



Couchbase



Getting Started with Couchbase

Tugdual “Tug” Grall
Technical Evangelist

email: tug@couchbase.com

twitter: @tgrall



Getting Started with Couchbase



Tugdual “Tug” Grall
Technical Evangelist

email: tug@couchbase.com

twitter: @tgrall

**Stefan Schneider**

@PLOTTIFY



Free [#Cologne](#) [#NoSQL](#) Developer Event
with [#CouchDB](#) Expert [@tgrall](#) 6.2., 19Uhr
via [@nerdhub_cgn](#) [nerdhub.de/events/17-
nosq...](http://nerdhub.de/events/17-nosq...) [@Bumi](#) [@HornCologne](#)

[← Reply](#) [↻ Retweet](#) [★ Favorite](#) [⋮ More](#) [📁 Pocket](#)

1
RETWEET



8:39 AM - 1 Feb 13

Couchbase

**Stefan Schneider**

@PLOTTIFY



Follow

Free [#Cologne](#) [#NoSQL](#) Developer Event
with [#Couchbase](#) Expert [@tgrall](#) 6.2., 19Uhr
via [@nerdhub_cgn](#) [nerdhub.de/events/17-
nosq...](http://nerdhub.de/events/17-nosq...) [@Bumi](#) [@HornCologne](#)



Reply



Retweet



Favorite



More

Pocket

1

RETWEET



8:39 AM - 1 Feb 13

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>

Couchbase Server Core Principles



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.

Couchbase

Couchbase 2.0 Launch!

