



Couchbase





Introduction à Couchbase Server 2.0

Tugdual Grall
Technical Evangelist
@tgrall
Couchbase



{“about” : “me”}

Tugdual “Tug” Grall

- Couchbase
 - Technical Evangelist
- eXo
 - CTO
- Oracle
 - Developer/Product Manager
 - Mainly Java/SOA
- Developer in consulting firms
- Web
 - @tgrall
 - <http://blog.grallandco.com>
 - Tgrall
- NantesJUG co-founder
- Pet Project :
 - <http://www.resultri.com>

Agenda

1. Introduction to Couchbase Server
2. Cluster Management
3. Operation and Document Design
4. Introduction to Views and Queries
5. Cross Data Center Replication
6. Q&A

Couchbase Developer Day Paris is this Friday! <http://couchbase.com/events>



Couchbase Server 2.0



Couchbase Server NoSQL Document Database



Couchbase

Couchbase Open Source Project

- Leading NoSQL database project focused on distributed database technology and surrounding ecosystem
- Supports both key-value and document-oriented use cases
- All components are available under the **Apache 2.0 Public License**
- Obtained as packaged software in both enterprise and community editions.



COUCHBASE
OPEN SOURCE PROJECT

Couchbase Server



Easy Scalability

Grow cluster without application changes, without downtime with a single click



Consistent High Performance

Consistent sub-millisecond read and write response times with consistent high throughput



Always On 24x365

No downtime for software upgrades, hardware maintenance, etc.



Flexible Data Model

JSON document model with no fixed schema.

