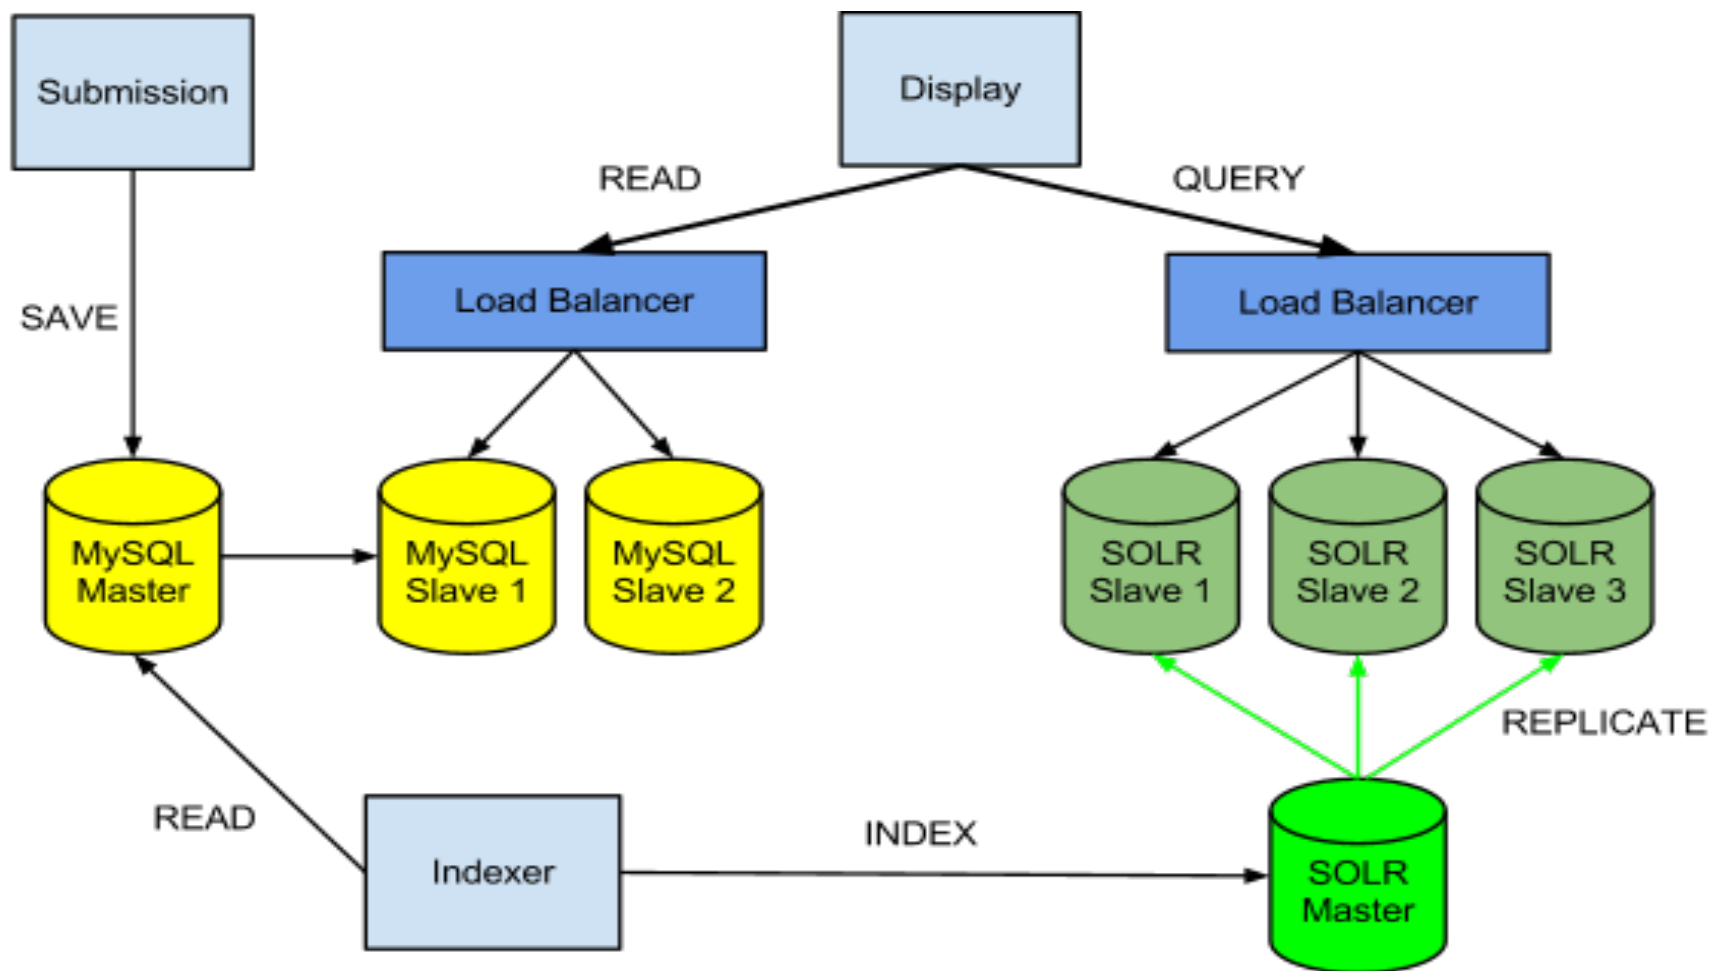
# Problem 3

- QPS Growing



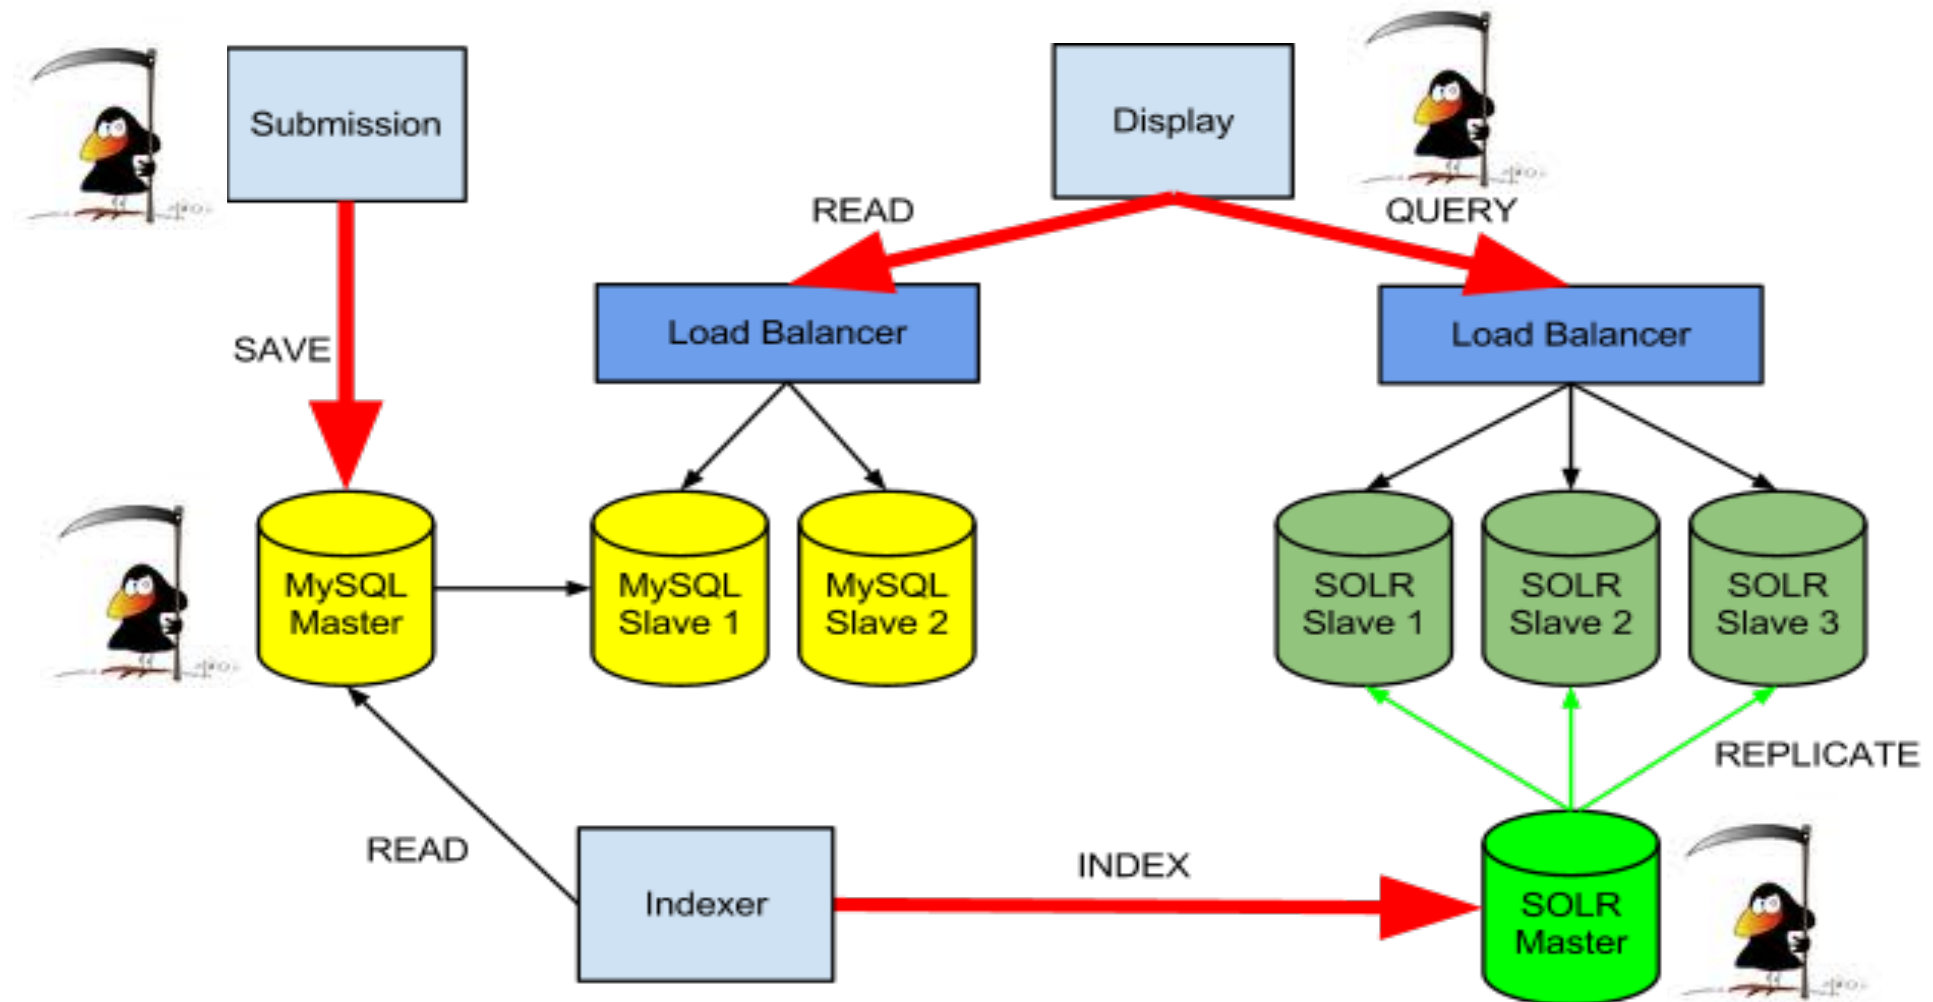
