

# Database Availability and Integrity in NoSQL

**Fahri Firdausillah**  
**[M031010012]**

# What is NoSQL

- Stands for **Not Only SQL**
- Mostly addressing some of the points: **non-relational, distributed, horizontal scalable, schema-free, easy replication support, eventually consistent**, and **huge data amount**
- This presentation will talk much about **replication** and **horizontal scalable** for database availability, then **eventually consistent** and **schema-free** for data integrity



# List of NoSQL Products

- **Cassandra** used on:
  - Digg, Facebook, Twitter, Reddit, Rackspace, Cloudkick, Cisco
- **Hadoop** used on:
  - Amazon Web Services, Pentaho, Yahoo!, The New York Times
- **CouchDB** used on:
  - CERN, BBC, Interactive Mediums
- **MongoDB** used on:
  - Foursquare, bit.ly, SourceForge, Fotopedia, Joomla Ads
- **Riak** used on:
  - Widescript, Western Communications, Ask Sponsored Listings

# Database Availability Outline

- Database Availability Means
- CAP Theorem (BASE vs ACID)
- Partitioning and Replication
- Replication Diagram
- “Ring” of Consistent Hashing
- Next .... → Database Integrity



# What is NoSQL

- Stands for **Not Only SQL**
- Mostly addressing some of the points: **non-relational, distributed, horizontal scalable, schema-free, easy replication support, eventually consistent, and huge data amount**
- This presentation will talk much about **horizontal scalable** for database availability and **eventually consistent** for data base integrity



# What is NoSQL

- Stands for **Not Only SQL**
- Mostly addressing some of the points: **non-relational, distributed, horizontal scalable, schema-free, easy replication support, eventually consistent, and huge data amount**
- This presentation will talk much about **horizontal scalable** for database availability and **eventually consistent** for data base integrity



# Database Availability Mean


- IBM divide database availability into 3 section:
  - **High Availability:** database and application is available in scheduled period, when maintenance period system is temporarily down.
  - **Continuous Operation:** system available all the time with no scheduled outages.
  - **Continuous Availability:** combination of HA & CO, data is always available, and maintenance is done without shutdown the system

\* Database Availability Considerations. IBM RedBook [2001]

# CAP Theorem (1)

## **Consistency, Availability and Partition Tolerance.**

A shared-data system can have at most two of those three.



\* *Visual Guide to NoSQL Systems*, Nathan Hurst, 2010 <sup>[8]</sup>



# CAP Theorem (2)

- **Consistent, Available (CA) Systems** have trouble with partitions and typically deal with it with replication.
- **Consistent, Partition-Tolerant (CP) Systems** have trouble with availability while keeping data.
- **Available, Partition-Tolerant (AP) Systems** achieve “eventual consistency” through replication and verification consistent across partitioned nodes.

# ACID and BASE

- **ACID**

**A**tomicity: All or nothing

**C**onsistency: Any transaction should result in valid tables

**I**solation: separate transactions

**D**urability: Database will survive a system failures.



# ACID and BASE cont'd

- **BASE**

**B**asically **A**vailable - system seems to work all the time

- **S**oft State - it doesn't have to be consistent all the time

- **E**ventually Consistent - becomes consistent at some later time

# Horizontal Scale

- Data explosion (especially in web application) force database system to scale

- **1st solution : Vertical scale**

Improving server specification by adding more processor, RAM, and storage device. Limited and expensive.

- **2<sup>nd</sup> solution : Horizontal scale**

Adding more cheap computer as server expansion. Do sharding and partitioning which is hard to implement and expensive using relational databases (RDBMS)

# Partitioning & Replication


- Partitioning

- Sharing the data between different nodes (data host)
- Each node placed on a ring
- Advantage : ability to scale incrementally
- Issues : non-uniform data distribution