New in 2.0


JSON support



Indexing and Querying



Incremental Map Reduce

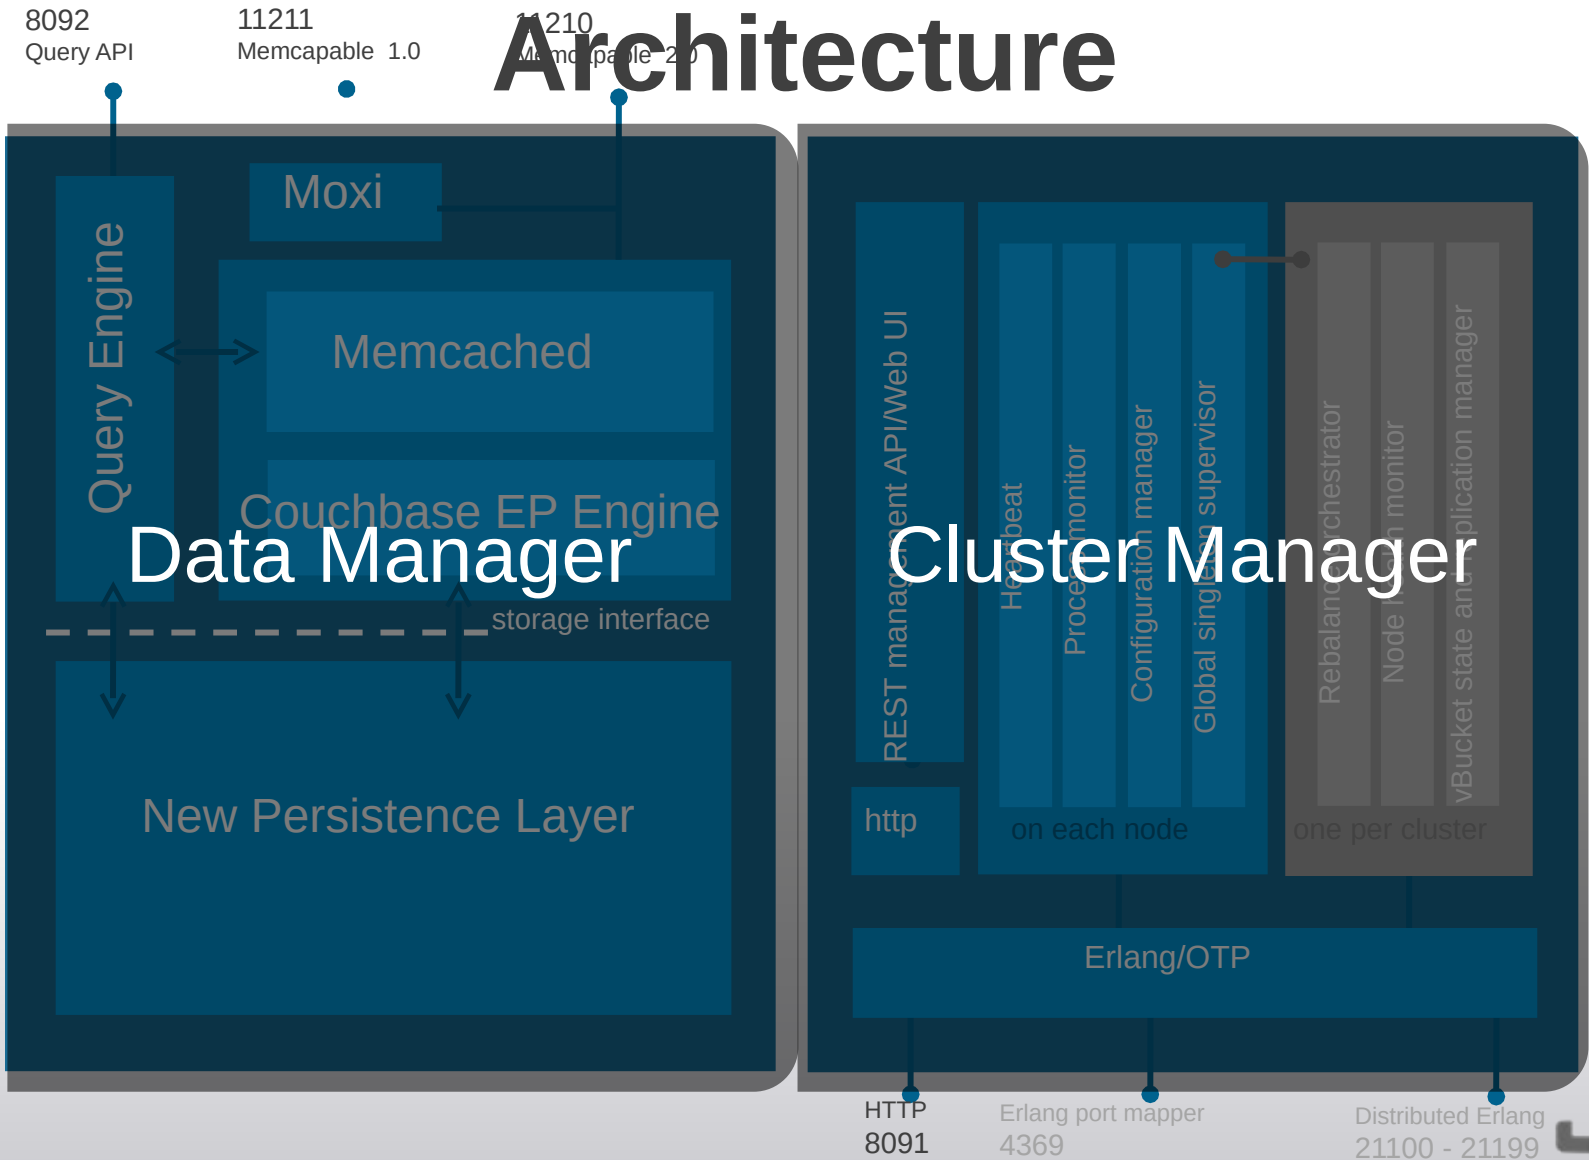


Cross data center replication

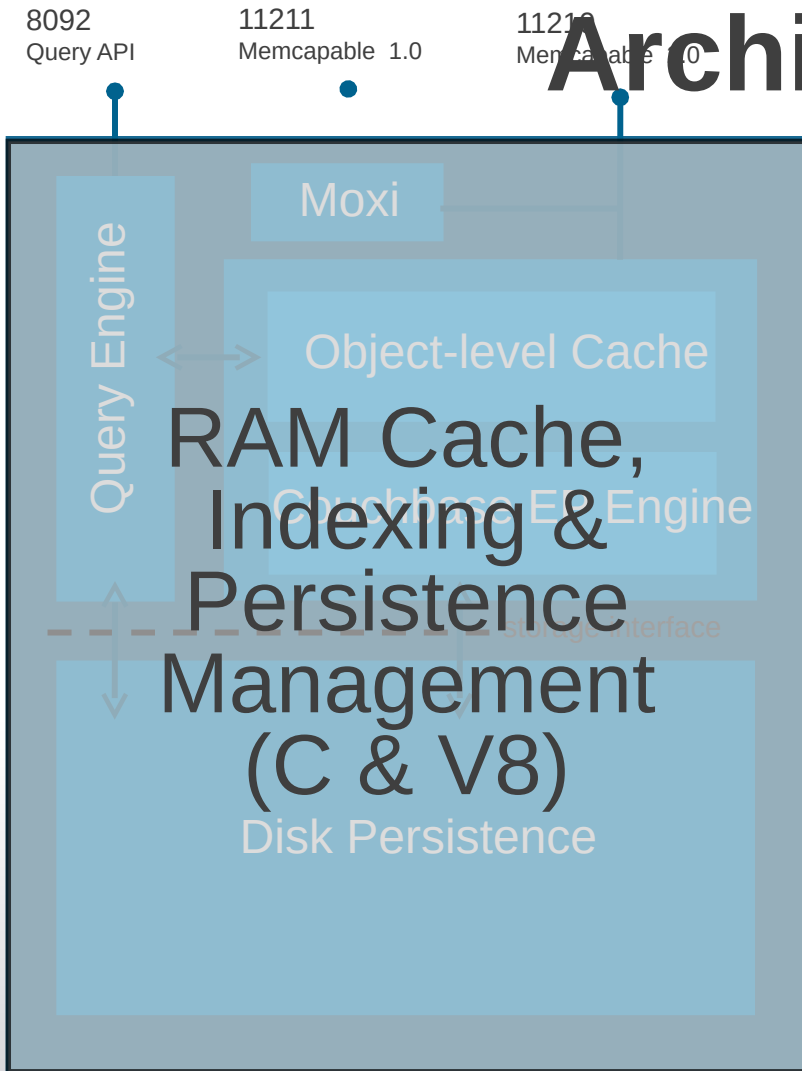


Couchbase Server 2.0

Architecture



Couchbase Server 2.0 Architecture



The Unreasonable Effectiveness of C by Damien Katz