# Solution 3

- Add more slaves





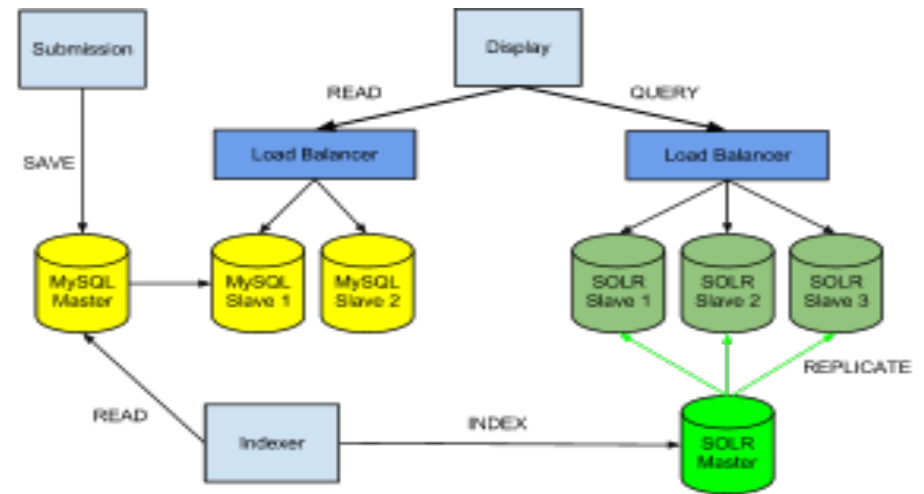
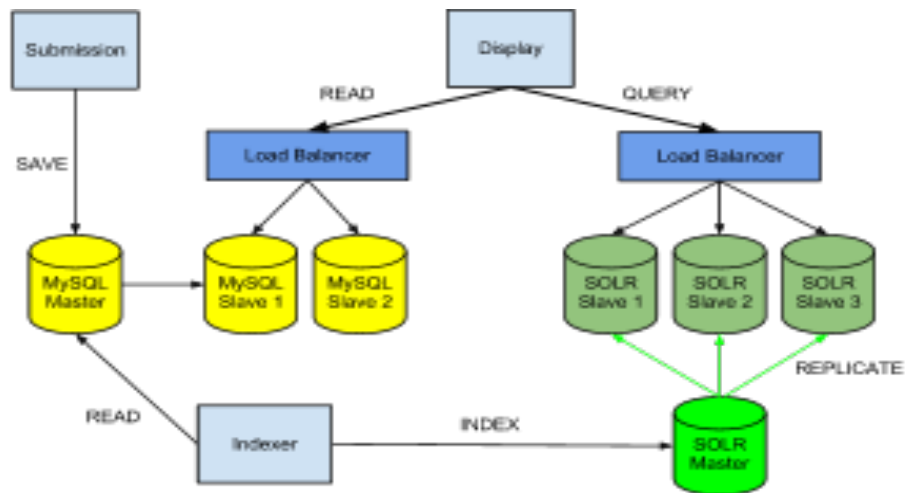
# Problem 4