12/12/12



Couchbase

Couchbase 2.0 New Features

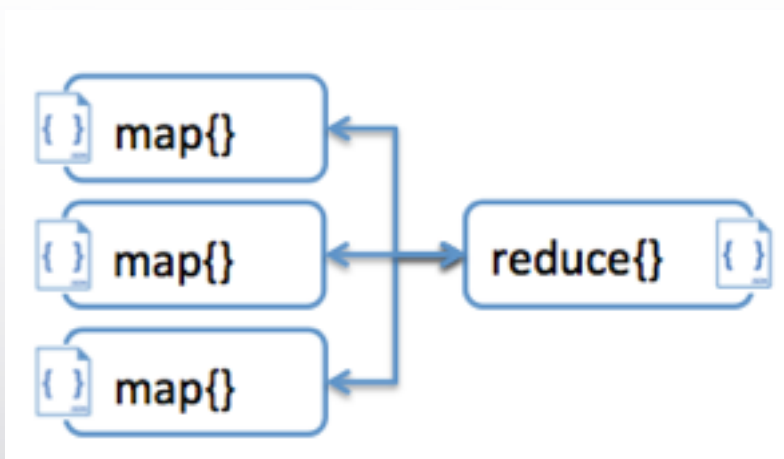
JSON support




Indexing and Querying



Incremental Map Reduce



Cross data center replication



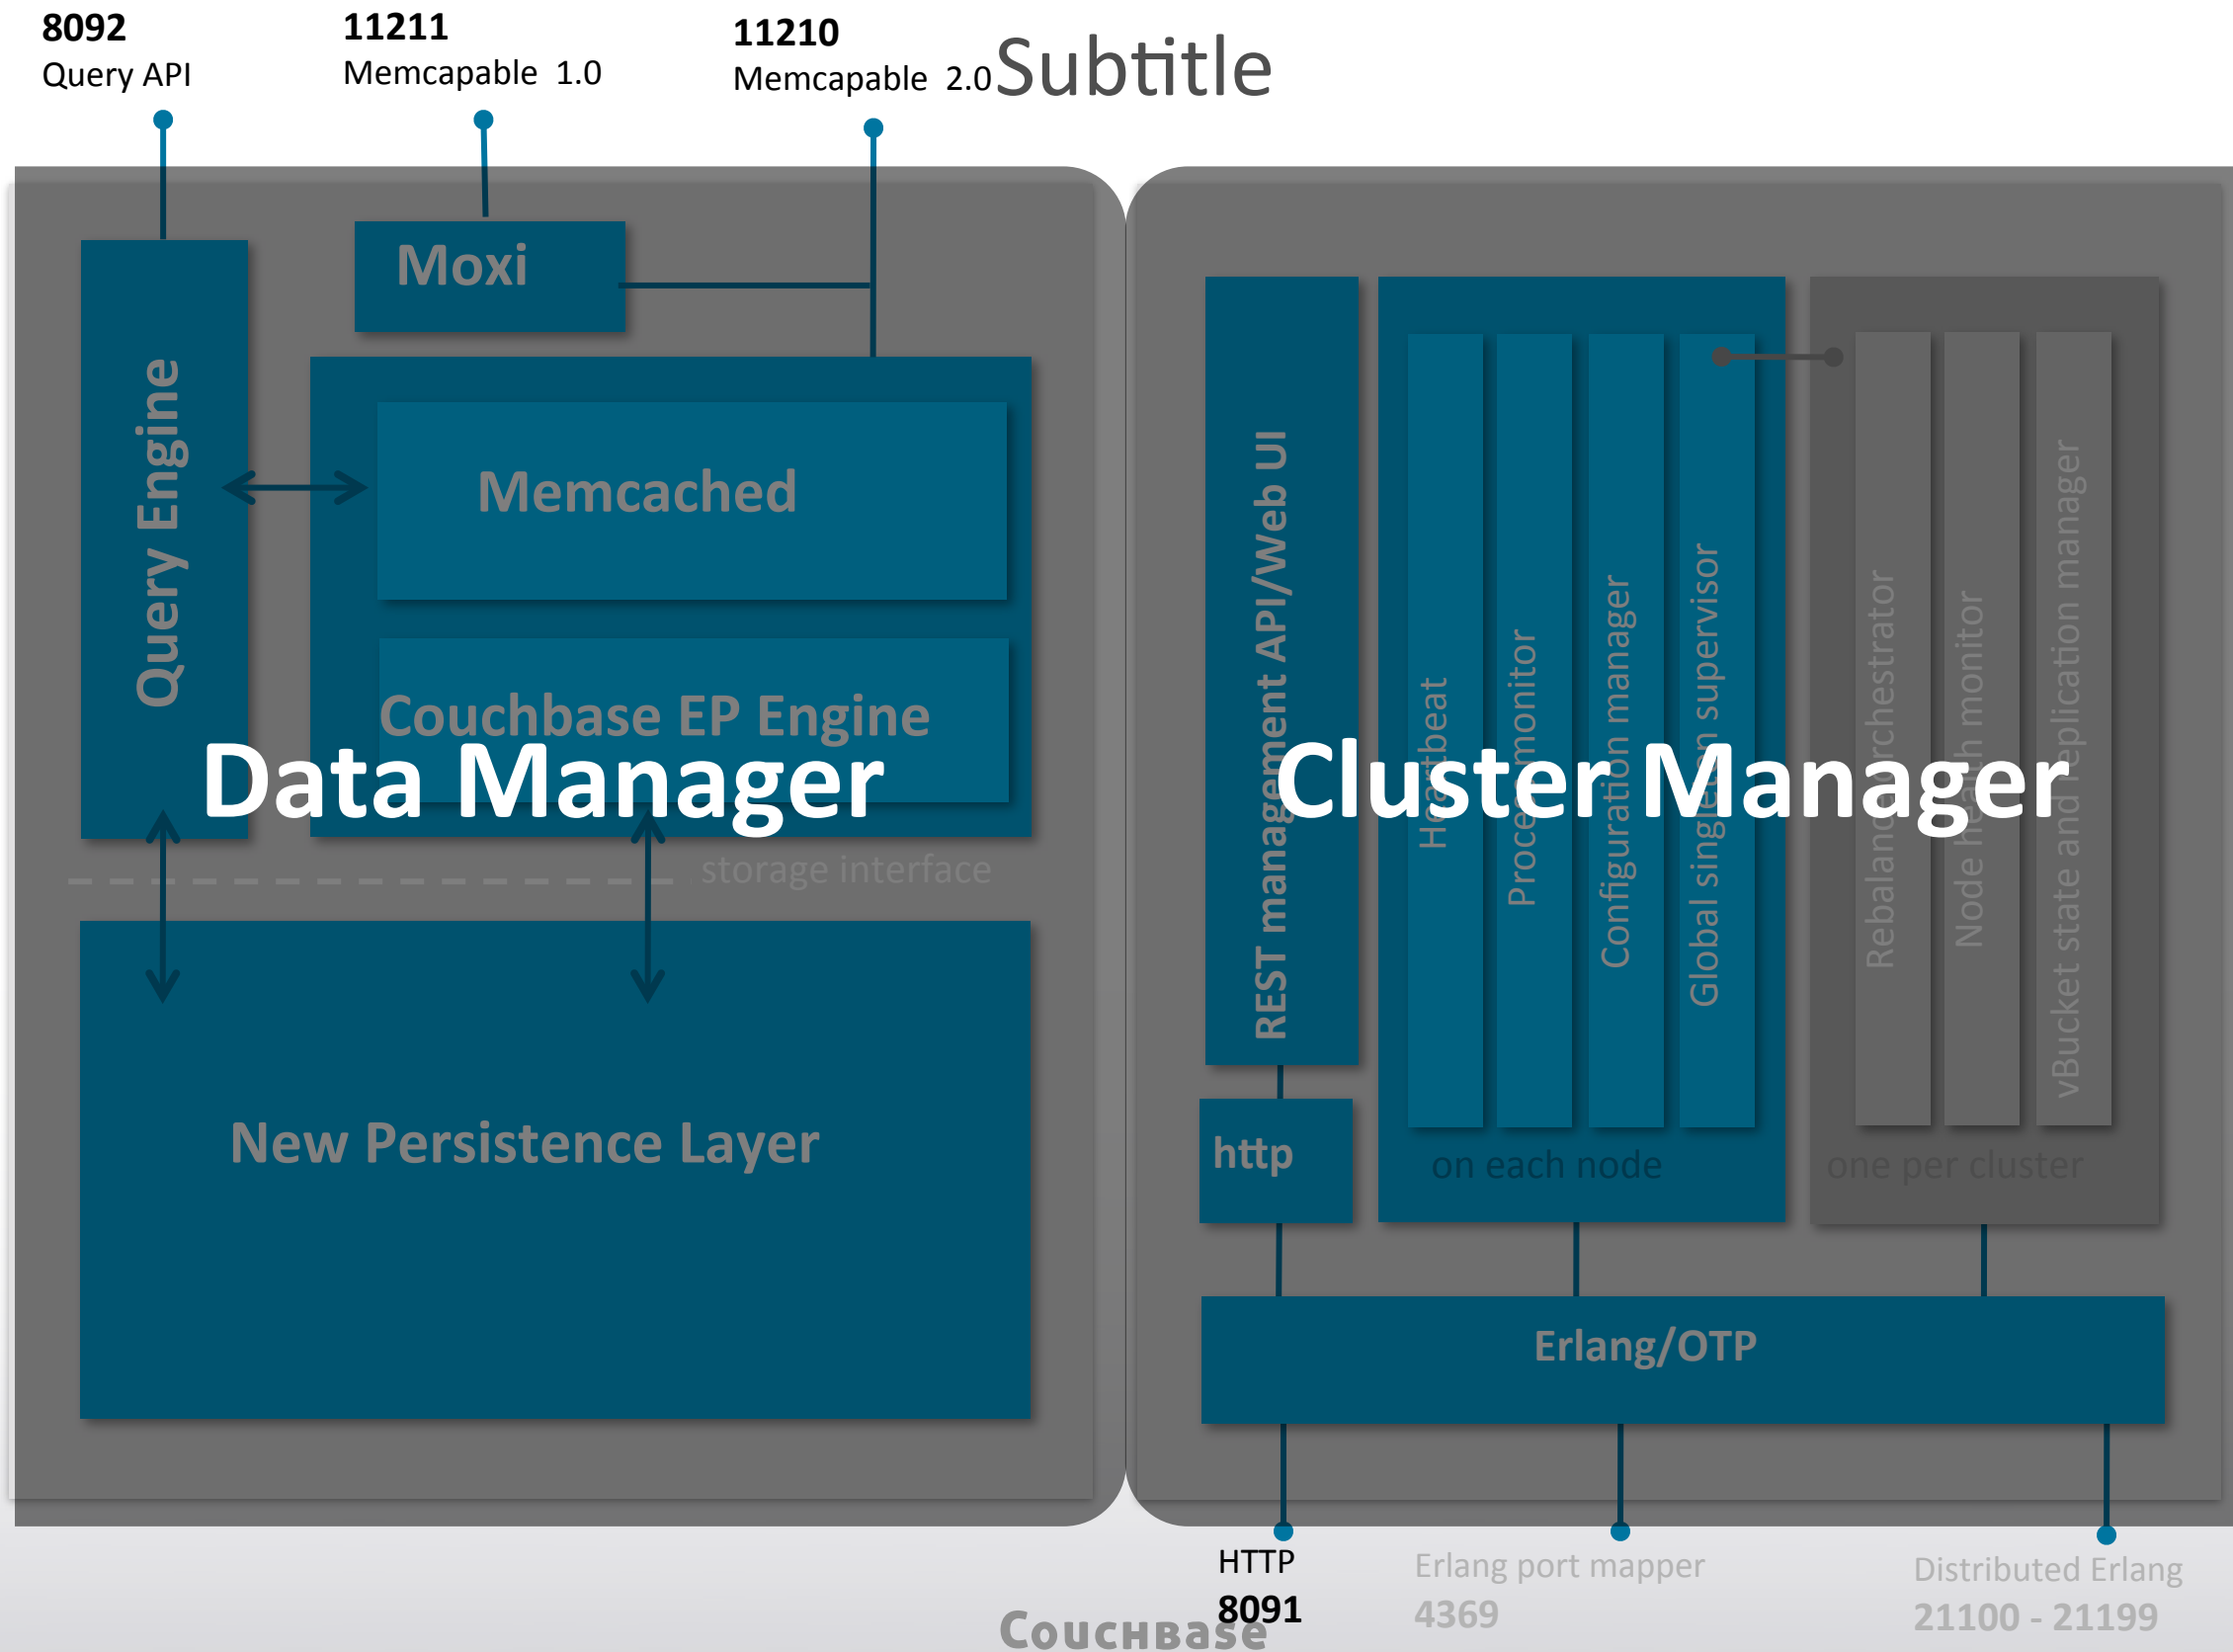
Couchbase

Couchbase Handles Real World Scale

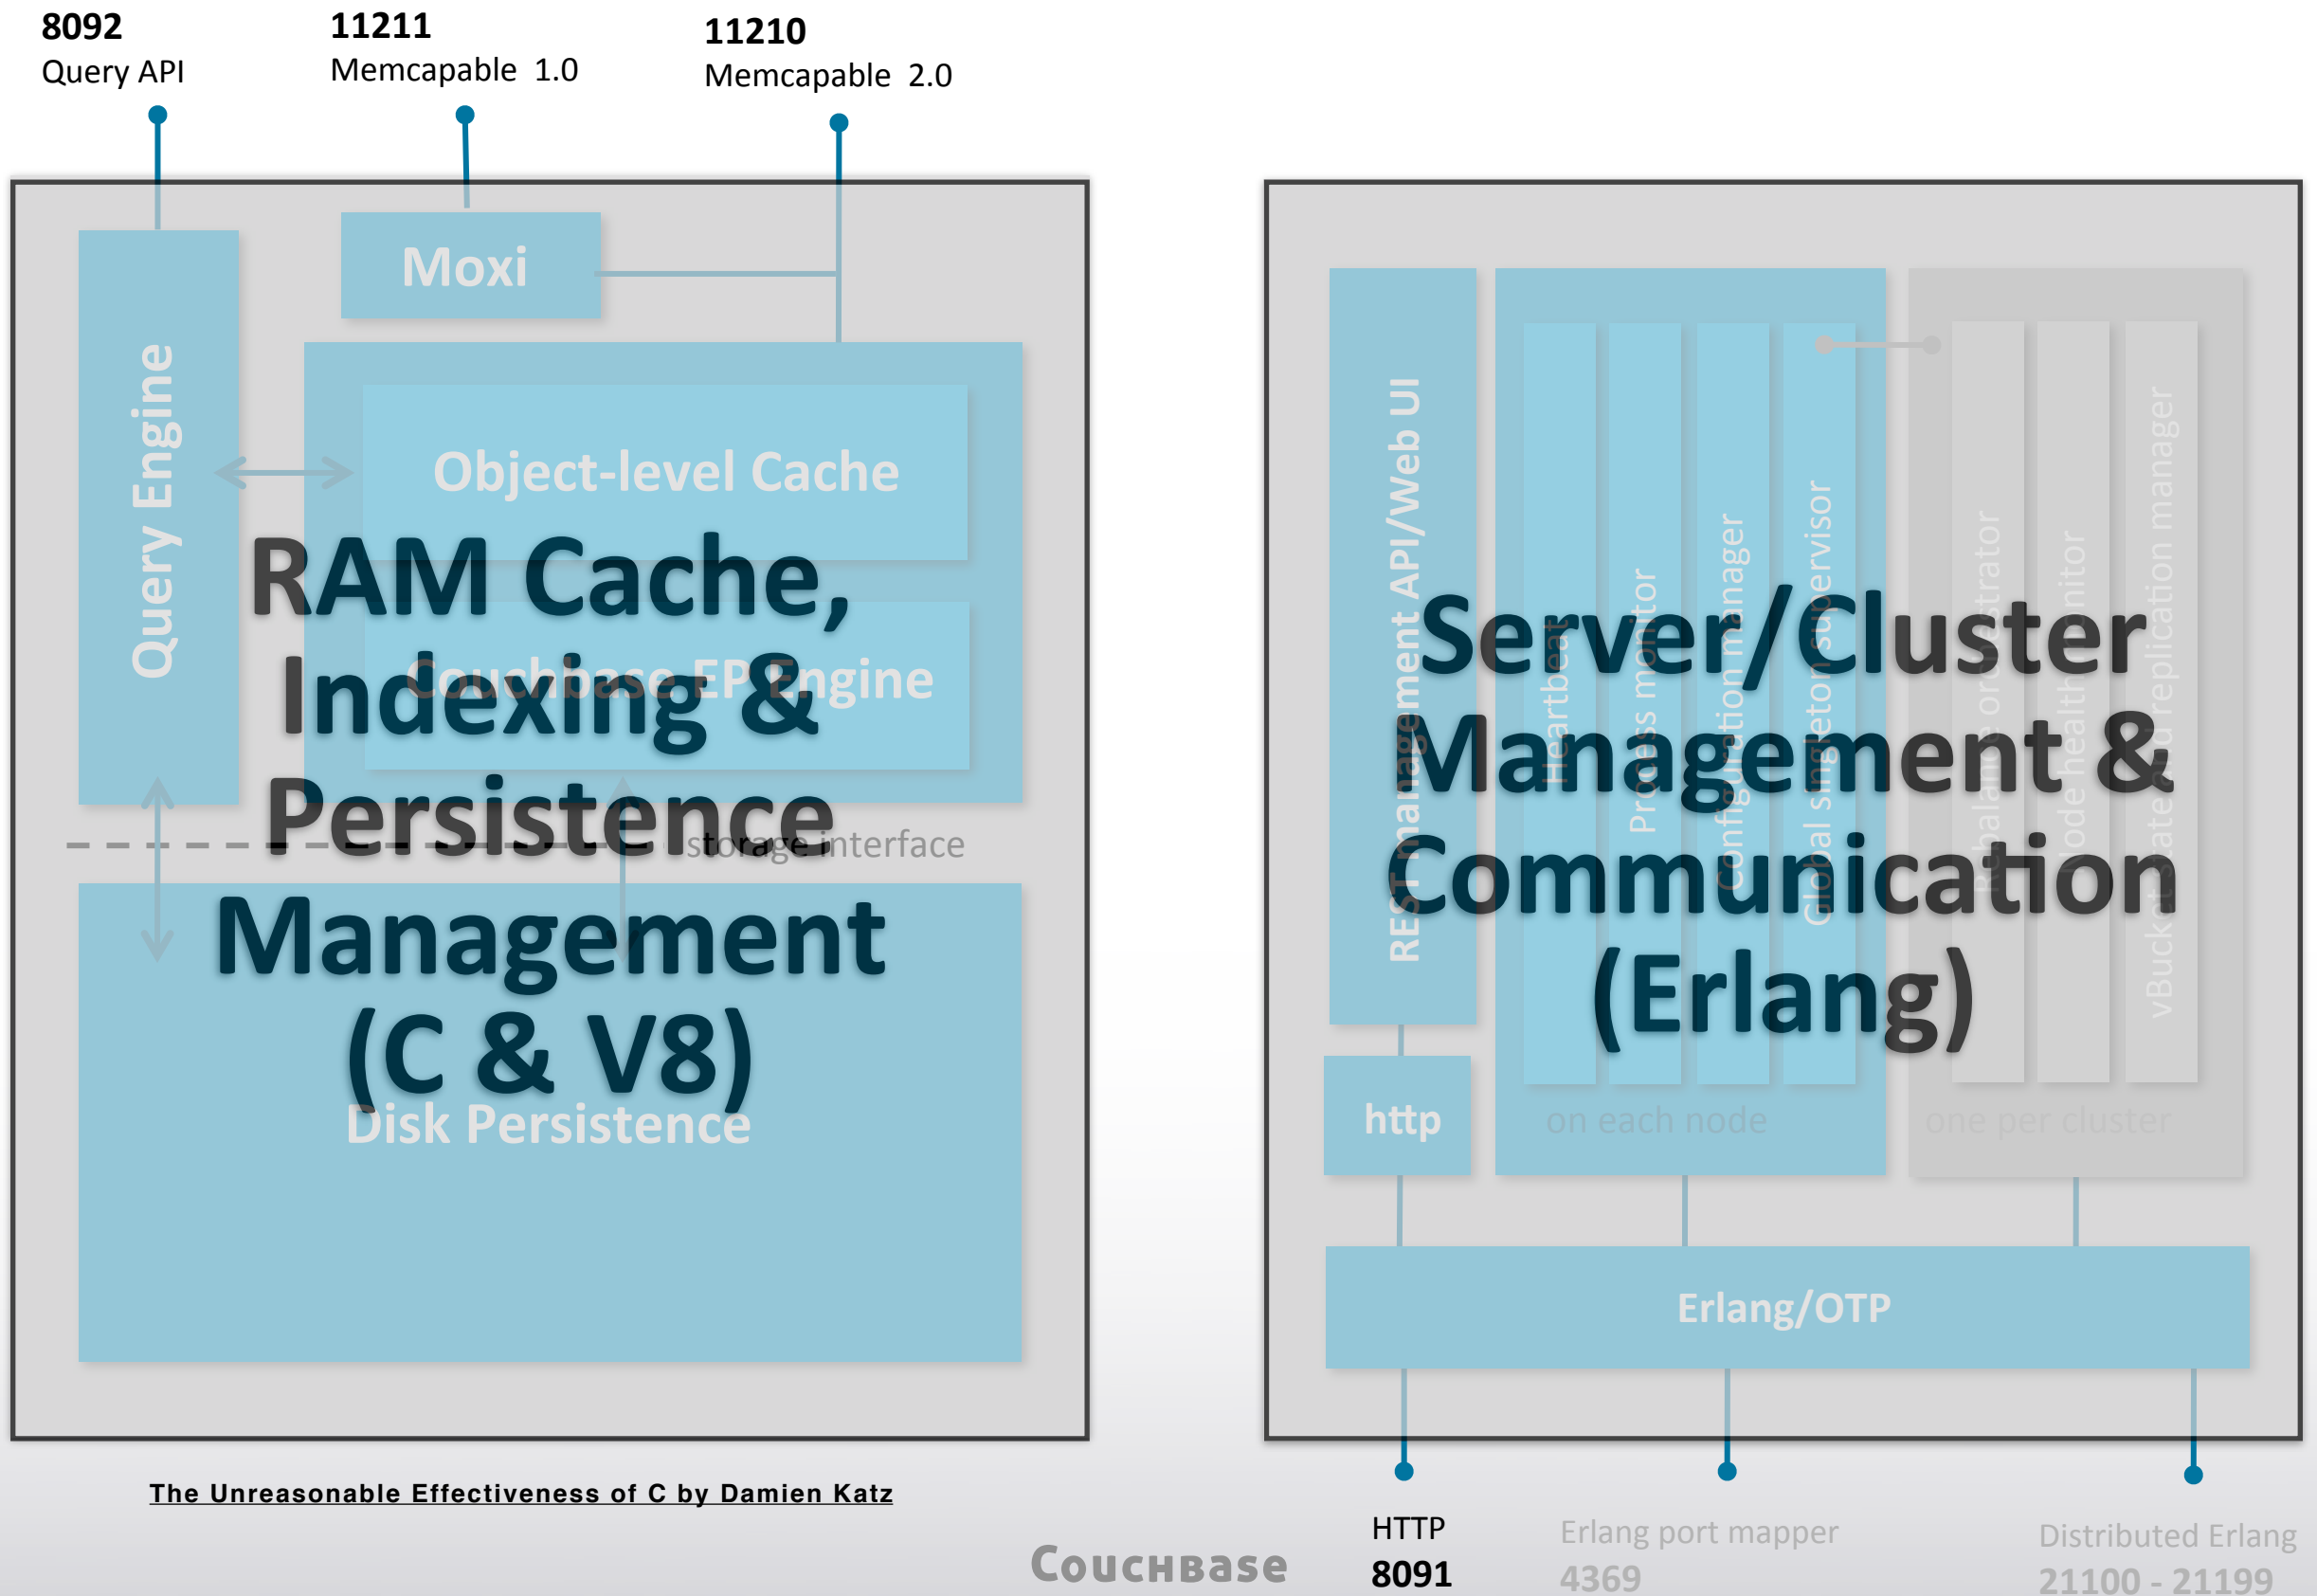


Couchbase

Couchbase Server 2.0 Architecture



Couchbase Server 2.0 Architecture



Open Source Project

Apache 2.0

Repositories

Find a Repository...

All Sources Forks Mirrors

testrunner Python ★ 9 | 15
forked from membase/testrunner
The TestRunner (Extracted from carin).
Last updated 9 hours ago

couchstore C ★ 15 | 7
forked from couchbaselabs/couchstore
couchbase storage file library
Last updated 13 hours ago

php-ext-couchbase C ★ 30 | 21
forked from couchbaselabs/php-ext-couchbase
A PHP extension wrapping libcouchbase
Last updated 4 days ago

manifest Ruby ★ 5 | 10
forked from membase/manifest
repo manifest
Last updated 2 days ago

couchbase-ruby-client C ★ 30 | 12
Couchbase Ruby client library (official) built atop libcouchbase
Last updated 2 days ago

couchbase
Joined on Feb 07, 2011
59 public repos 37 members

status:open

Search for status:open

ID	Subject	Owner	Project	Branch	Updated	V	R
1dcb1102	PCBC-83 - Add persist_to and replicate_to	Toni Nordby	php-ext-couchbase	master	12:17 PM	-1	