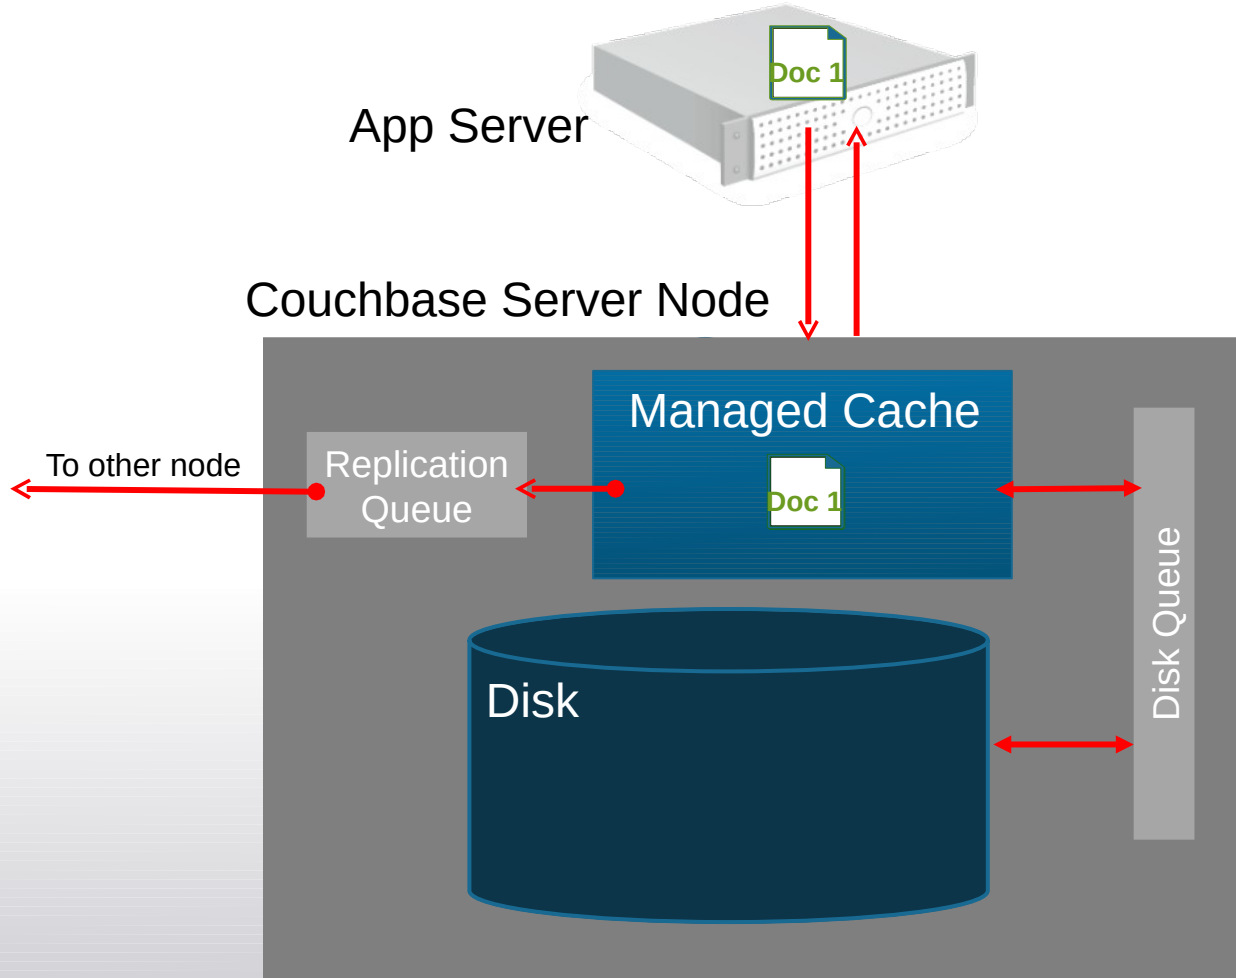
HTTP 8091

Erlang port mapper 4369

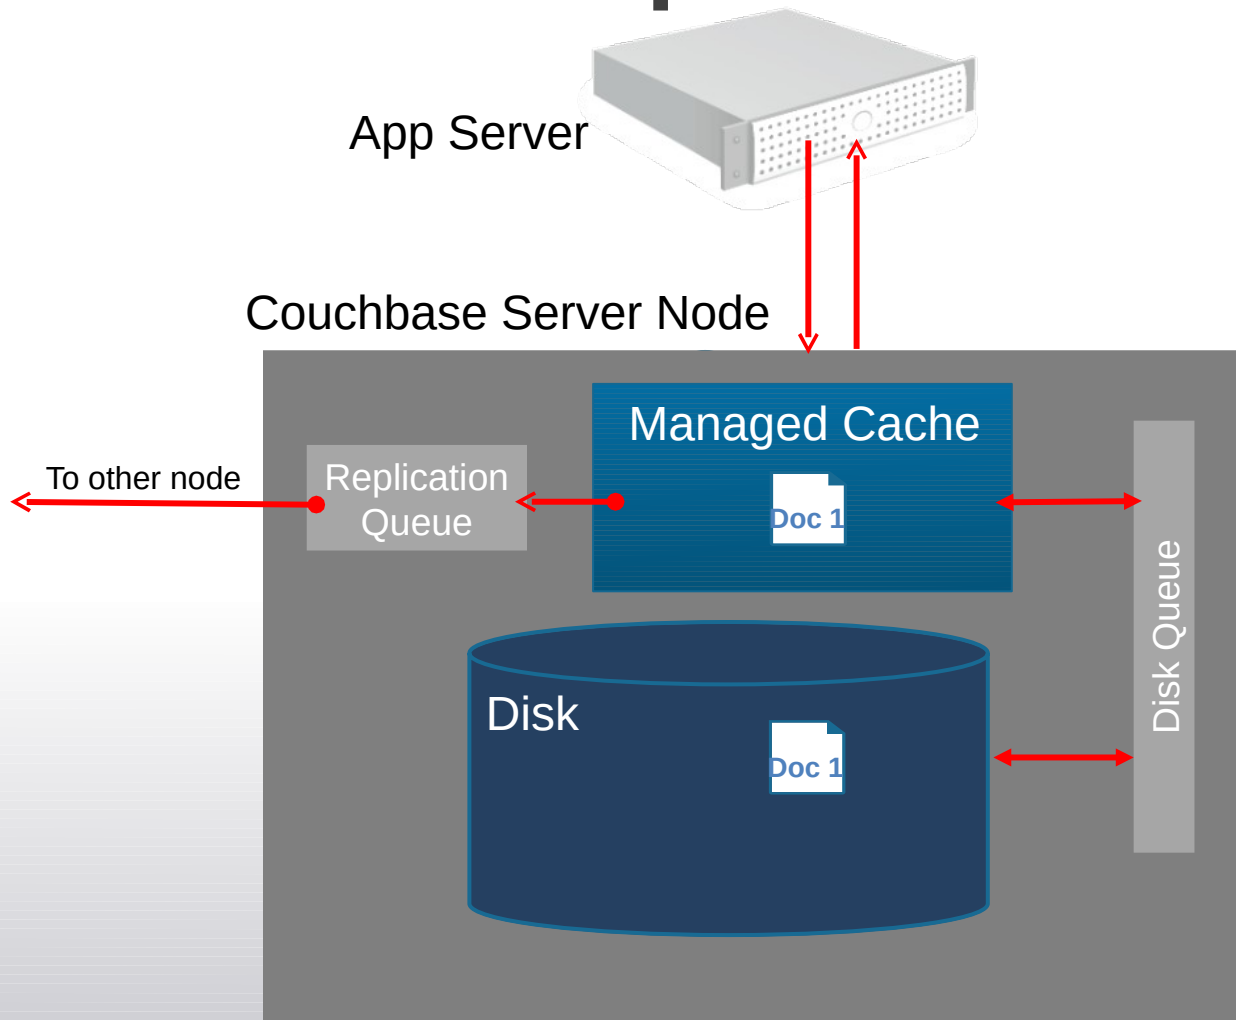
Distributed Erlang 21100 - 21199



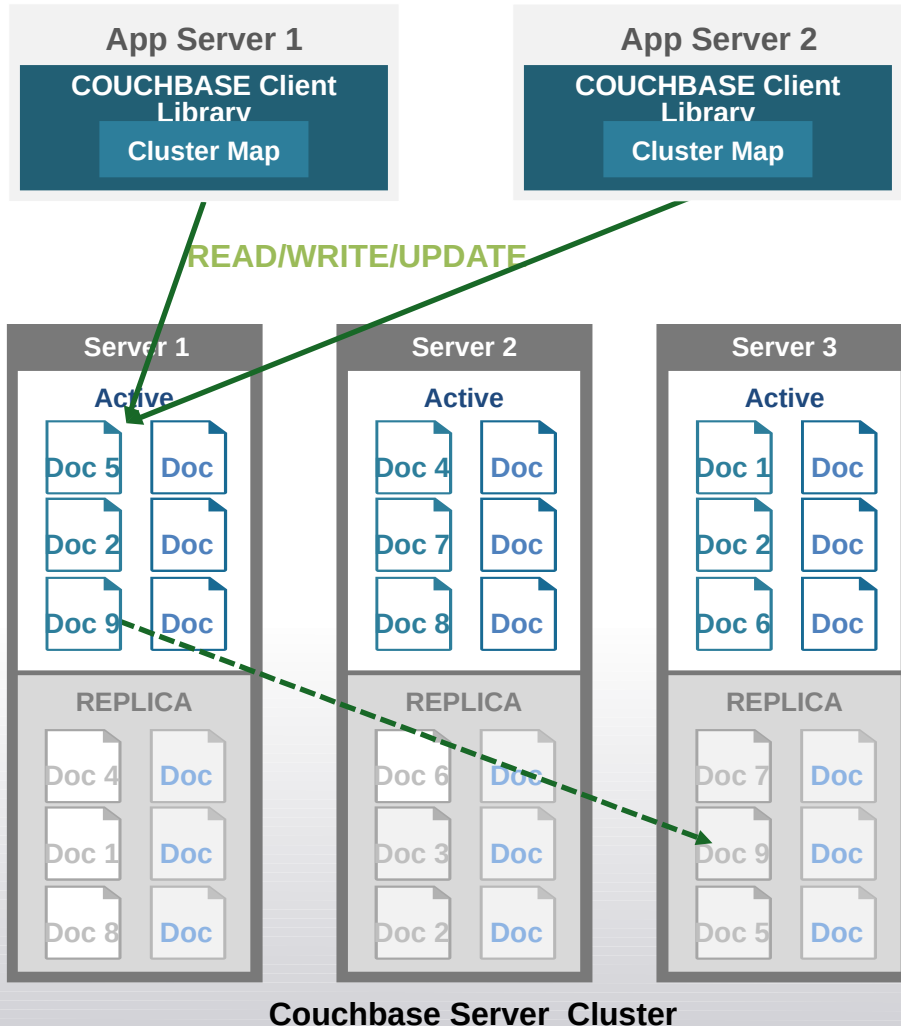
Single node - Couchbase Write Operation



Single node - Couchbase Read Operation

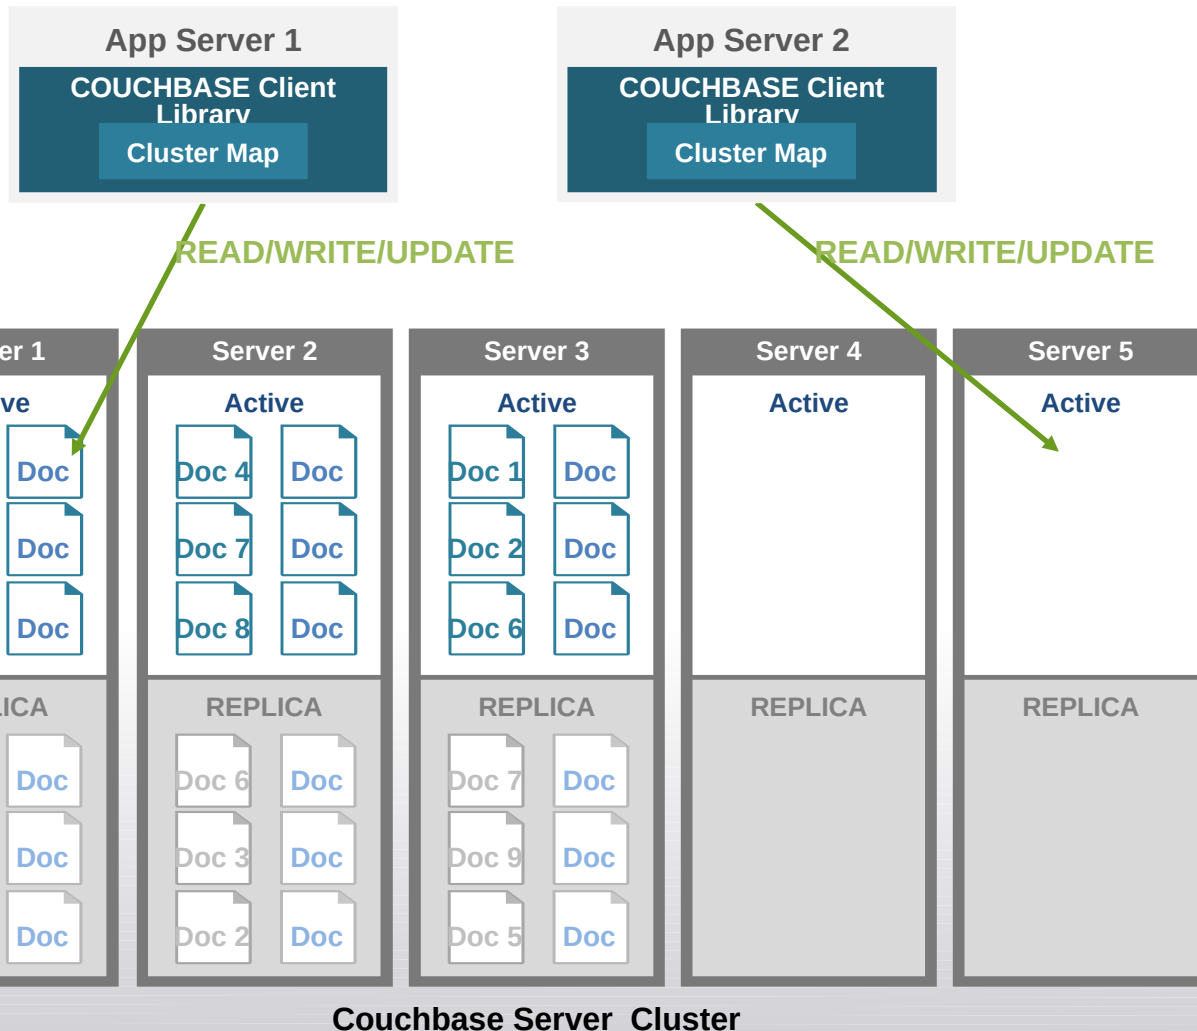


Basic Operation