- Number of Clients is growing



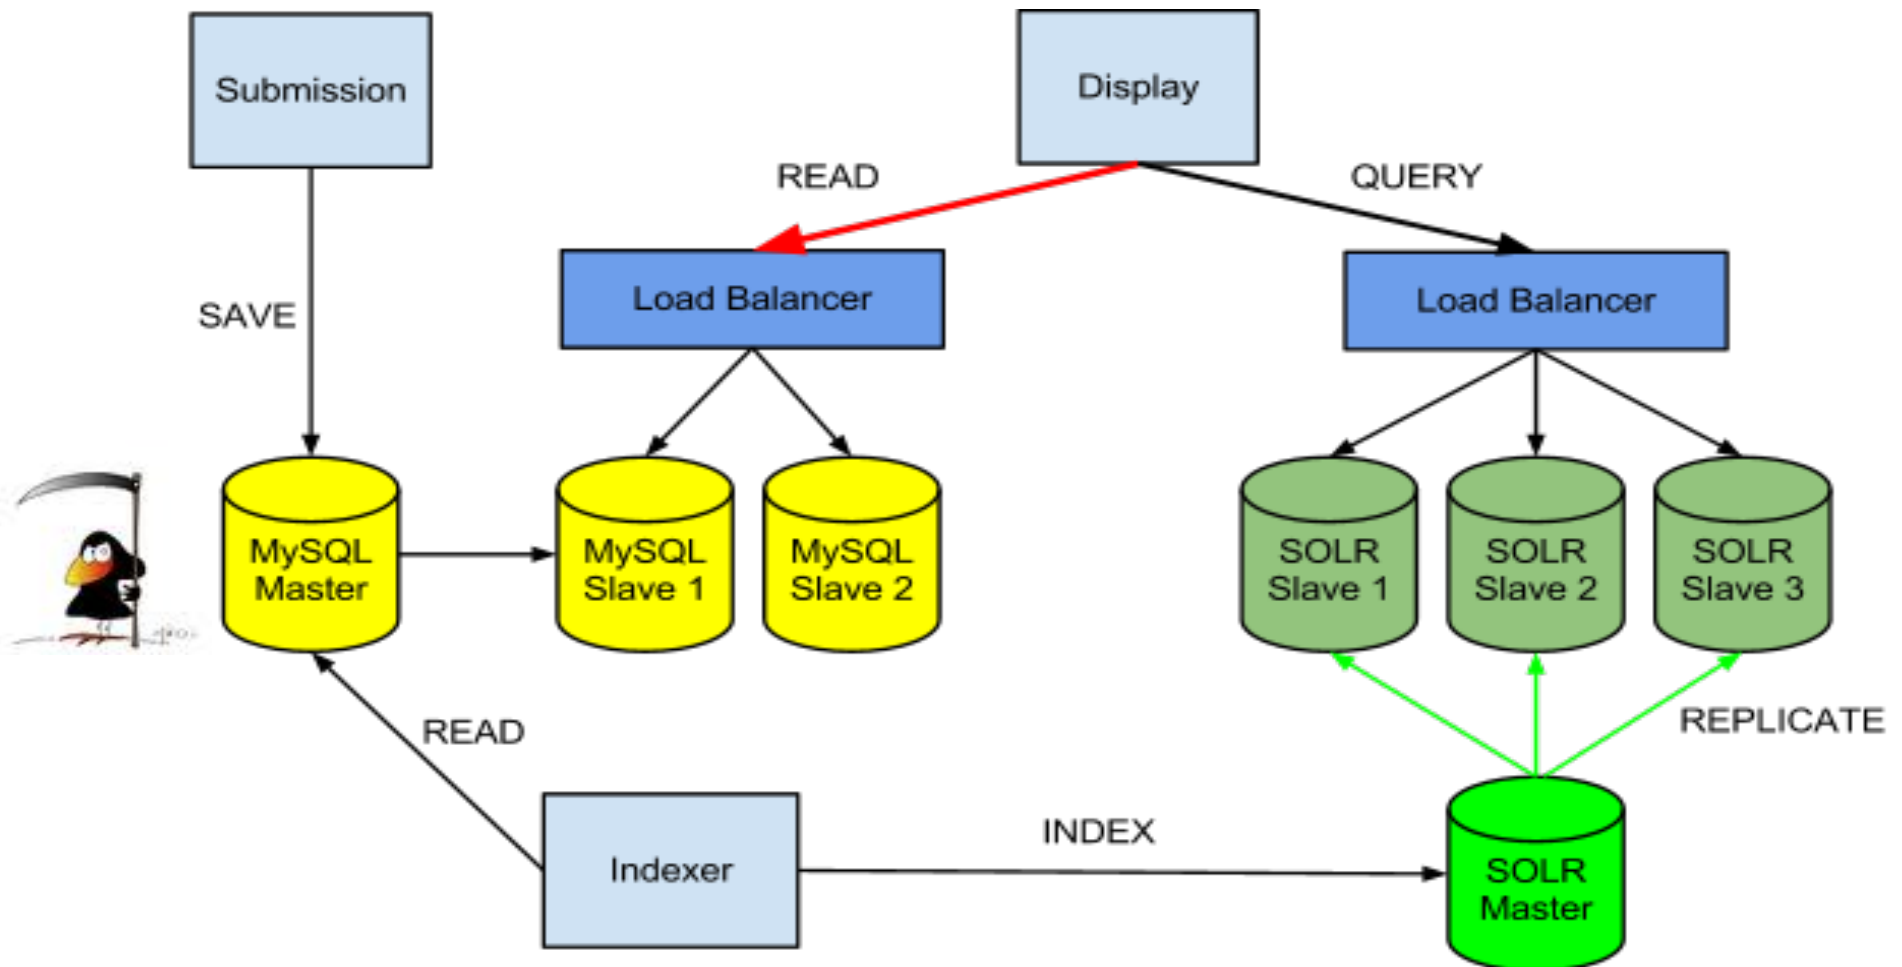
# Solution 4

- Multiple Deployment Groups



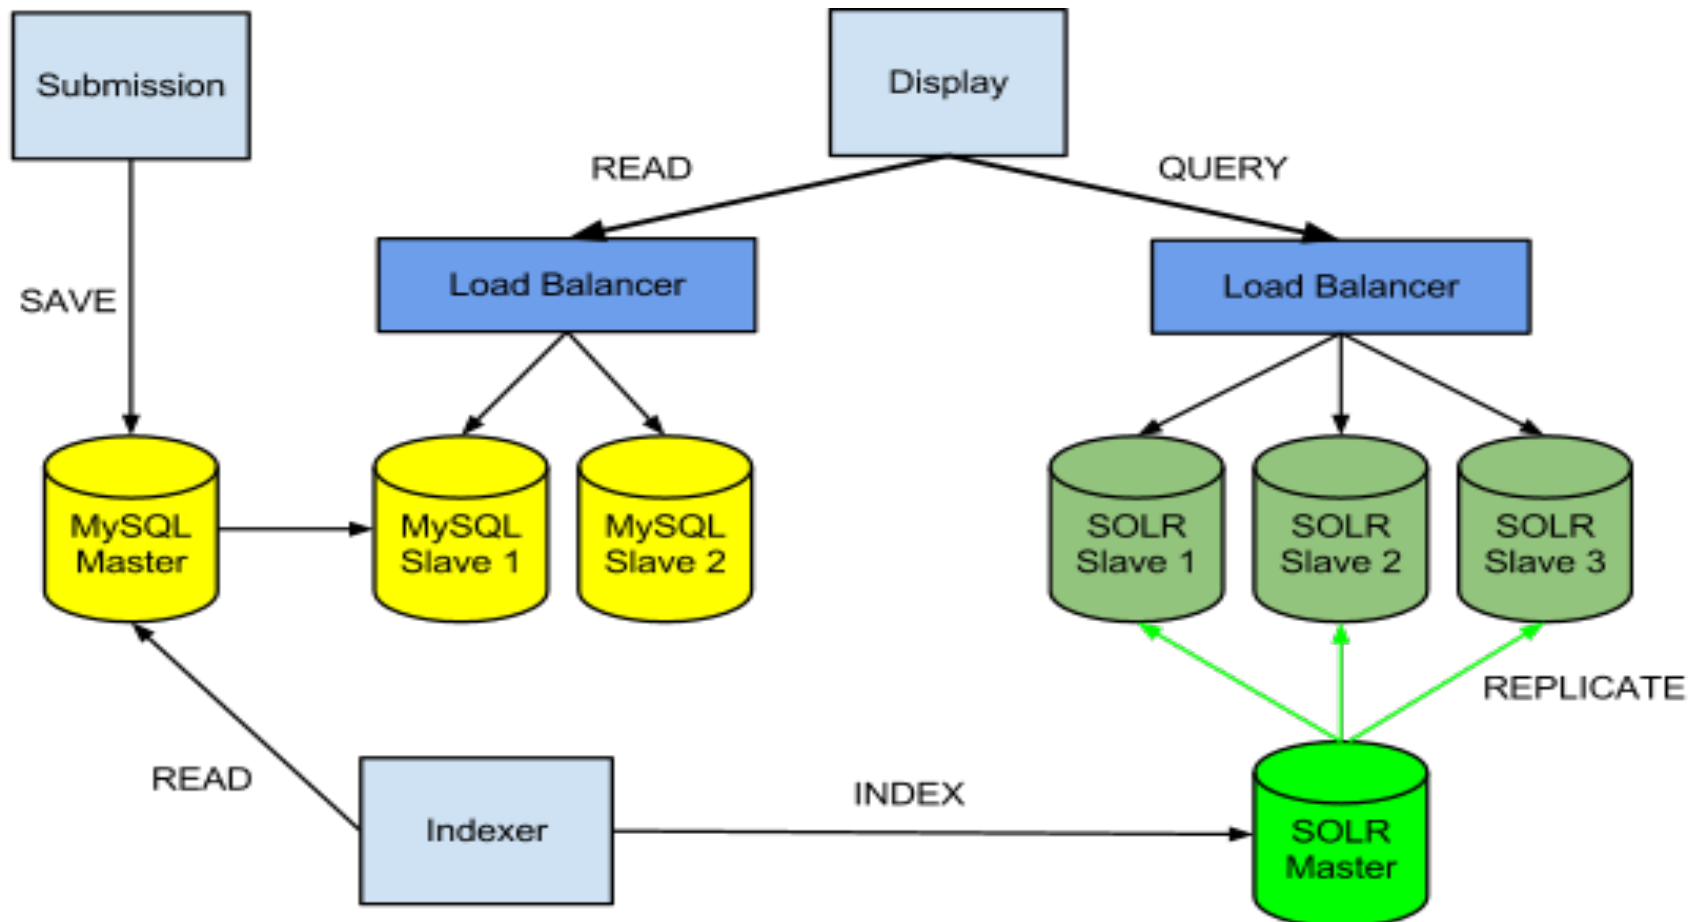
# Problem 5

- MySQL still in