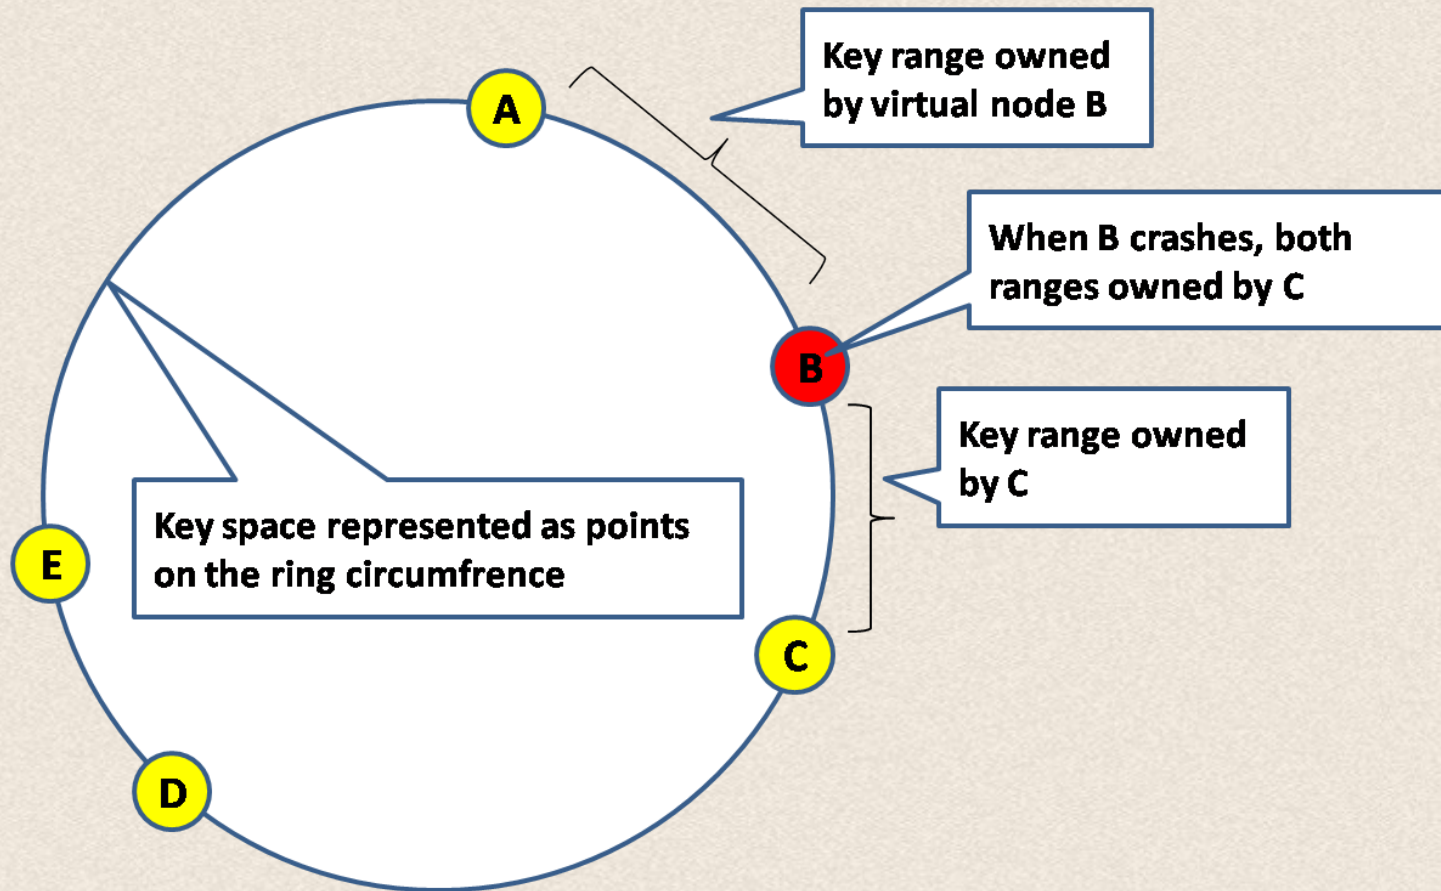
- Replication

- Multiple nodes
- Multiple datacenters
- High availability and durability

# Data Replication

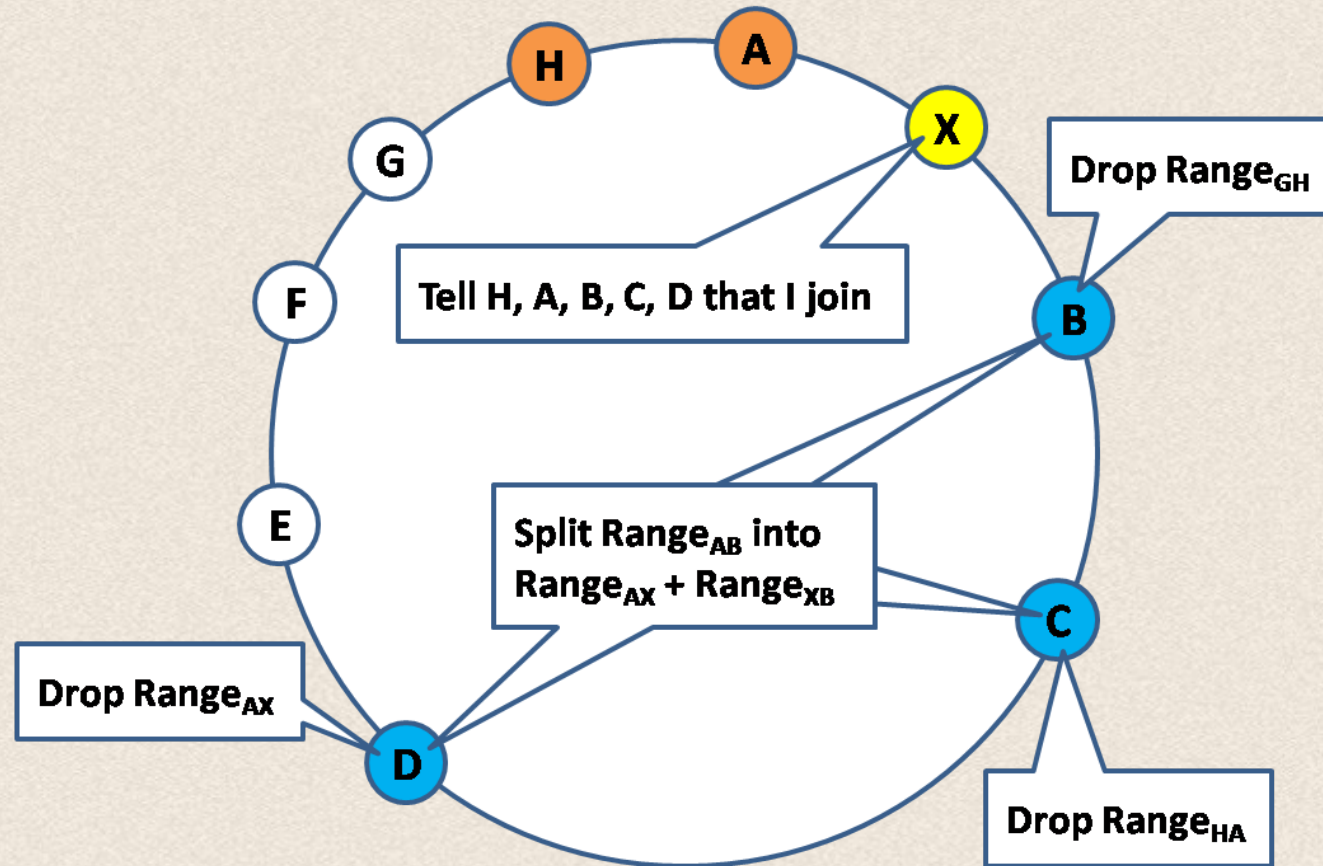


# Ring of Consistent Hashing



# When a New Node Join Network

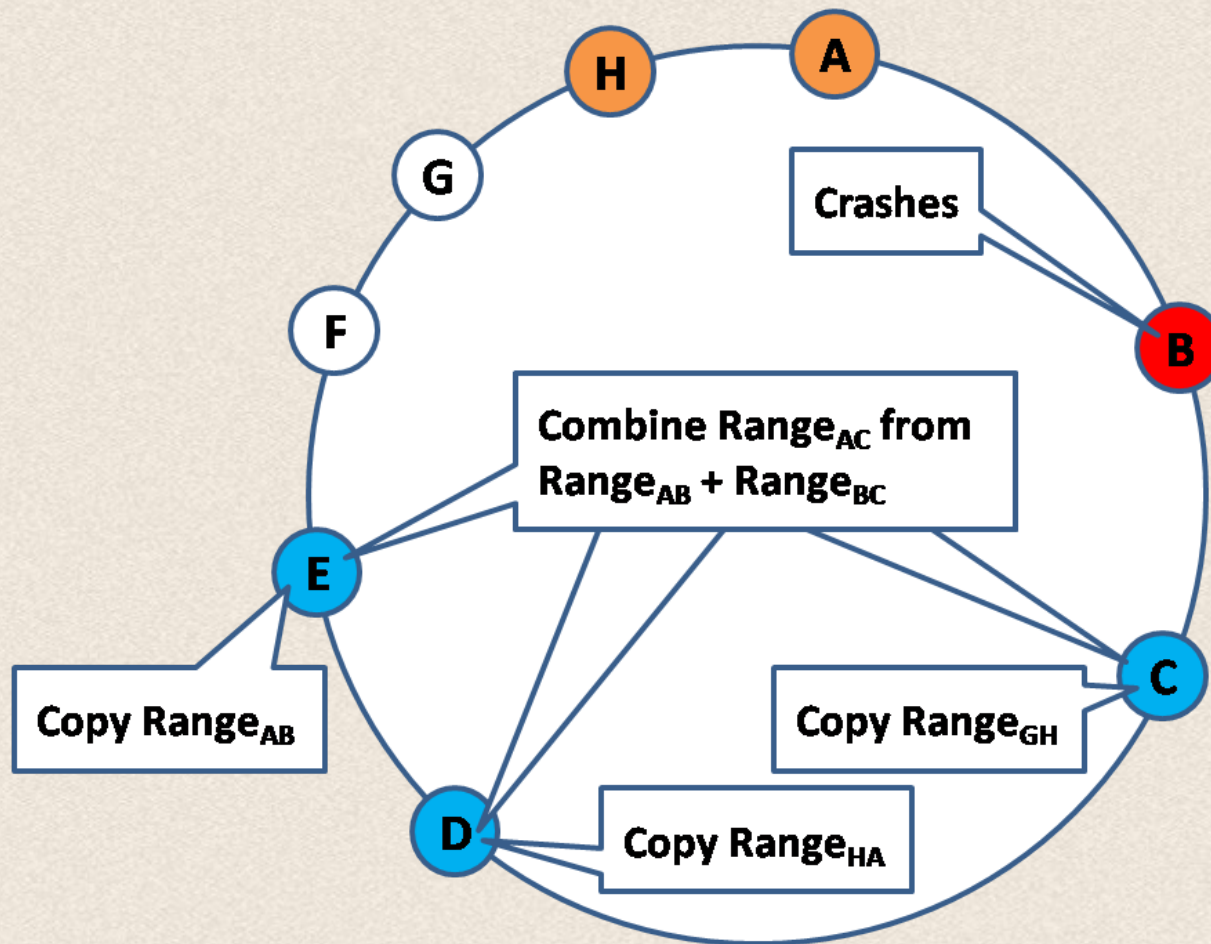
H, A, X, B, C, D will update the membership synchronously  
And then asynchronously propagate the membership changes to other nodes





# When Existing Node Leaves Network

Asynchronously propagate the membership changes to other nodes



# Database Integrity Outline

- Database Integrity Means
- Do We Really Need Consistency?
- Eventually Consistent
- Variations of Eventually Consistency
- Problem in Strict Schema
- Schema-Free

# Database Integrity Means

- Ensure data entered into the database is **accurate, valid, and consistent**. *Three basic types of integrity constraints:*
  - **Entity integrity**, allowing no two rows to have the same identity within a table.
  - **Domain integrity**, restricting data to predefined data types.
  - **Referential integrity**, requiring the existence of a related row in another table, e.g. a customer for a given customer ID.