- Docs distributed evenly across servers
- Each server stores both active and replica docs
 - Only one server active at a time
- Client library provides app with simple interface to database
- Cluster map provides map to which server doc is on
 - App never needs to know
- App reads, writes, updates docs
- Multiple app servers can access same document at same time

Add Nodes to Cluster

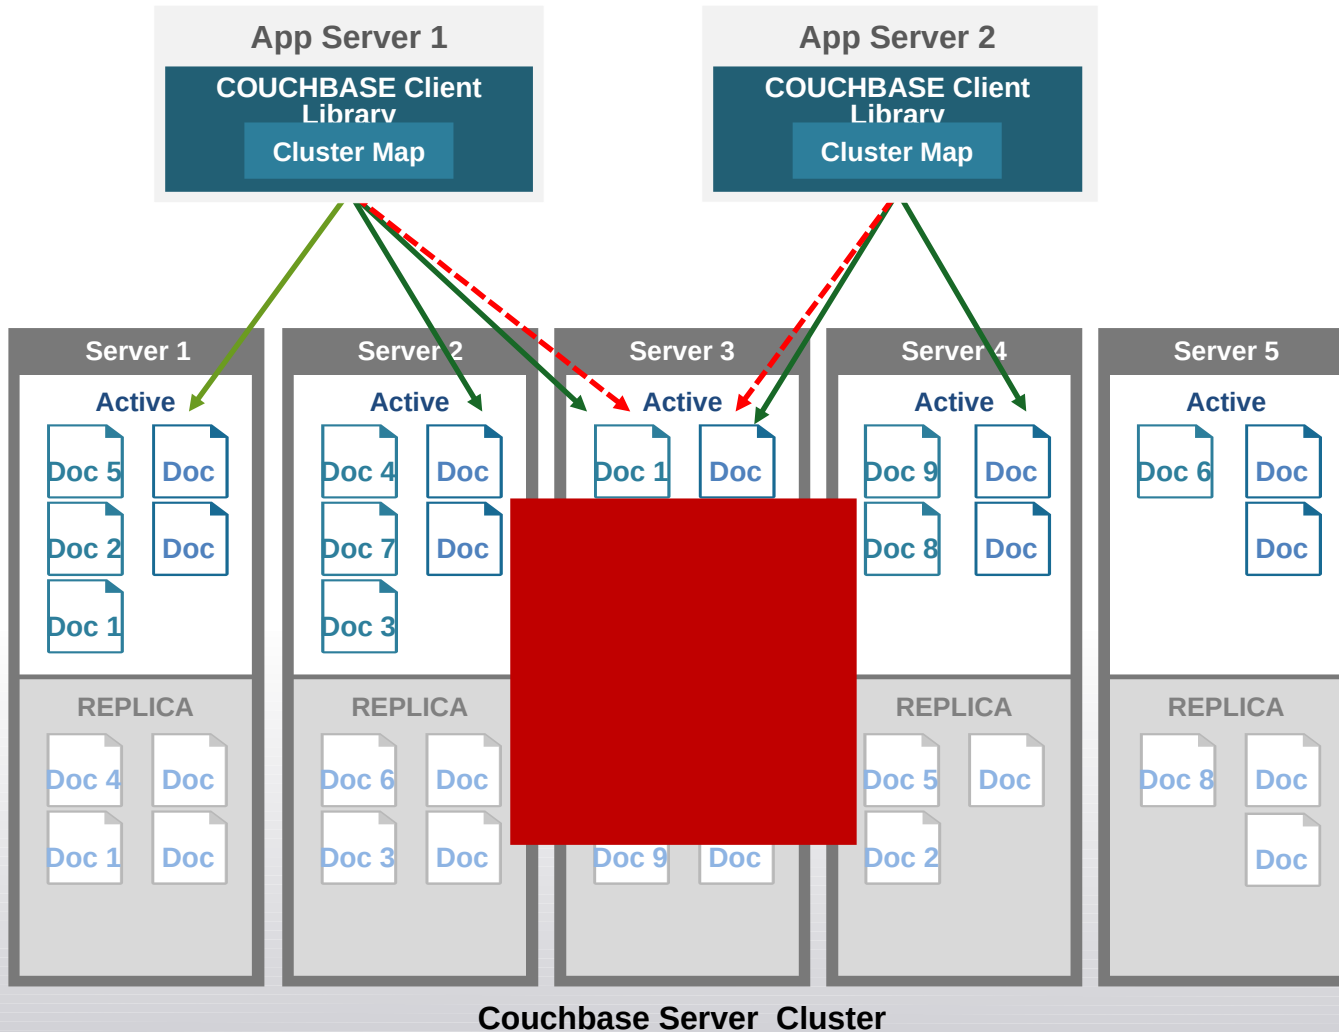


- Two servers added
One-click operation
- Docs automatically
rebalanced across
cluster
 - Even distribution of docs
 - Minimum doc movement
- Cluster map
updated
- App database
calls now
distributed
over larger number
of servers