trouble: Too many Indexes, Aggregates



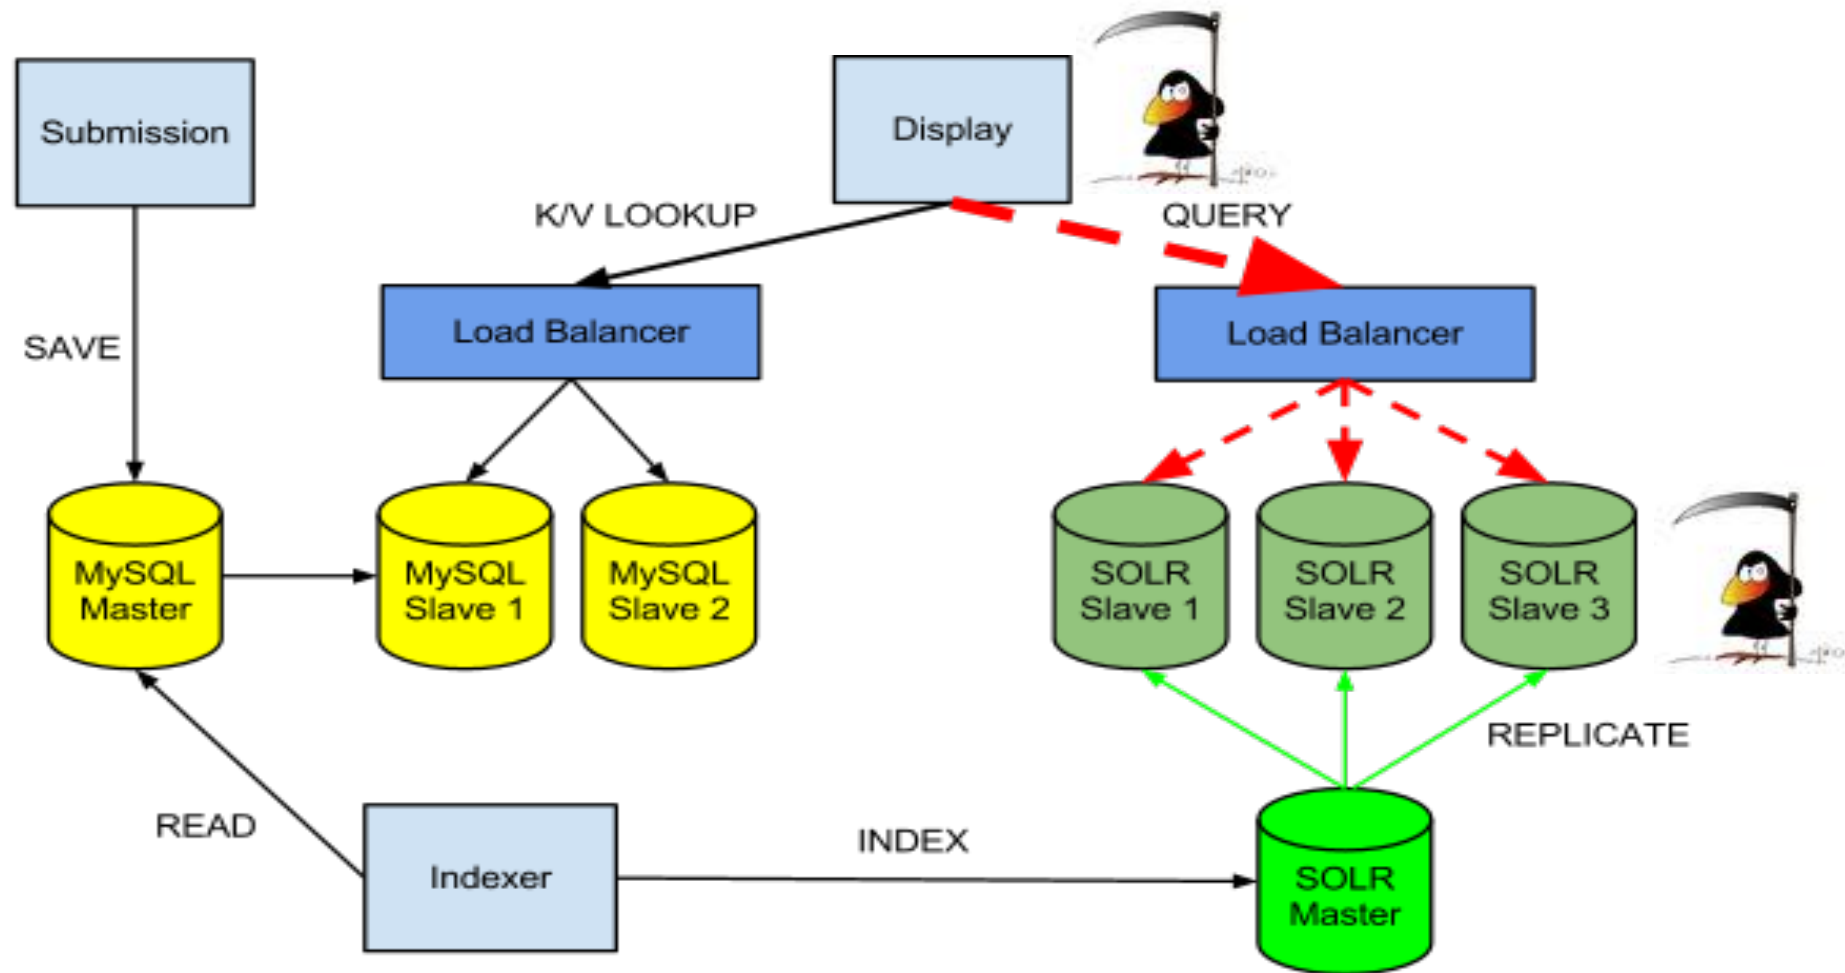
# Solution 5

- Offload filtering and sorting to SOLR
- Add more slaves
- Add more RAM