1e8ec0da	Reef up options code	Toni Nordby	php-ext-couchbase	master	8:29 AM		
1ecb45638	Remove debugger development dependency	Sergey Avseyev	couchbase-ruby-client	master	12:05 AM		✓
1881769a7	CBDE-848 Print out server response after do_mset/mget	Chibeng Hong	testrunner	master	Dec 1		✓
10817215e	Add manpages for all of cbc-* commands	Toni Nordby	libcouchbase	master	Dec 1		
1e7bf786e	Added support for libevent in win32	Brett Harrison	libcouchbase	master	Dec 1		
1c453a709	MB-7283 - docs for upgrade when hostnames configured	Steve Yen	voltron	couchbase	Dec 1		
16857c94a	MB-7271 - docs on restoring via moxi	Steve Yen	membase-cli	master	Dec 1		
117bd4c0	Sort file listing	Toni Nordby	php-ext-couchbase	master	Dec 1		✓
1677e0f4b	refactor move version to separate file	Toni Nordby	php-ext-couchbase	master	Dec 1		✓
10ec0bfea	move stats to a separate file	Toni Nordby	php-ext-couchbase	master	Dec 1		✓
12c550897	Simplify flush logic	Toni Nordby	php-ext-couchbase	master	Dec 1		✓
1d663eb35	wip using wait_backfill complete	Alaksyey Kandratsenka	ns_server	master	Nov 30		
1840b348c	libbucketmigrator added ability to wait for end of backfill	Alaksyey Kandratsenka	ns_server	master	Nov 30		
13189b67d	MB-6726 added internal settings for vbucket scheduling variables	Alaksyey Kandratsenka	ns_server	master	Nov 30		
13bb5c065	MB-100 store just binaries in master activity events keeper	Alaksyey Kandratsenka	ns_server	master	Nov 30		✓