Couchbase

Fail Over Node



- App servers accessing docs
- Requests to Server 3 fail
- Cluster detects server failed
 - Promotes replicas of docs to active
 - Updates cluster map
- Requests for docs now go to appropriate server
- Typically rebalance would follow





Demo Time





Operation & Document Design



Basic Store & Retrieve Operations

- **get(key)**
Retrieve a document
- **set(key, value)**
Store a document or replace if it exists
- **add(key, value)**
Store a document and fail if it exists
- **replace(key, value)**
Replace a document and fail if it does not exist

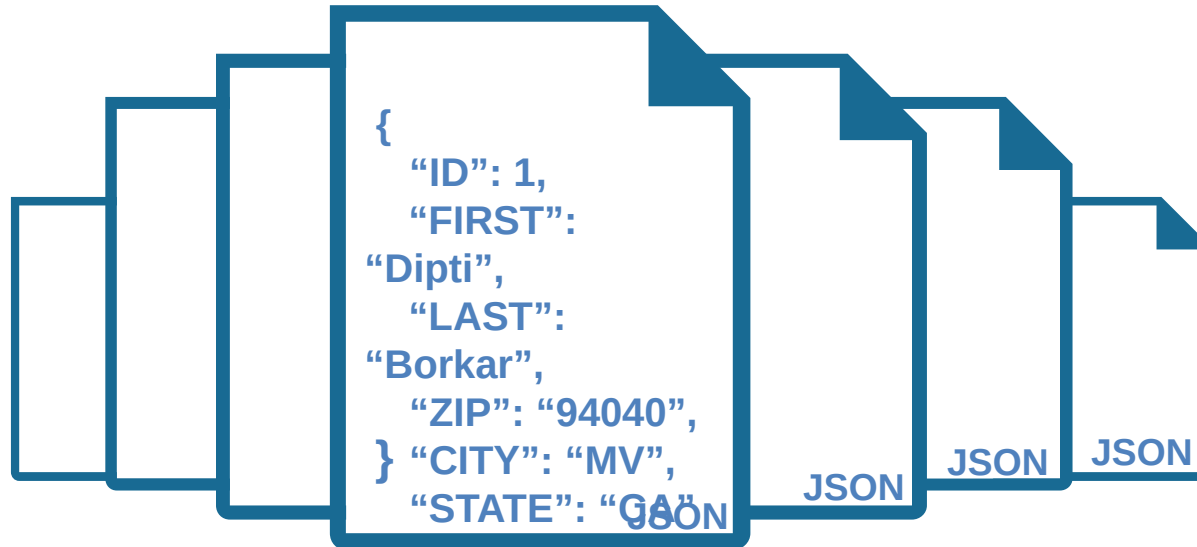
Lots of other Operations

- **Atomic Counters (increment, decrement)**
- **Append/Prepend**
- **CAS (Compare and Set)**
 - Optimistic Locking
- **“Get with Lock”**
 - Write Lock on Objects
- **Bulk Operations**
 - Saves network overhead
- **Stats**
- **View operations**

Couchbase 2.0 Bonus Points: JSON

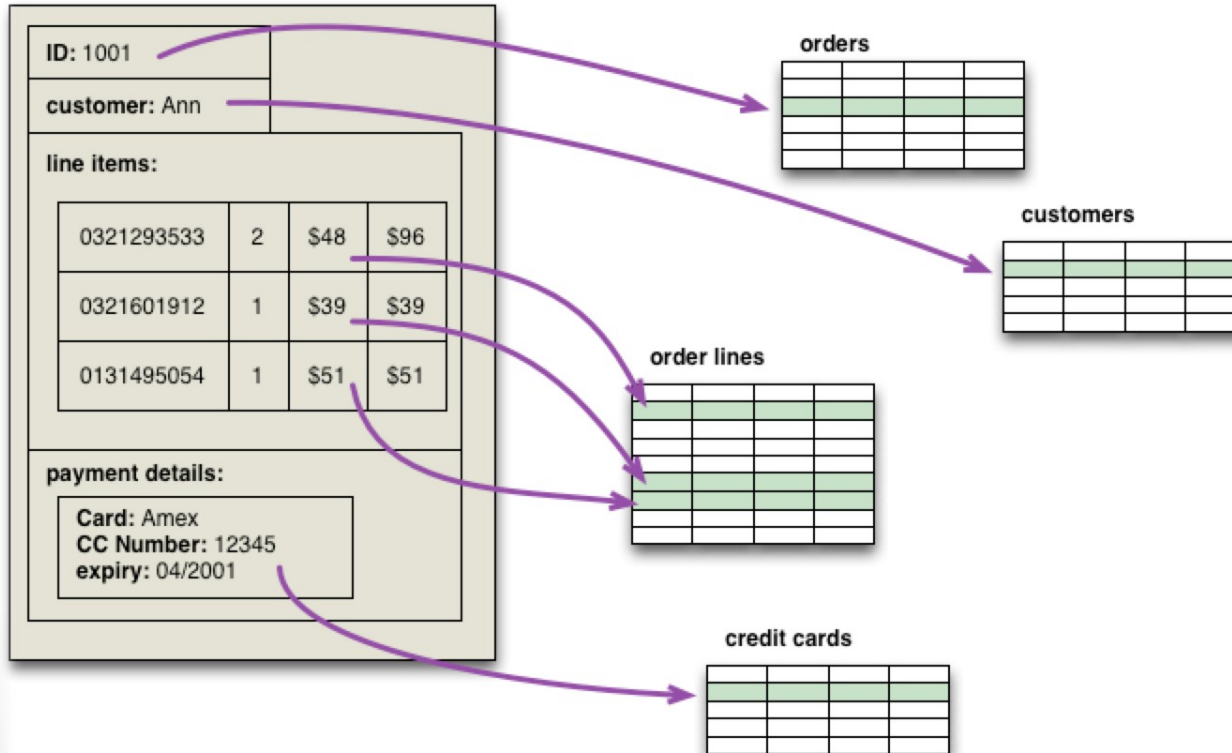
- JSON is a **lightweight format** to represent document structure in a language-independent manner.
- If JSON documents are stored, the **View engine** can be used.
- Allows to build **secondary indexes** on your datasets.
- Makes it possible to ask questions like “Give me all user documents by lastname”.