# Problem 6

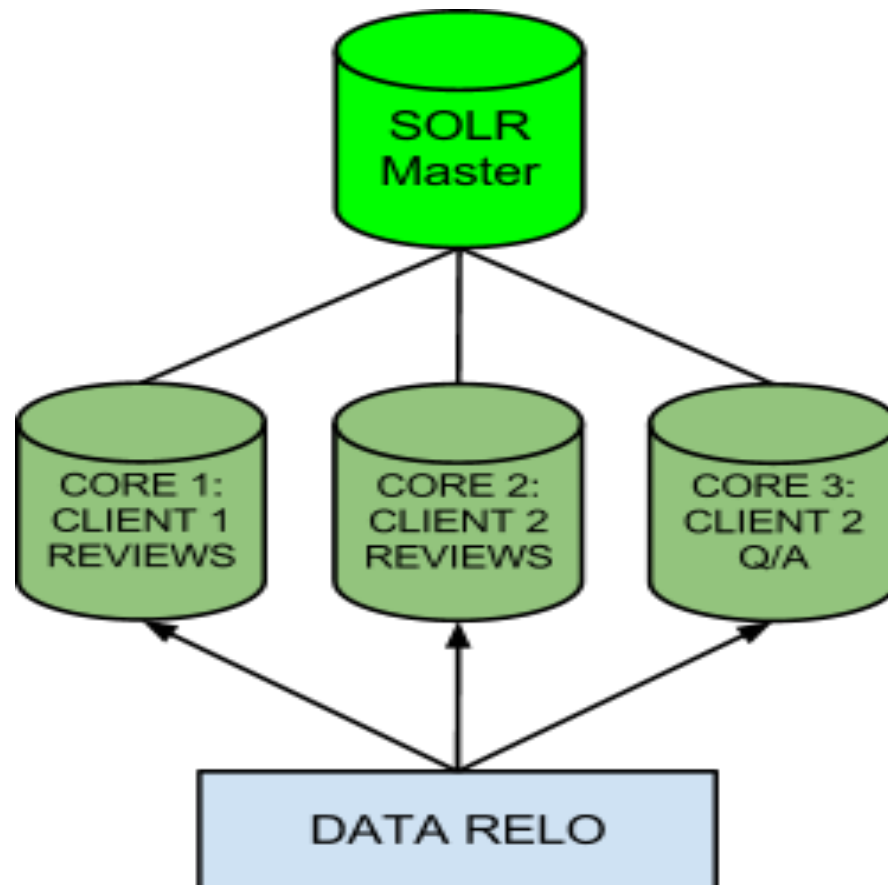
- Query latency is getting worse





# Solution 6

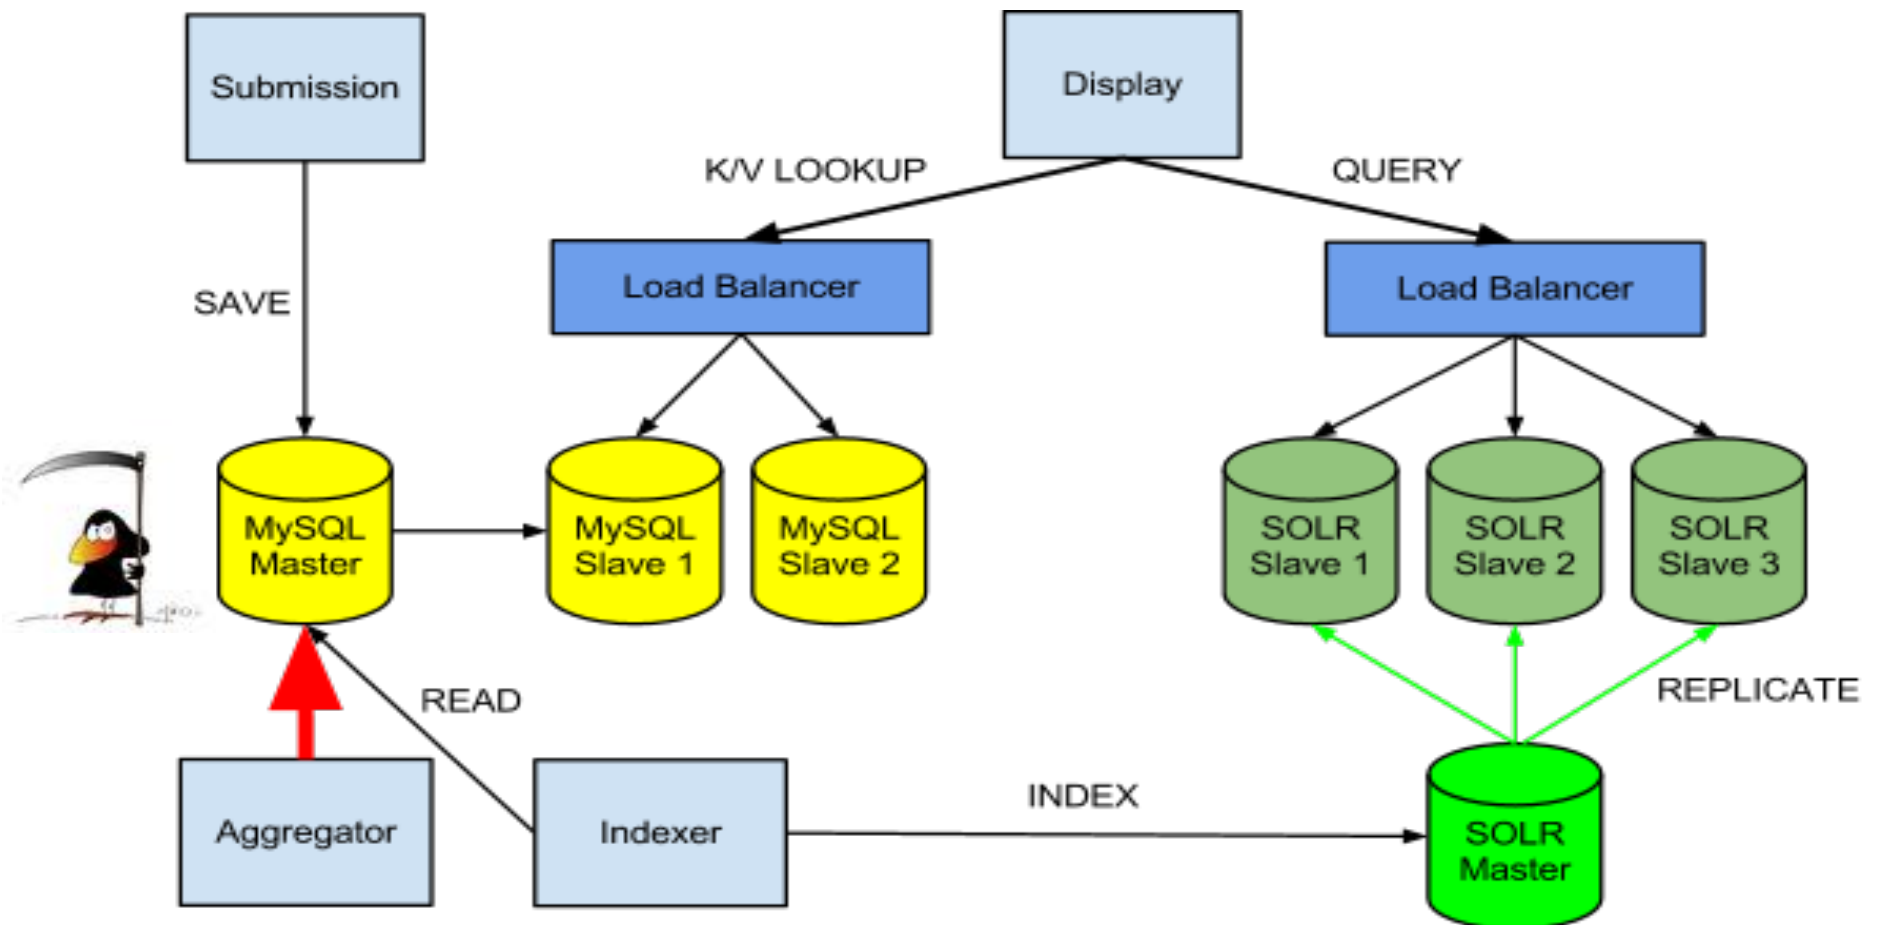
- Multiple cores
- Core management tools
- Transparent data migration