15c76ea3b	MB-100 added scripts/reformat-master-events.rb	Alaksyey Kandratsenka	ns_server	master	Nov 30		
15410d114	MB-6726 added script to draw timeline of new style rebalance	Alaksyey Kandratsenka	ns_server	master	Nov 30		
1d213a884	MB-6726 added vbucket scheduling simulation/visualization tool	Alaksyey Kandratsenka	ns_server	master	Nov 30		
1ac0db1b	MB-6726 added master activity events for in-rebalance compaction	Alaksyey Kandratsenka	ns_server	master	Nov 30		
1b4433876	MB-6726 added rebalance ignore_view_compactions setting	Alaksyey Kandratsenka	ns_server	master	Nov 30		
1d66aa7d9	MB-6726 avoid doing obviously duplicate set view config changes	Alaksyey Kandratsenka	ns_server	master	Nov 30		
1a26da1be	MB-6726.999 implemented index aware vbucket move scheduling	Alaksyey Kandratsenka	ns_server	master	Nov 30		
1a5154979	MB-100 added master events support for new style rebalance	Alaksyey Kandratsenka	ns_server	master	Nov 30		
19cc836fa	MB-7285 Fix badmatch in index updater on cleanup	Elipe David Boba Manana	couchdb	2.0.0	Nov 30	✓	✓

Next: Press 'T' to view keyboard shortcuts
Powered by [Gerrit Code Review](#) (2.4.2) | [Report Bug](#)


<https://github.com/couchbase/>

<https://github.com/couchbaselabs/>

Gerrit: <http://review.couchbase.org/>

Couchbase

There is an elephant in the room



Couchbase


What about Couch**?