Flexible Data Model



- No need to worry about the database when changing your application
- Records can have different structures, there is no fixed schema
- Allows painless data model changes for rapid application development

Aggregate View of Data



<http://martinfowler.com/bliki/AggregateOrientedDatabase.html>

Store and Retrieve Aggregates

ID: 1001			
customer: Ann			
line items:			
0321293533	2	\$48	\$96
0321601912	1	\$39	\$39
0131495054	1	\$51	\$51
payment details:			
Card: Amex CC Number: 12345 expiry: 04/2001			



```
order::1001
{
  uid: ji22jd,
  customer: Ann,
  line_items: [
    { sku: 0321293533, quan: 3, unit_price: 48.0
    },
    { sku: 0321601912, quan: 1, unit_price: 39.0
    },
    { sku: 0131495054, quan: 1, unit_price: 51.0
    }
  ],
  payment: { type: Amex, expiry: 04/2001,
    last5: 12345 }
}
```

- Easier to Distribute Data
- More Flexibility
- Reduced Latency



Demo Time




Official SDKs



www.couchbase.com/develop

Community SDKs



Couchbase



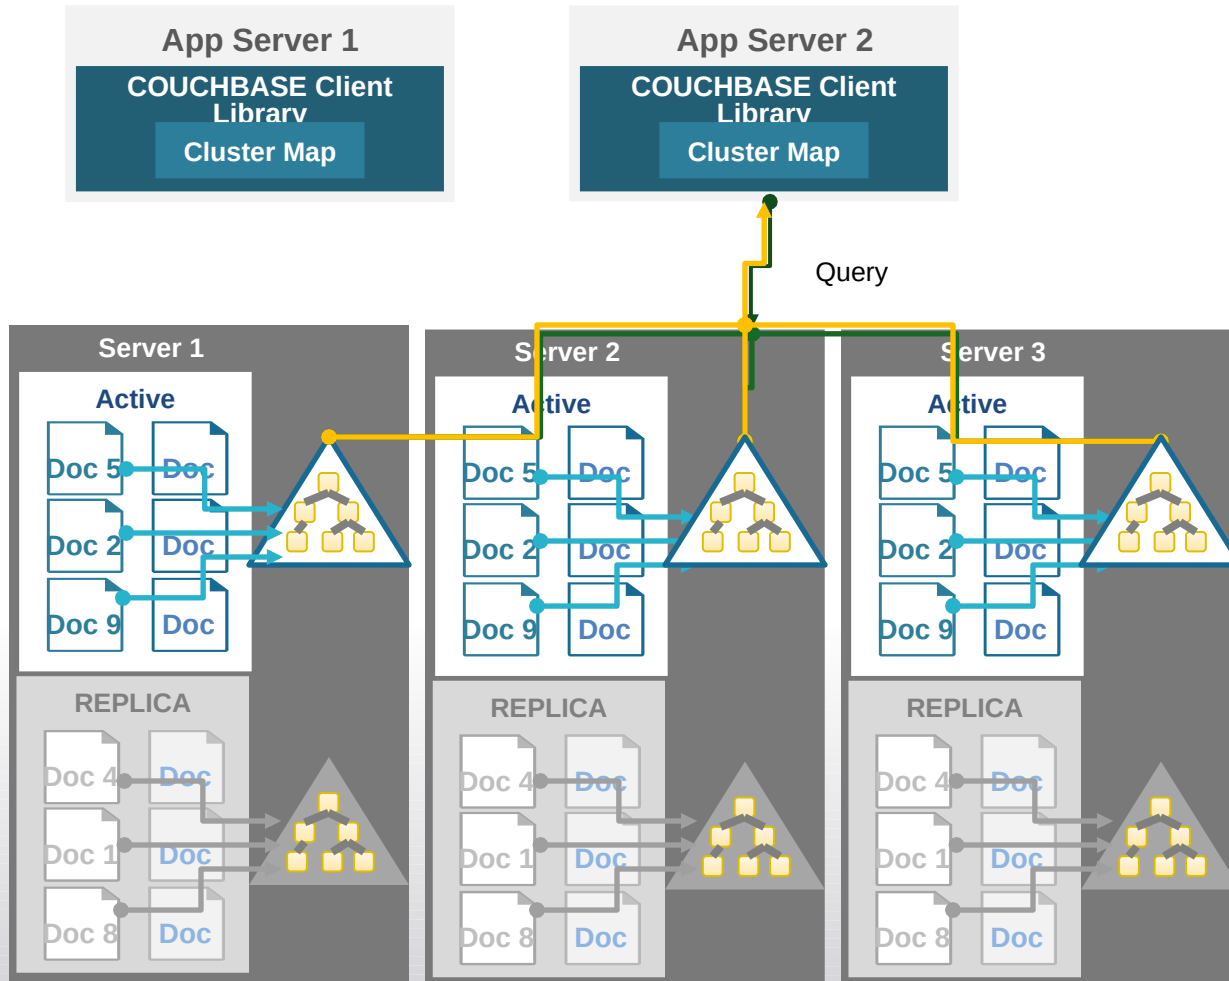
Indexing and Querying



Indexing and Querying – The basics

- Define materialized views on JSON documents and then query across the data set
- Using views you can define
 - Primary indexes
 - Simple secondary indexes (most common use case)
 - Complex secondary, tertiary and composite indexes
 - Aggregations (reduction)
- Indexes are *eventually indexed*
- Queries are *eventually consistent*
- Built using Map/Reduce technology
 - Map and Reduce functions are written in Javascript