# Problem 7

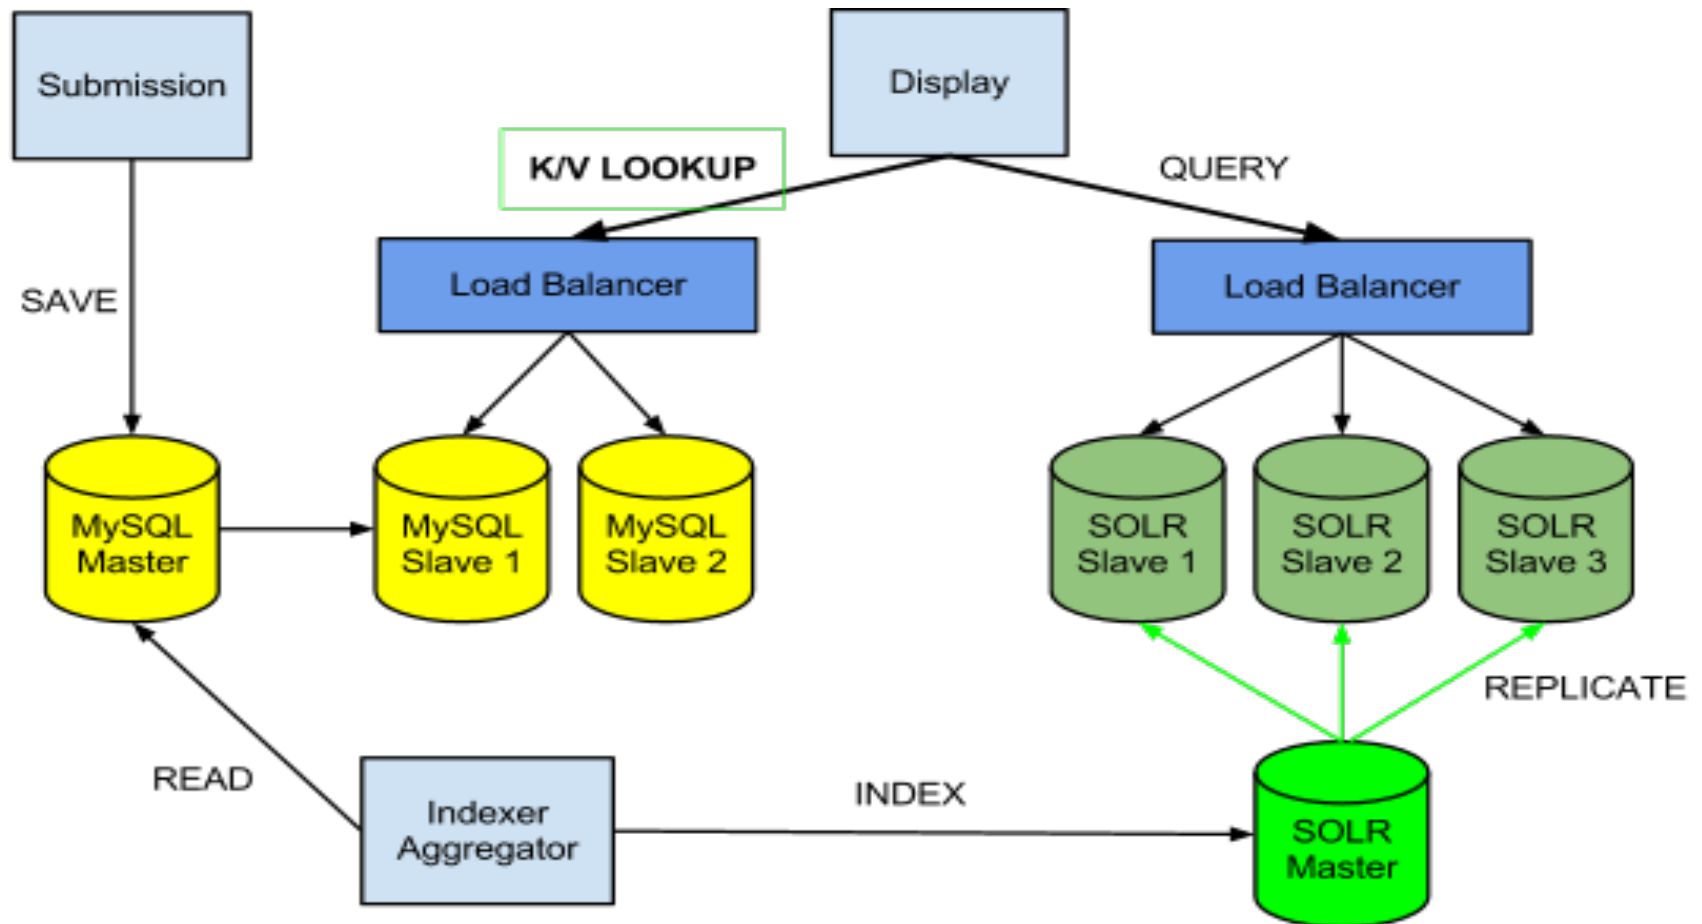
MySQL is still in trouble

- Too many aggregate tables
- Error-prone aggregation logic



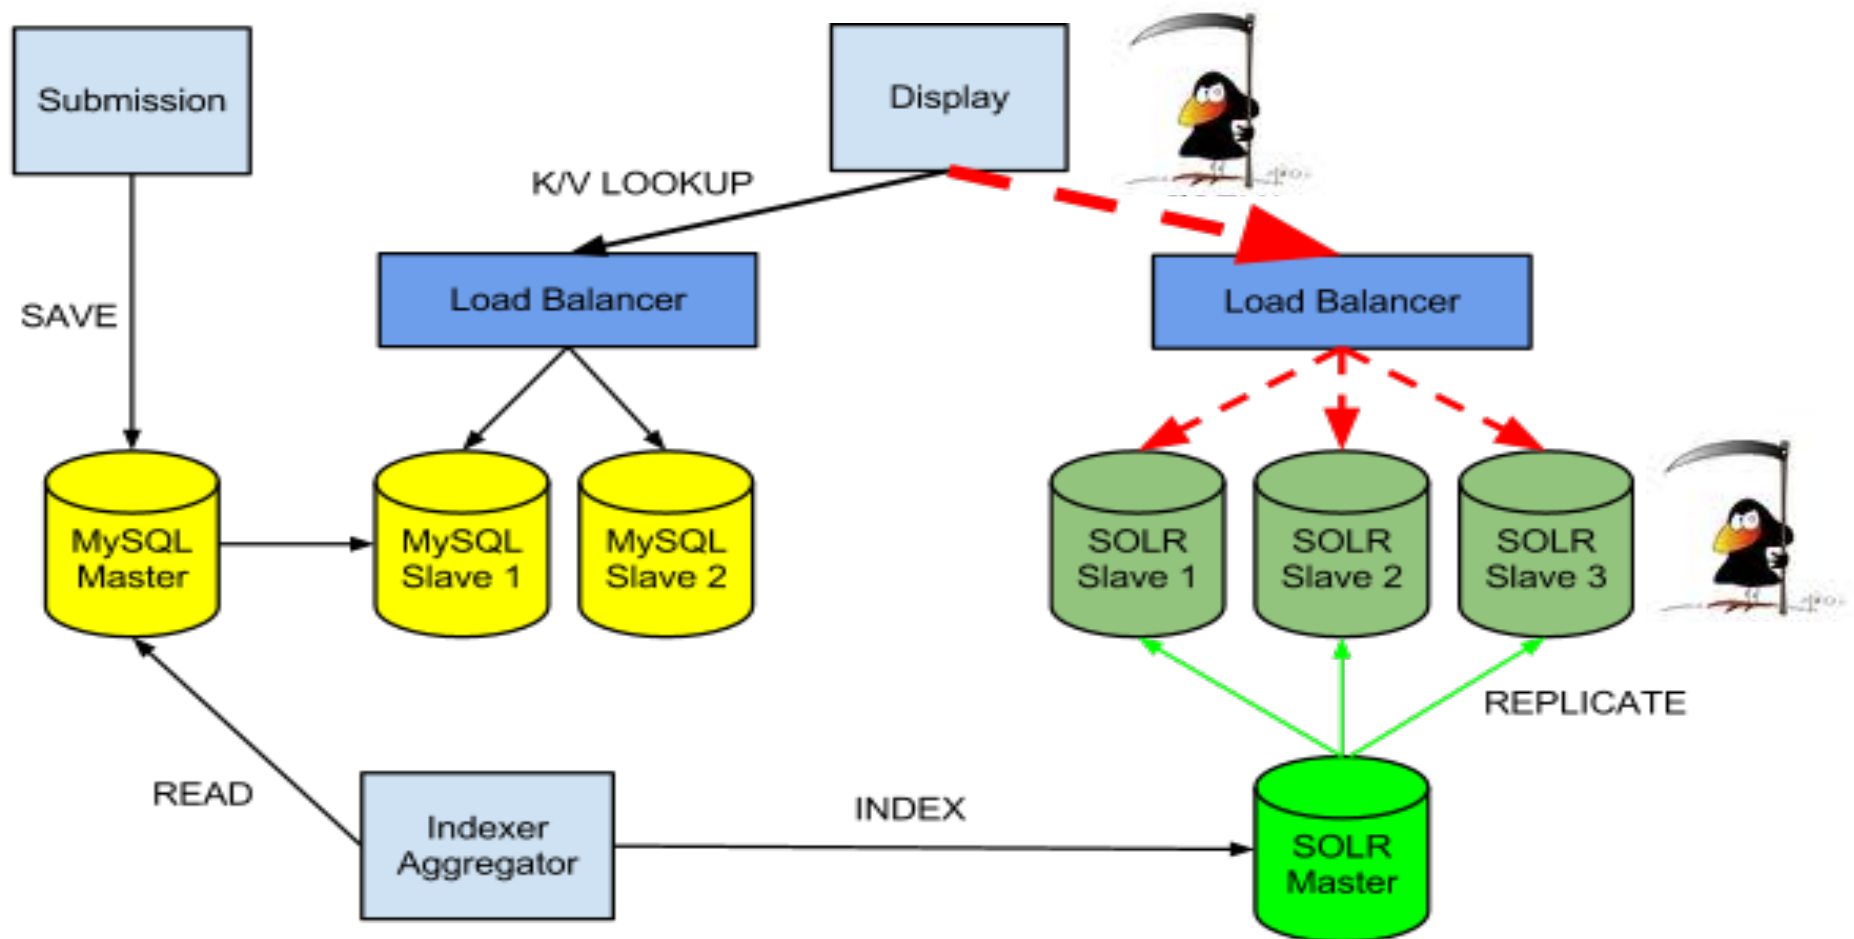
# Solution 7

- Off-load simple aggregates to SOLR
- Use Faceting