- **Couchbase & CouchDB are different products**

- Different Open Source Projects
- Different architecture and goals

More info at :

<http://www.couchbase.com/couchdb>

Comparing Apache CouchDB and Couchbase		
	 CouchDB	Couchbase
Document-oriented (JSON) database	●	●
100% open source software	●	●
Incremental JavaScript map-reduce	●	●
Data synchronization <small>CouchDB to Couchbase synchronization not supported</small>	●	●
Highly-reliable storage architecture	●	●
HTTP API	●	
CouchApps	●	
Auto-sharding cluster technology		●
Non-stop operation <small>Live cluster reconfiguration and database upgrades</small>		●
Professional SDKs		●
Memcached-compatible API		●
Built-in data caching		●

Couchbase

Couchbase



Couchbase

SETTING UP TO DEVELOP

Couchbase

Install Couchbase Server 2.0



Ubuntu



RedHat



Mac OS X



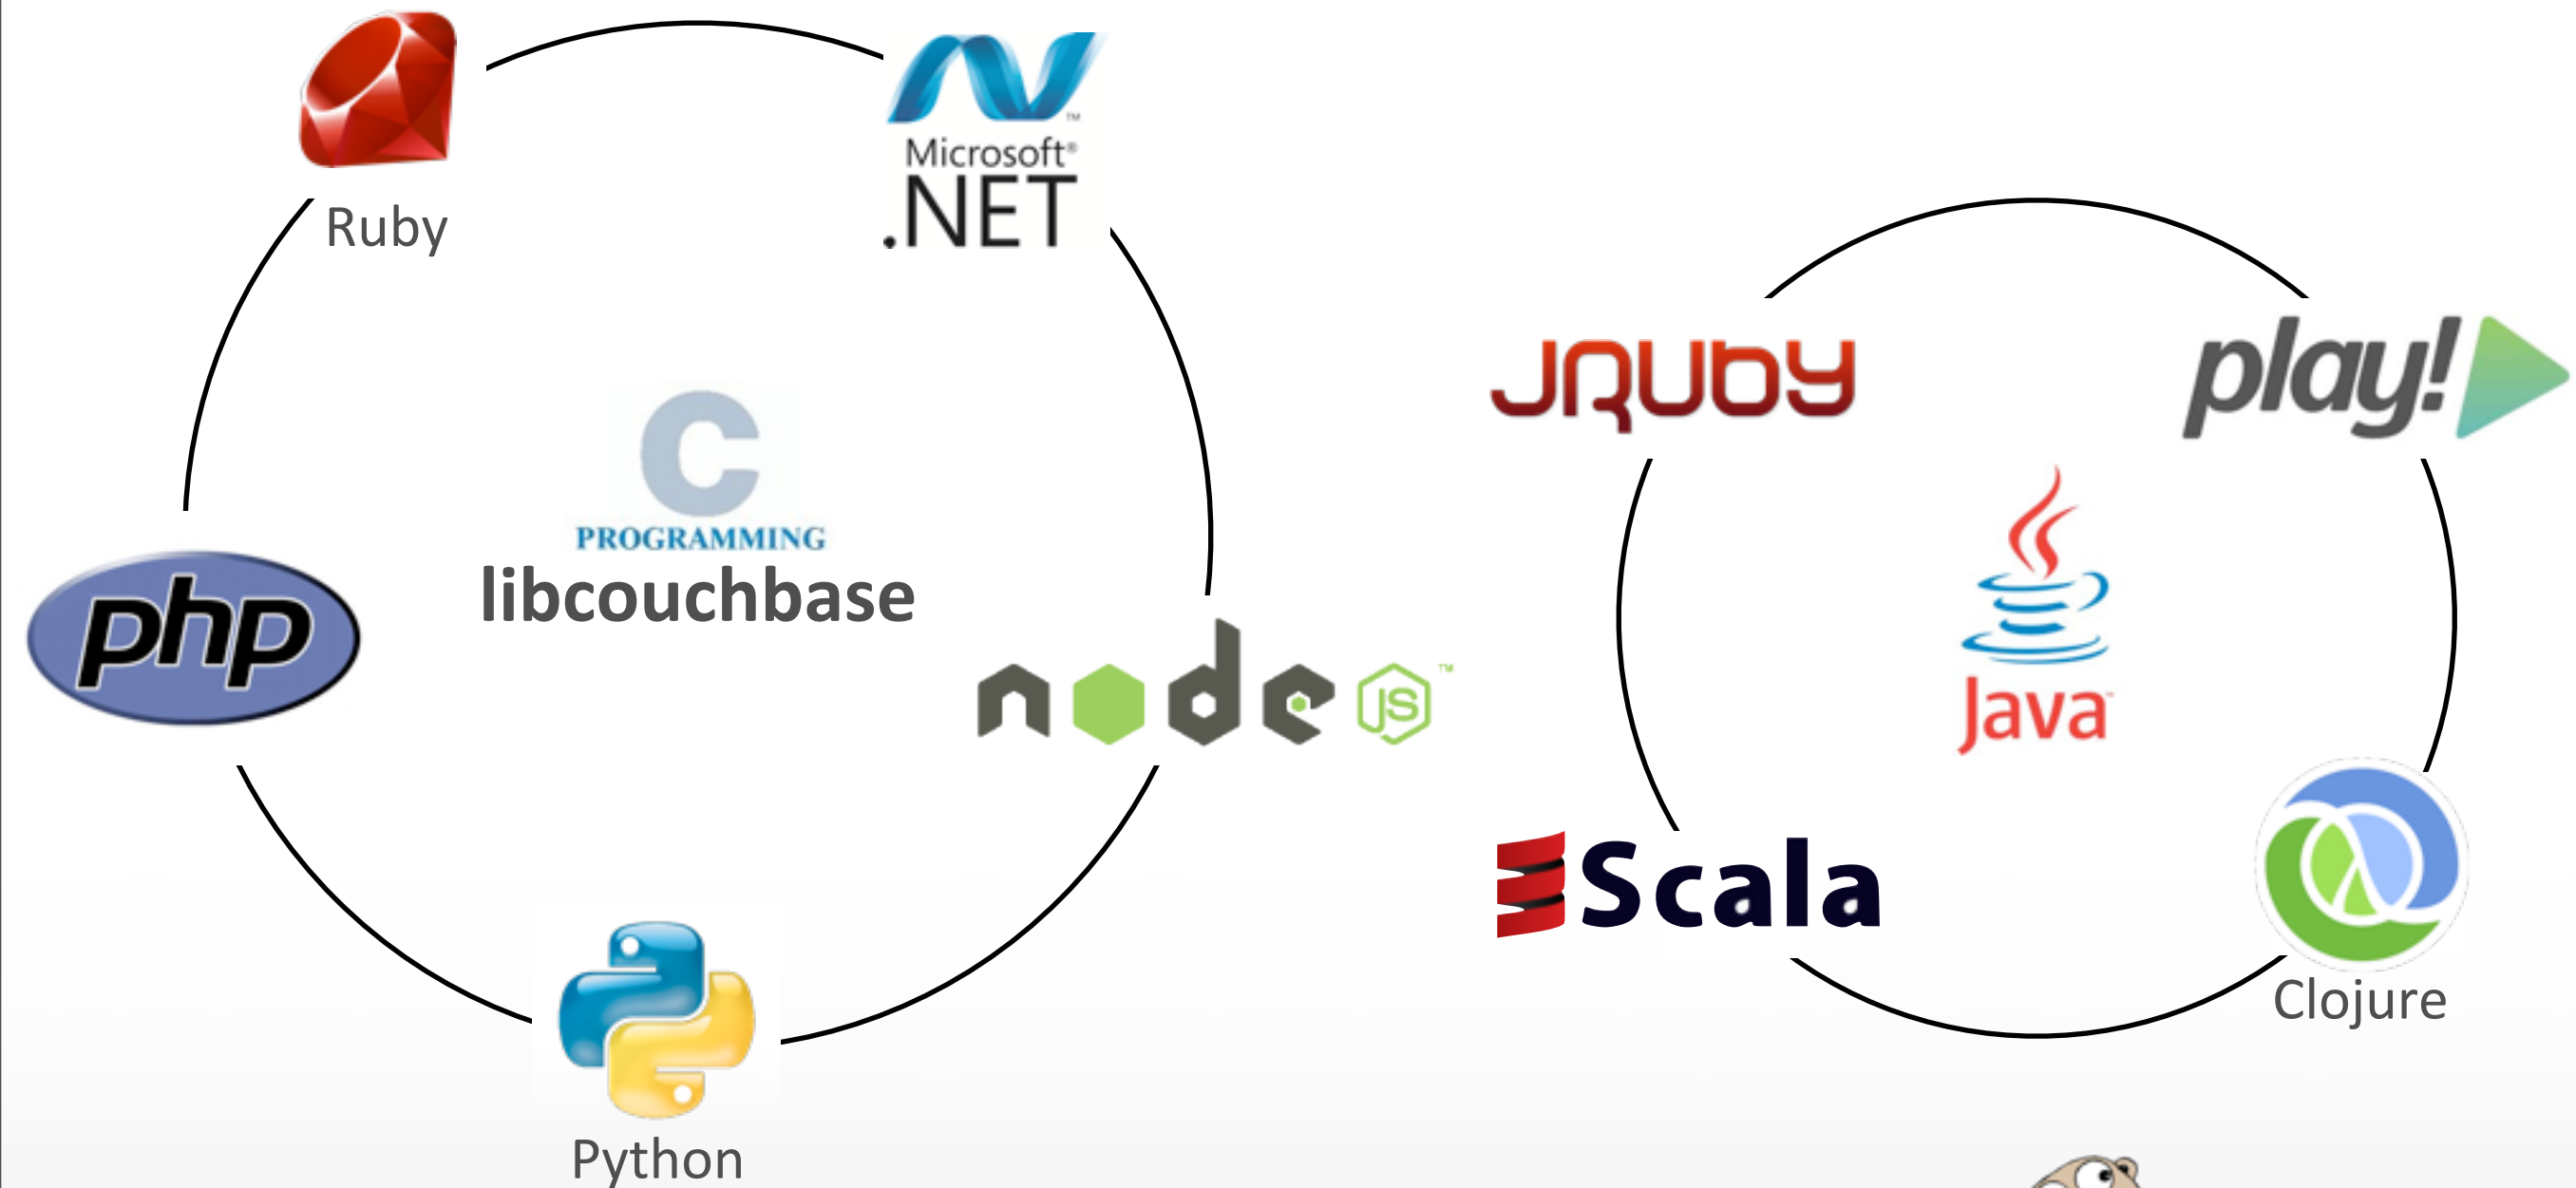
Windows

or build from sources



Couchbase

Official SDKs



www.couchbase.com/develop



Go

Couchbase

Client SDKs

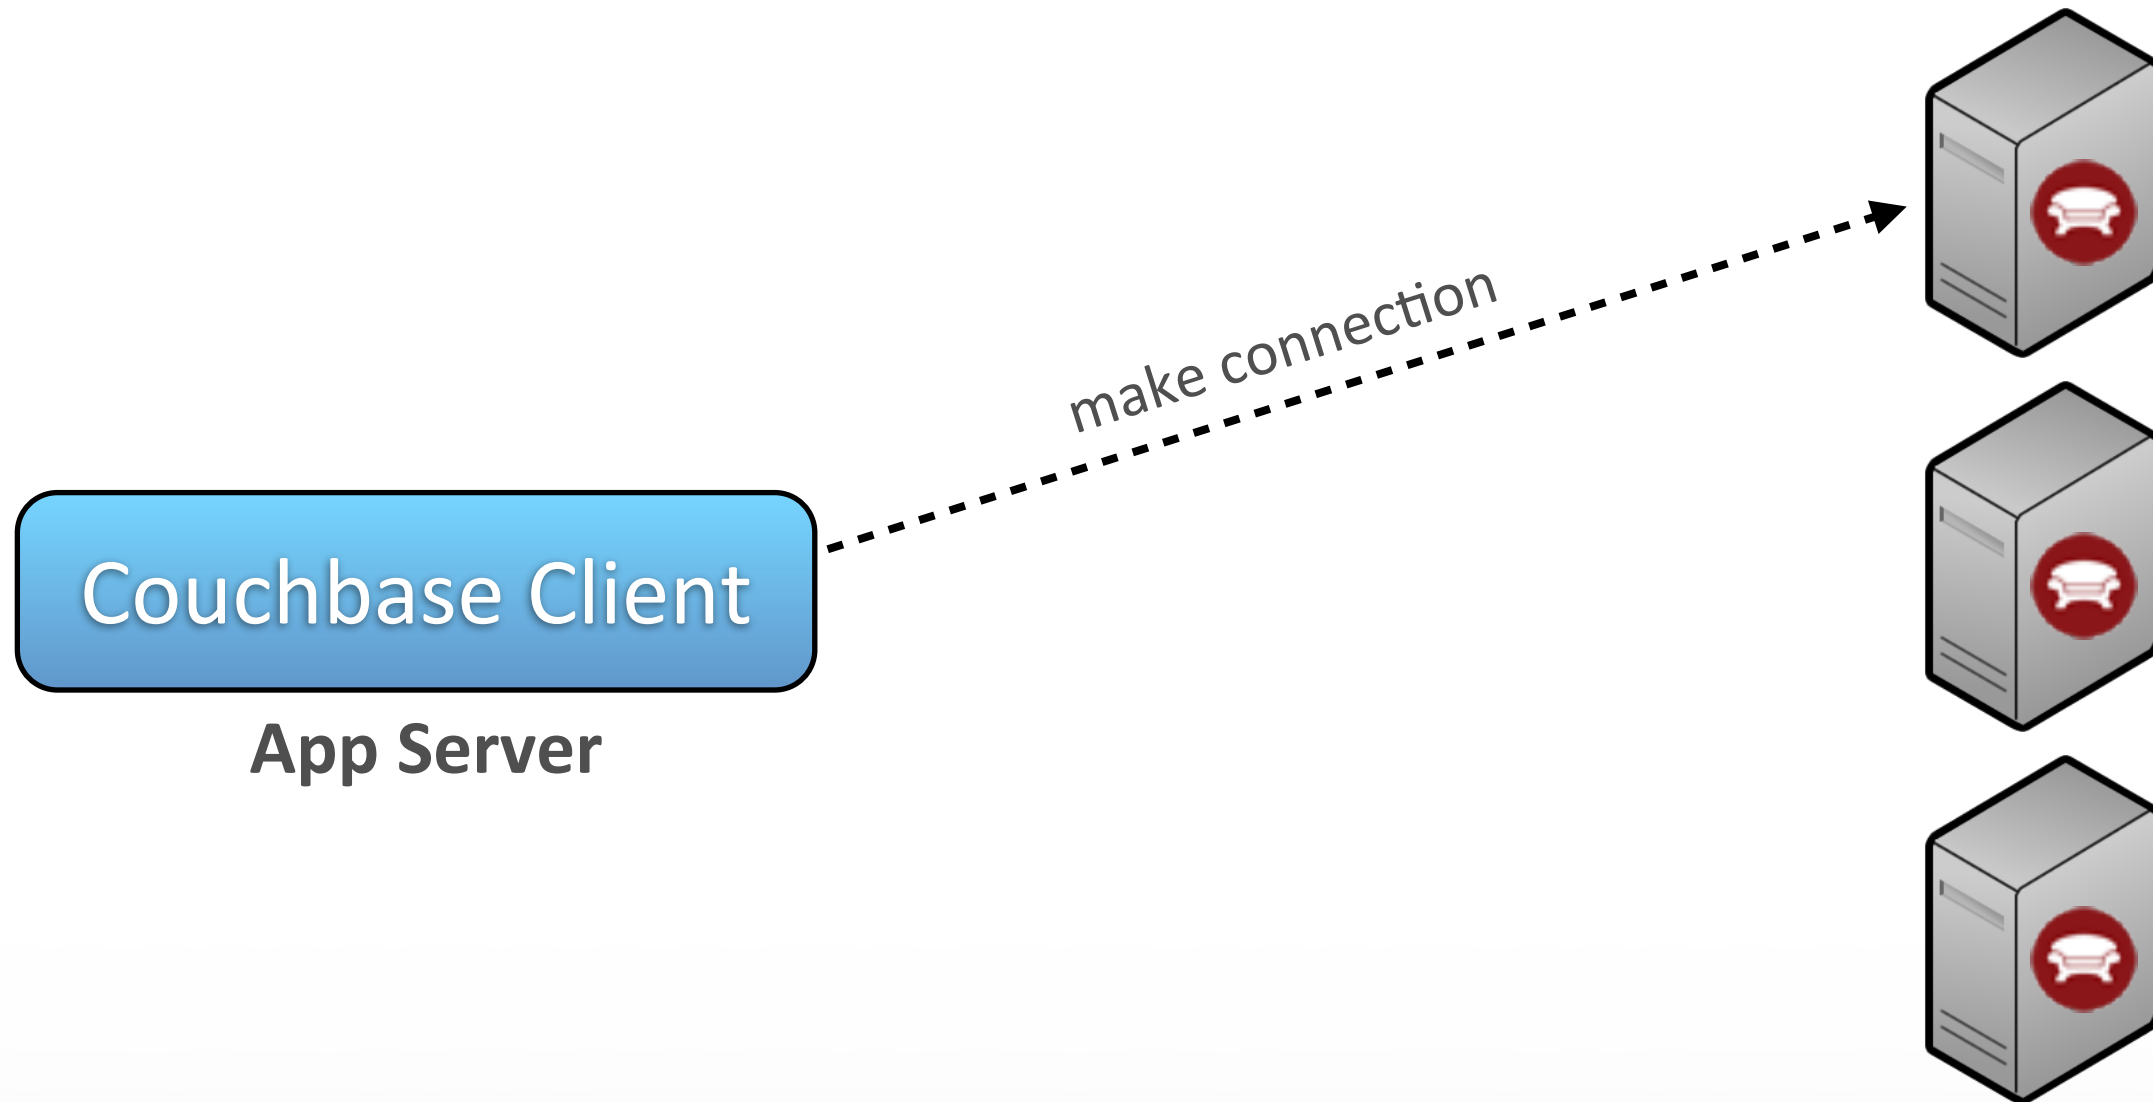
Couchbase Client