# Do We Really Need Consistency?

- In strict OLTP environment (e.g. banking and ERP) data consistency is heart of the system.
- But even in Amazon (e-commerce) real-time consistency is not really needed.
- In large shared data environment such Facebook, Digg, Yahoo, Google, etc. data consistency can be relaxed
- Systems with strong ACID have poor performance.

# Eventually Consistent

- Specific form of weak consistency
- If no new updates are made, eventually all accesses will return the last updated value.
- System does not guarantee subsequent accesses will return the updated value.
- A number of conditions need to be met before the value will be returned.

# Variations of Eventually Consistency

- Causal consistency
- Read-your-writes consistency
- Session consistency
- Monotonic read consistency
- Monotonic write consistency

# Problem in Strict Schema

- Agile methodology is about changing adoption
- Dynamic Frameworks (e.g. Ruby on Rails, Django, and Grails, Symfony) are now widely used
- In many cases it is hard to migrate across database
- Adding more columns to an existing table is a full version

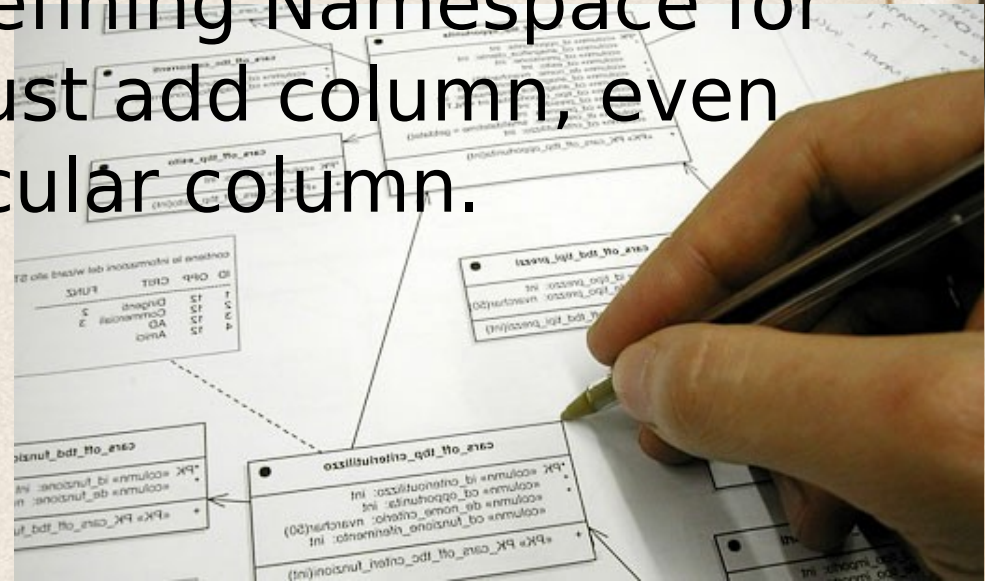
**symfony**

**django**



# Schema-Free

- Enable to add column in row level. Not restricted to column level.
- Each rows only use column they need (saving space).
- All we need to do is defining Namespace for tables. Then we can just add column, even another table in particular column.
- No more integration **headache**





# Conclusion & (not a) Summary

- NoSQL is yet another form of database.
- NoSQL don't intend to replace RDBMS.
- It is database alternative in Large data shared environment.
- Relaxing consistency will boost database availability and performance.
- There is no *Free Lunch* and *Silver Bullet* in database technologies.

A scroll of parchment is unrolled, showing the text "Thank You" in the center. The scroll is set against a dark, textured background. The parchment is light beige with a slightly aged appearance. The text is written in a simple, dark brown font. The scroll is rolled up at the top and bottom edges, with the ends of the parchment visible on the left and right sides. There are some decorative flourishes in the top left corner of the parchment.

Thank You