# Problem 8

- SOLR is GC'ing itself to death
- 99p Latency = :-(



# Solution 8

- GC tuning
  - Concurrent Low Pause Collector
  - Parallel Young Generation Collector
  - Incremental mode



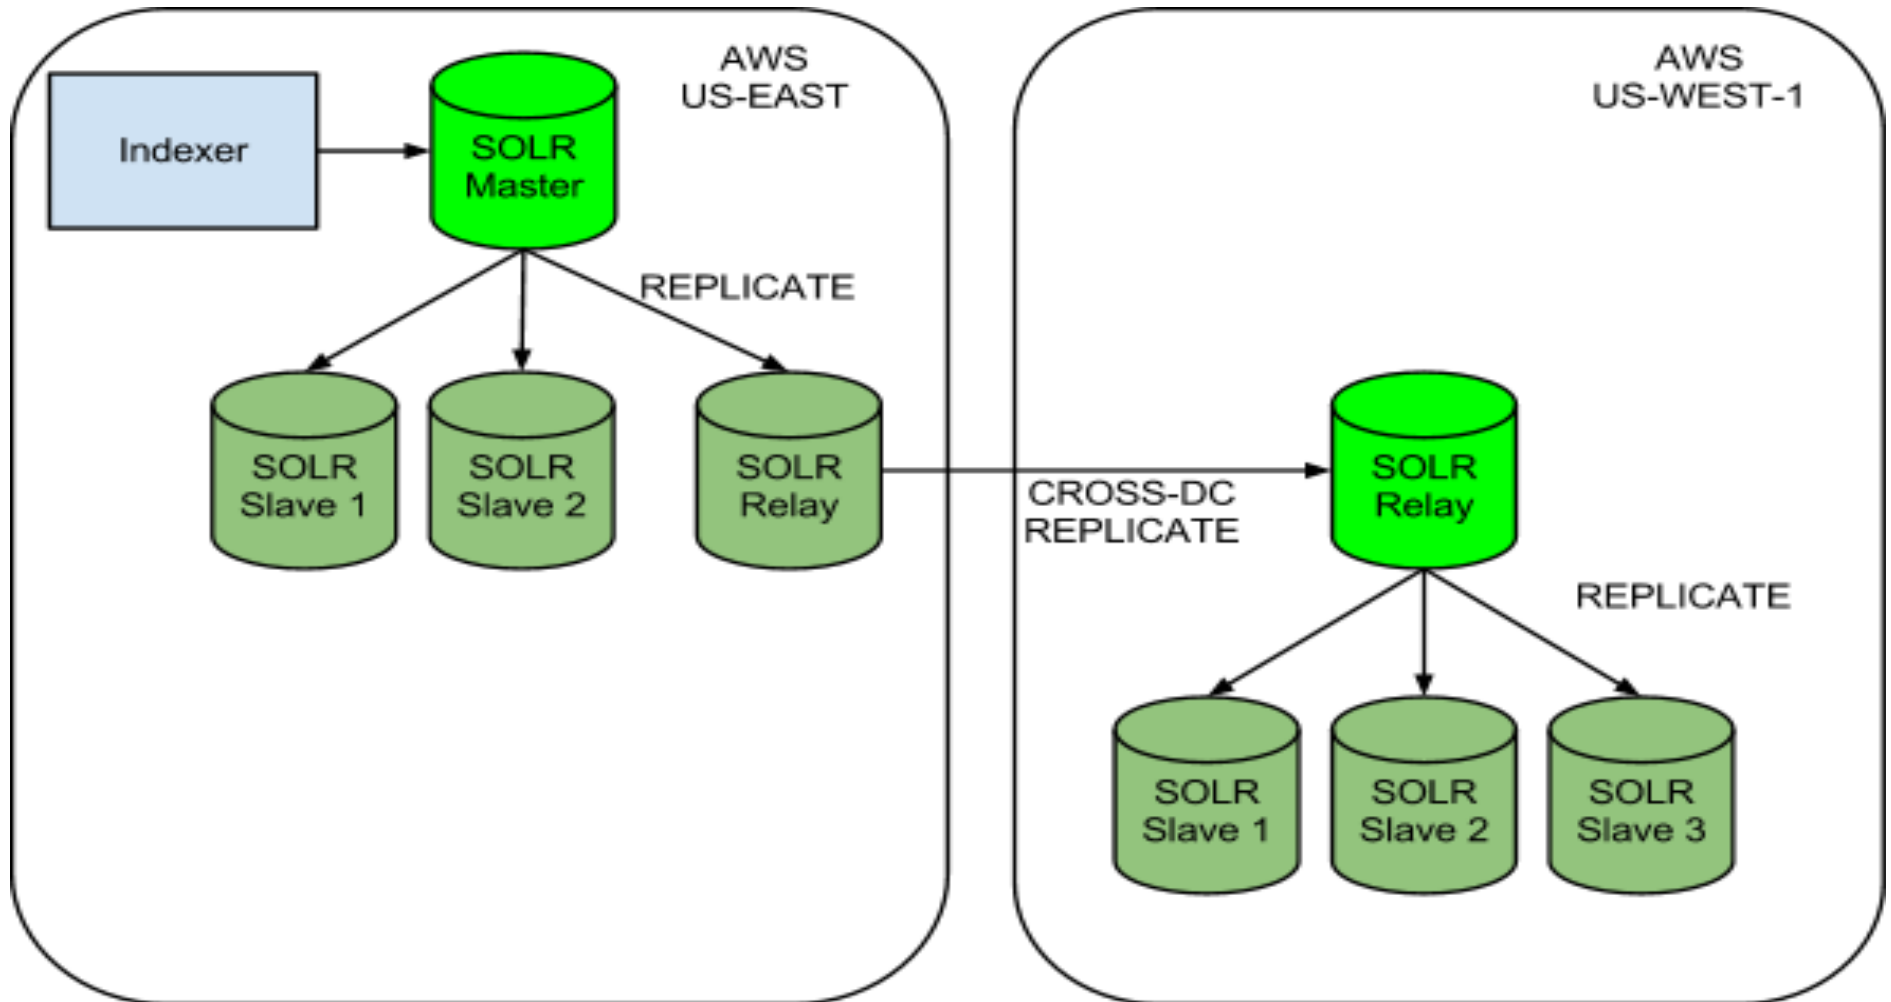
# Problem 9

- Disaster Recovery (Multi-DC)