Indexing and Querying



Couchbase Server Cluster

- Indexing work is distributed amongst nodes
- Large data set possible
- Parallelize the effort
- Each node has index for data stored on it
- Queries combine the results from required nodes



Demo Time





Cross Data Center Replication (XDCR)



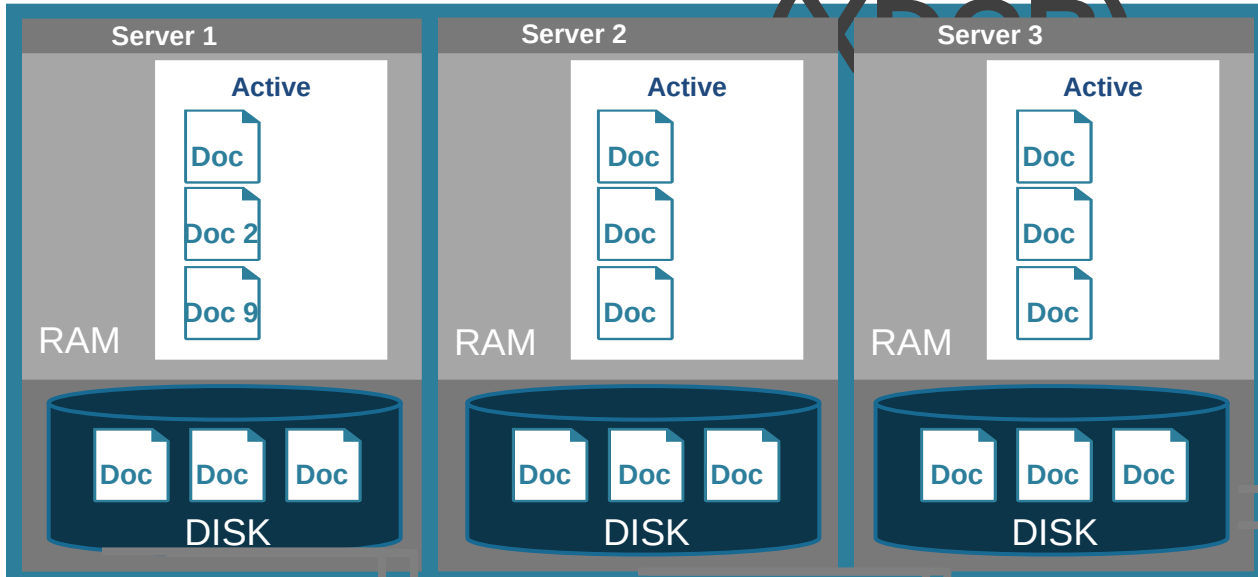
Cross Data Center Replication – The basics

- Replicate your Couchbase data across clusters
- Clusters may be spread across geos
- Configured on a per-bucket (per-database) basis
- Supports unidirectional and bidirectional operation
- Application can read and write from both clusters
 - Active – Active replication
- Replication throughput scales out linearly
- Different from intra-cluster replication

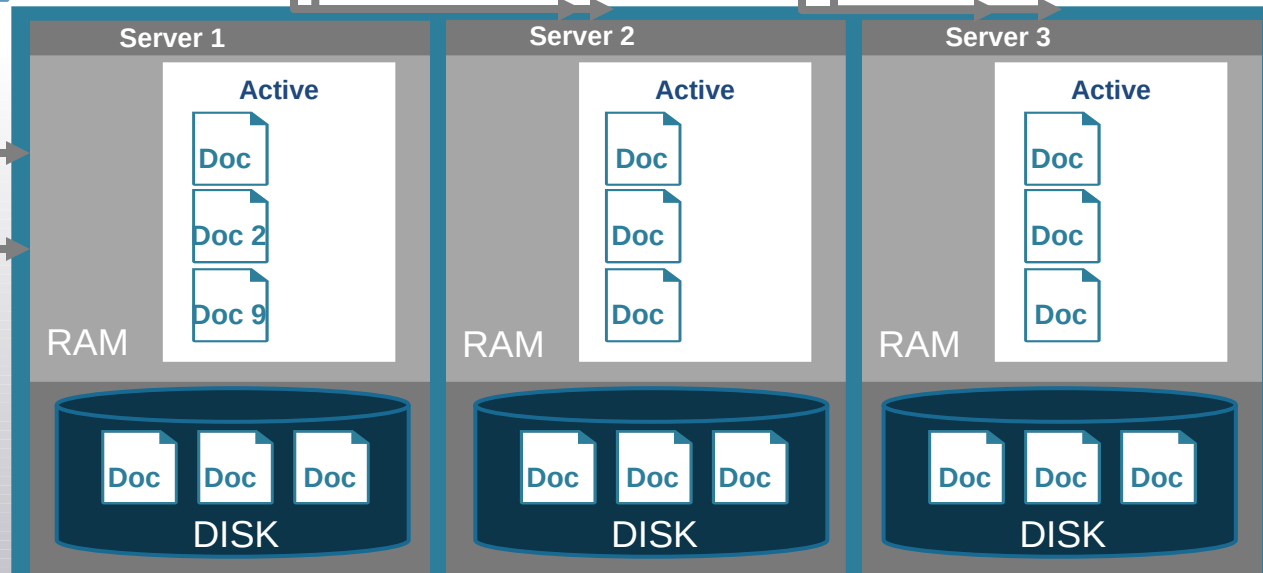
Cross Data Center Replication

(YDOP)

Couchbase Server Cluster
US



Couchbase Server Cluster
EMEA



Couchbase Handles Real World





Q & A





[Dev Day]

- **Couchbase DeveloperDays Paris**
 - This Friday !
 - Presentations & Labs
 - Information
 - <http://www.couchbase.com/events>
 - Use the promo code TUG

Thank you!

tug@couchbase.com
@tgrall



Get Couchbase Server at
<http://www.couchbase.com/download>





Couchbase