# Solution 9

- Cross-DC replication



# Problem 10

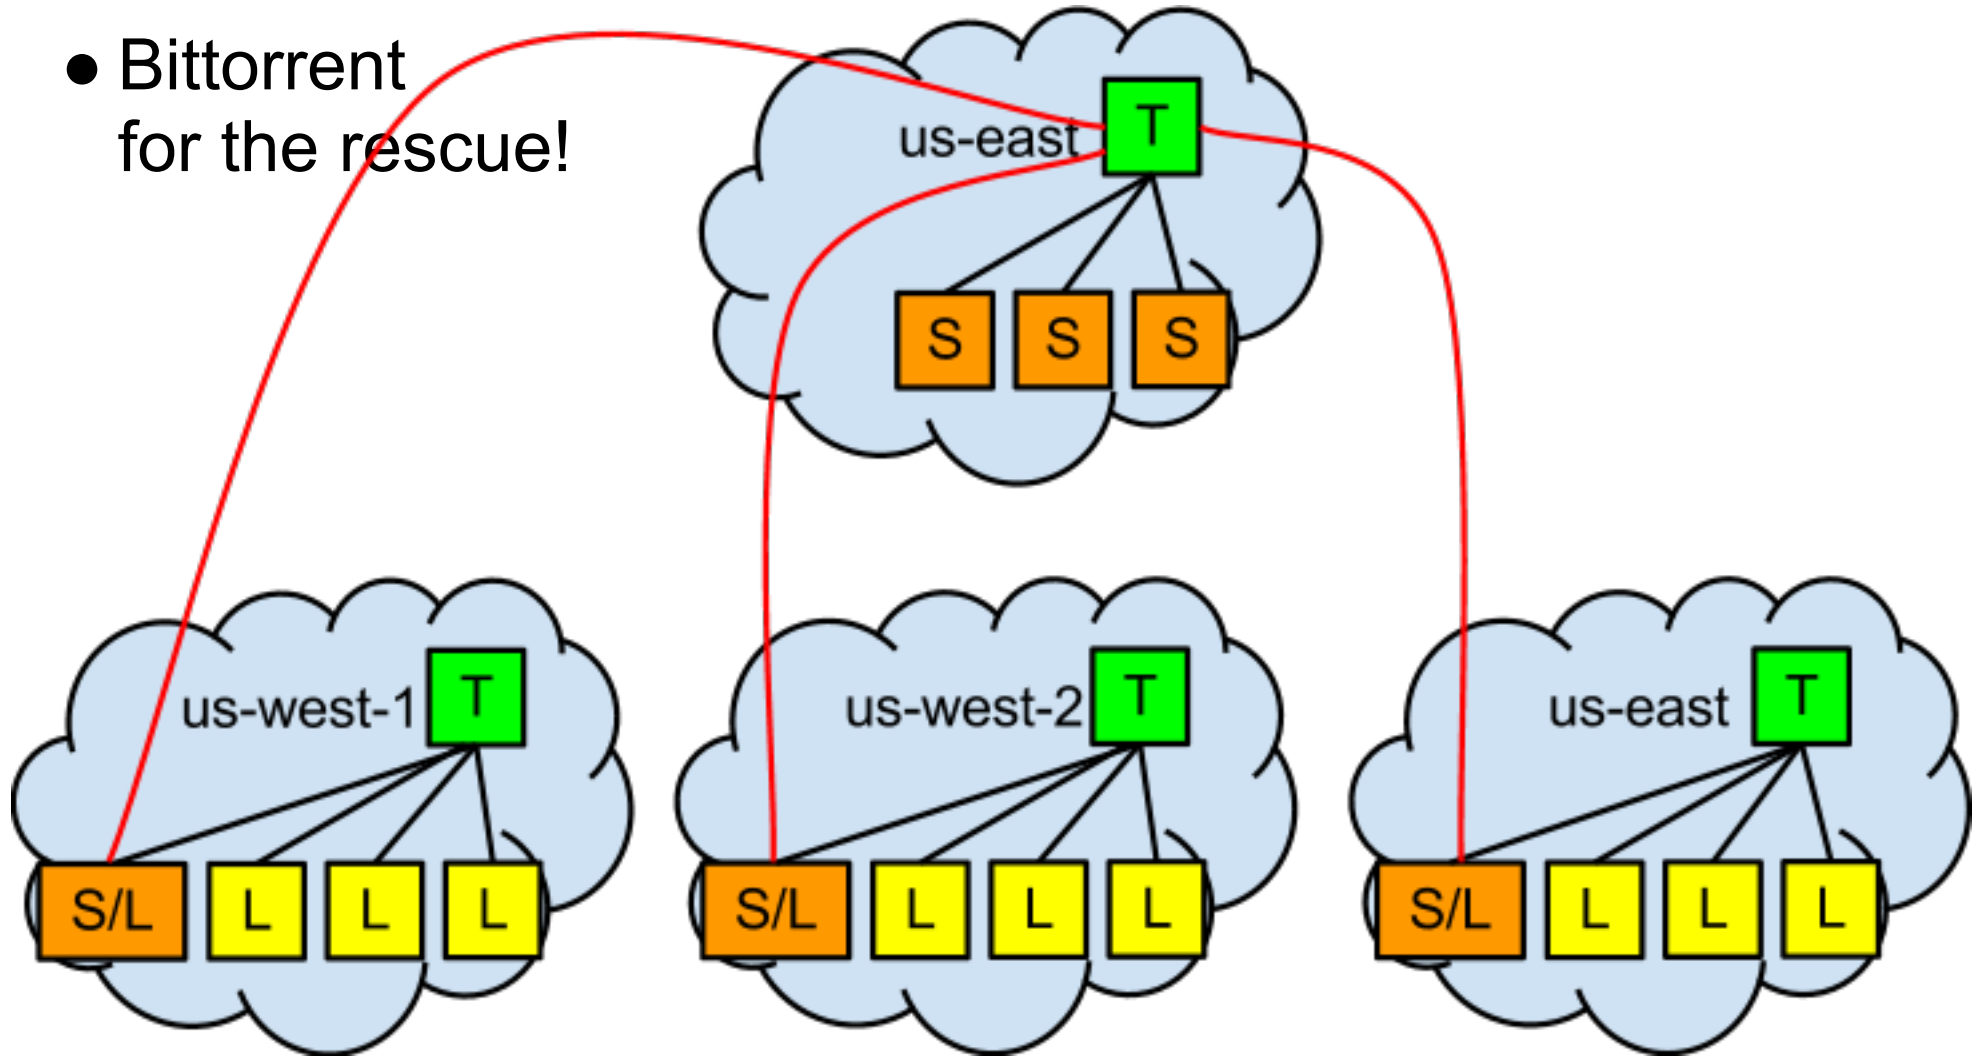
- Schema Changes...

"ALTER TABLE OF DOOM" @josh\_wills



# Solution 10

- Bittorrent for the rescue!



# Remaining Challenges

- Elasticity
- HA: failover
- Schema changes

# Questions



@apinkin