App Server



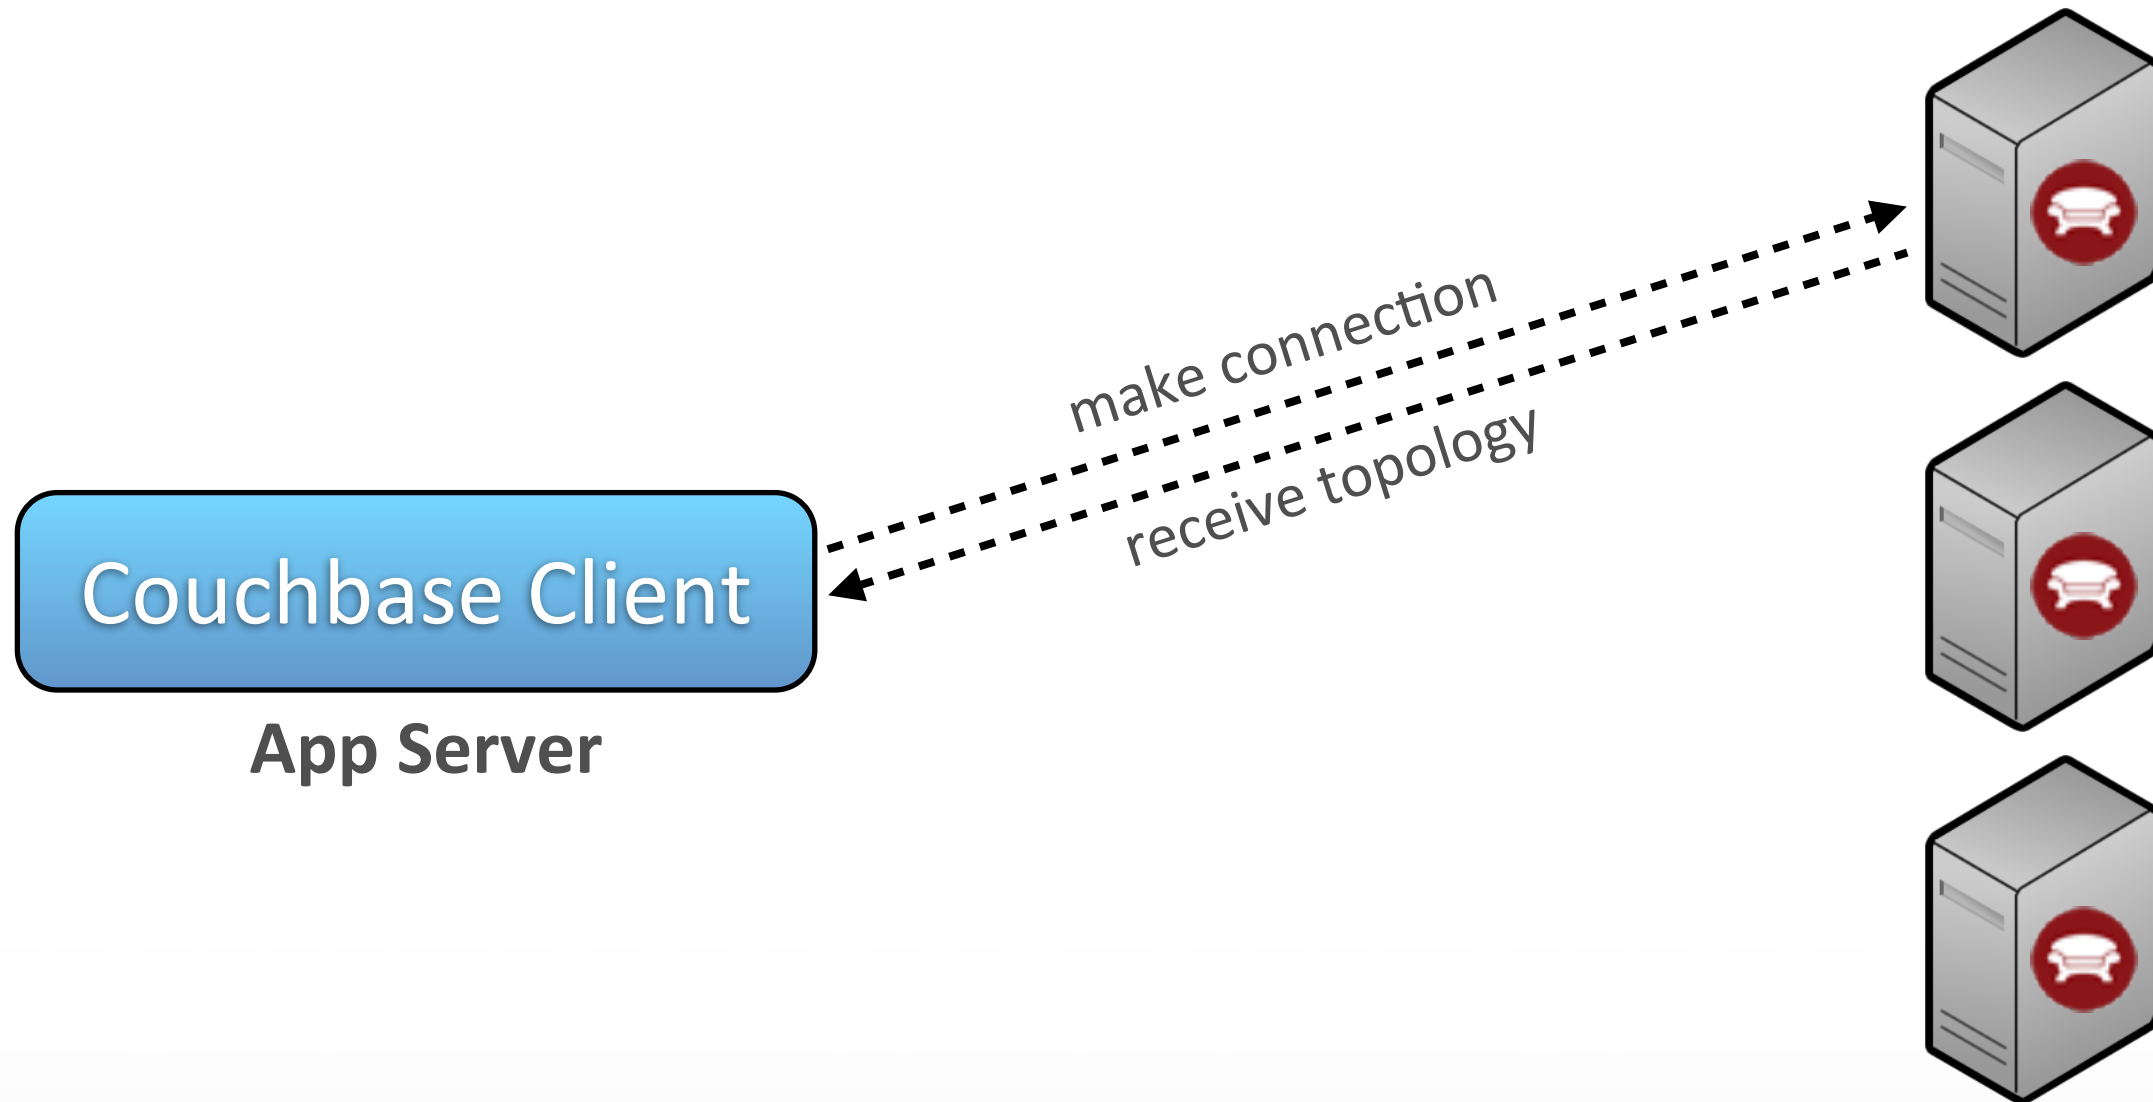
Couchbase

Client SDKs



Couchbase

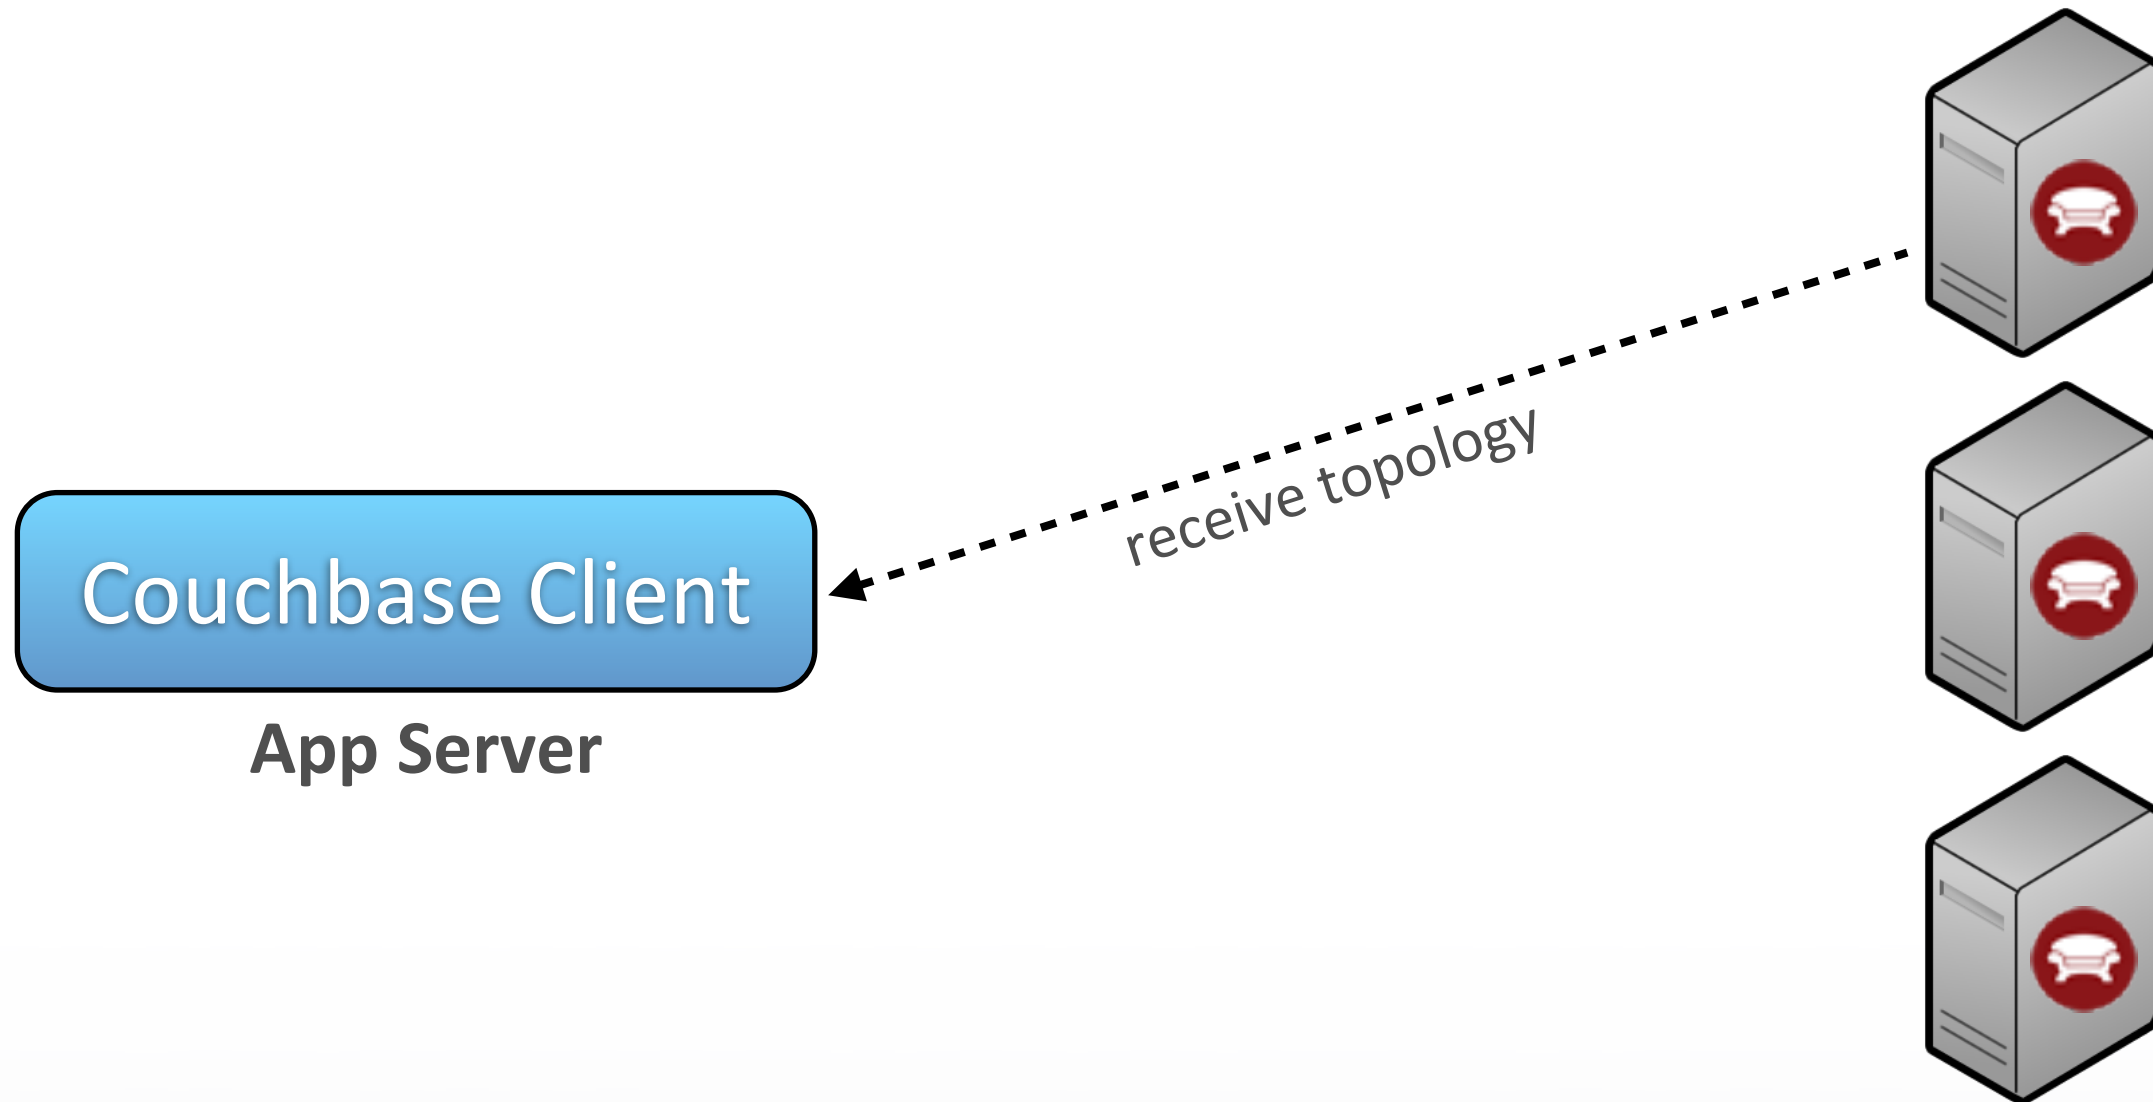
Client SDKs



App Server

Couchbase

Client SDKs



Couchbase

Client SDKs

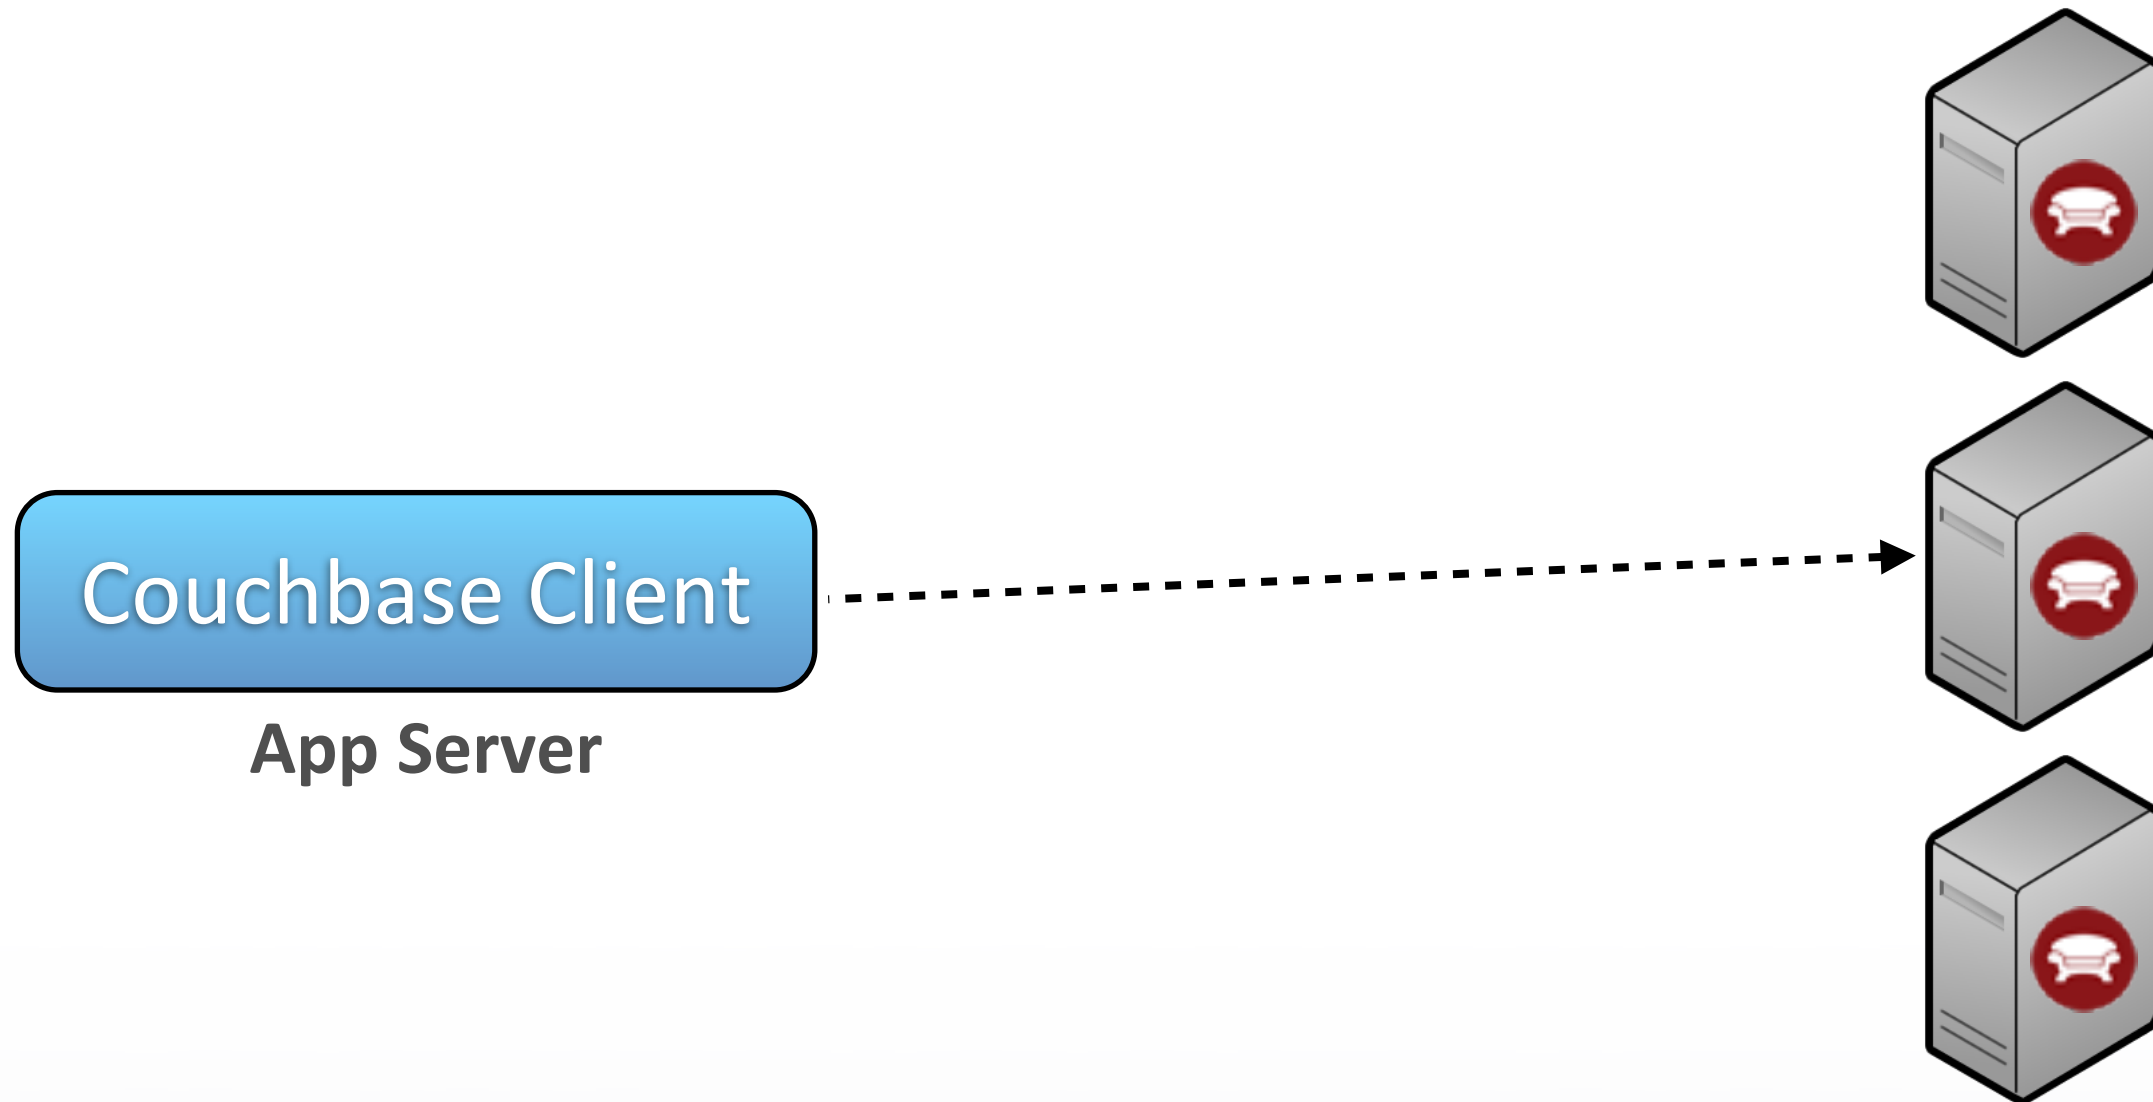
Couchbase Client

App Server



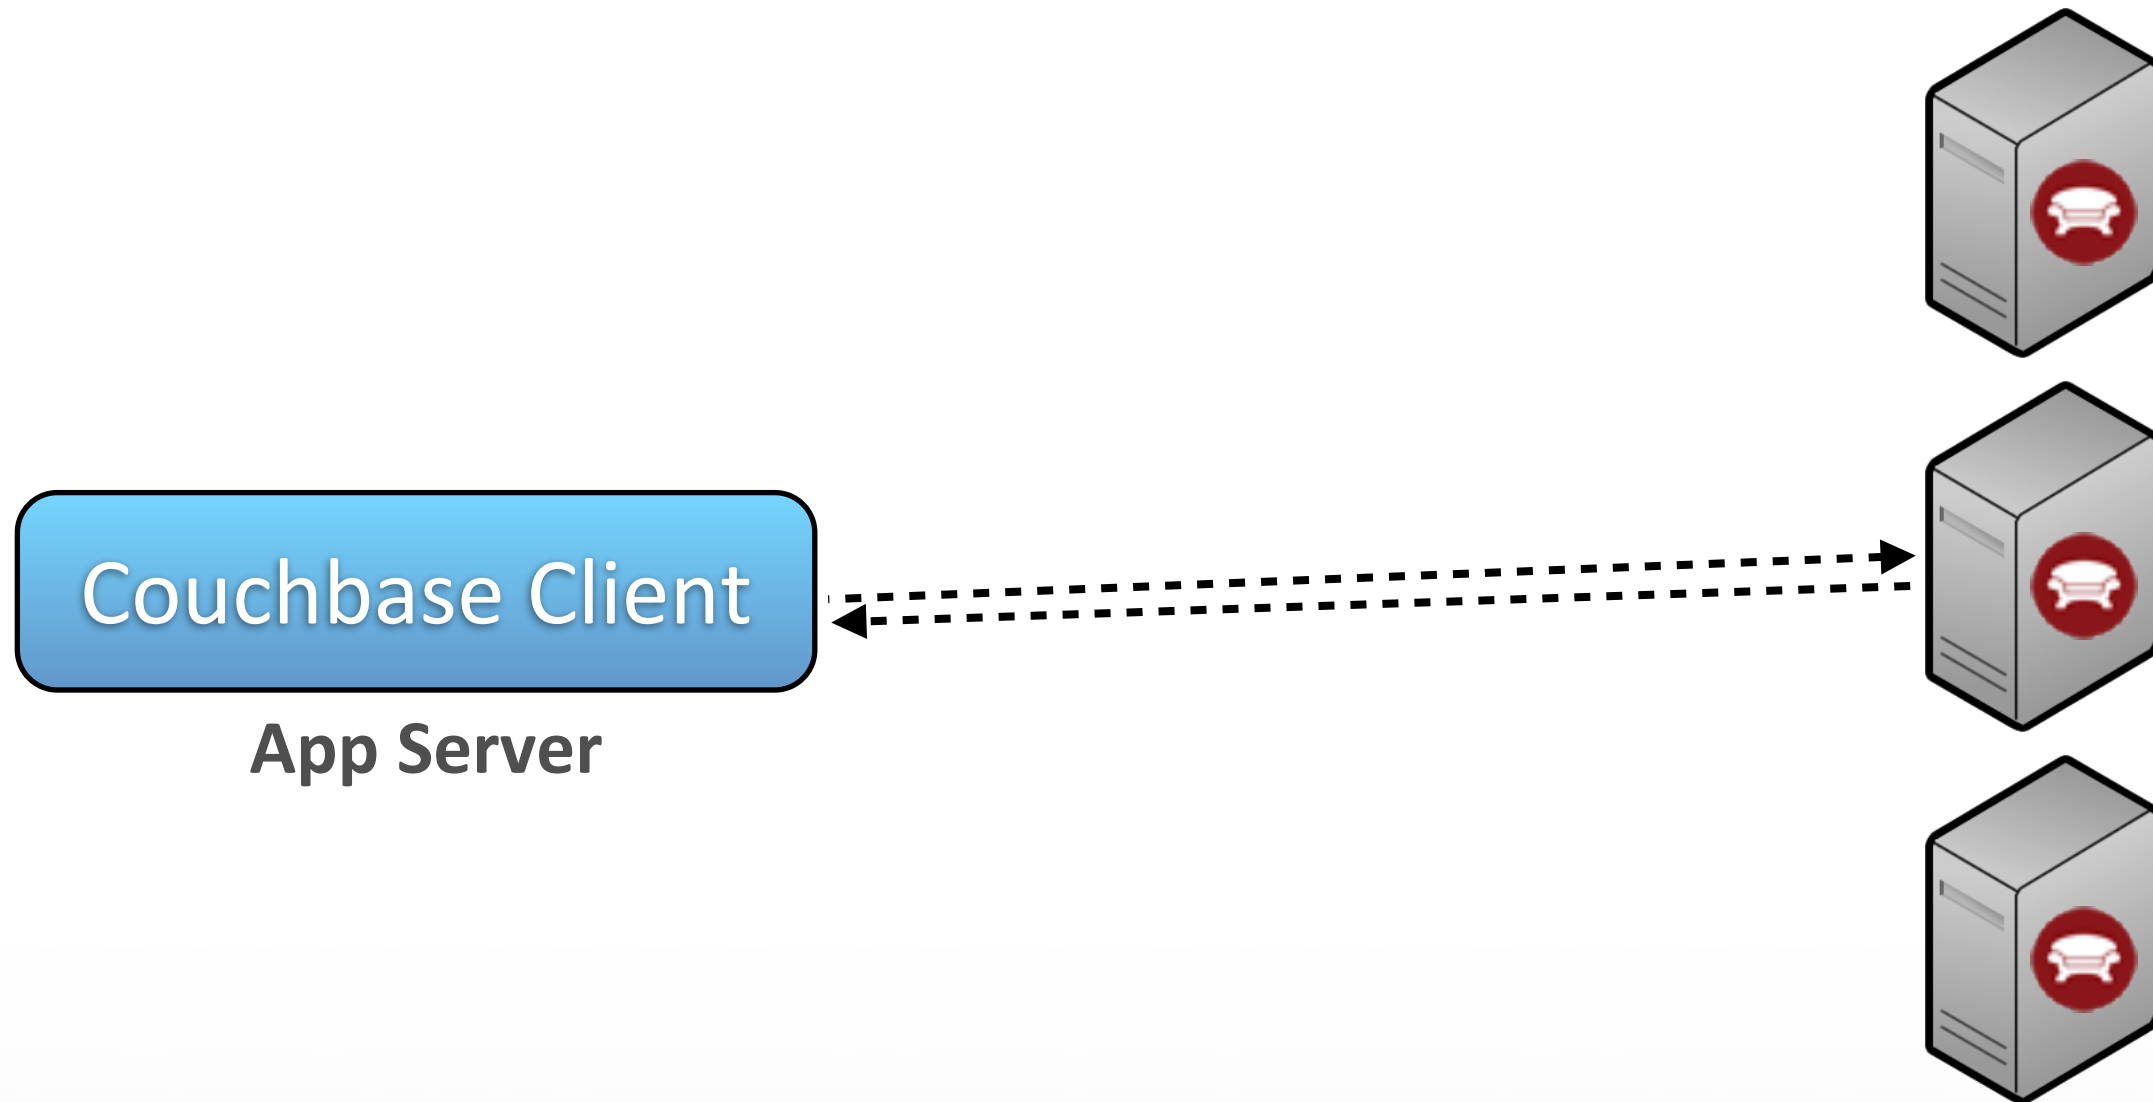
Couchbase

Client SDKs



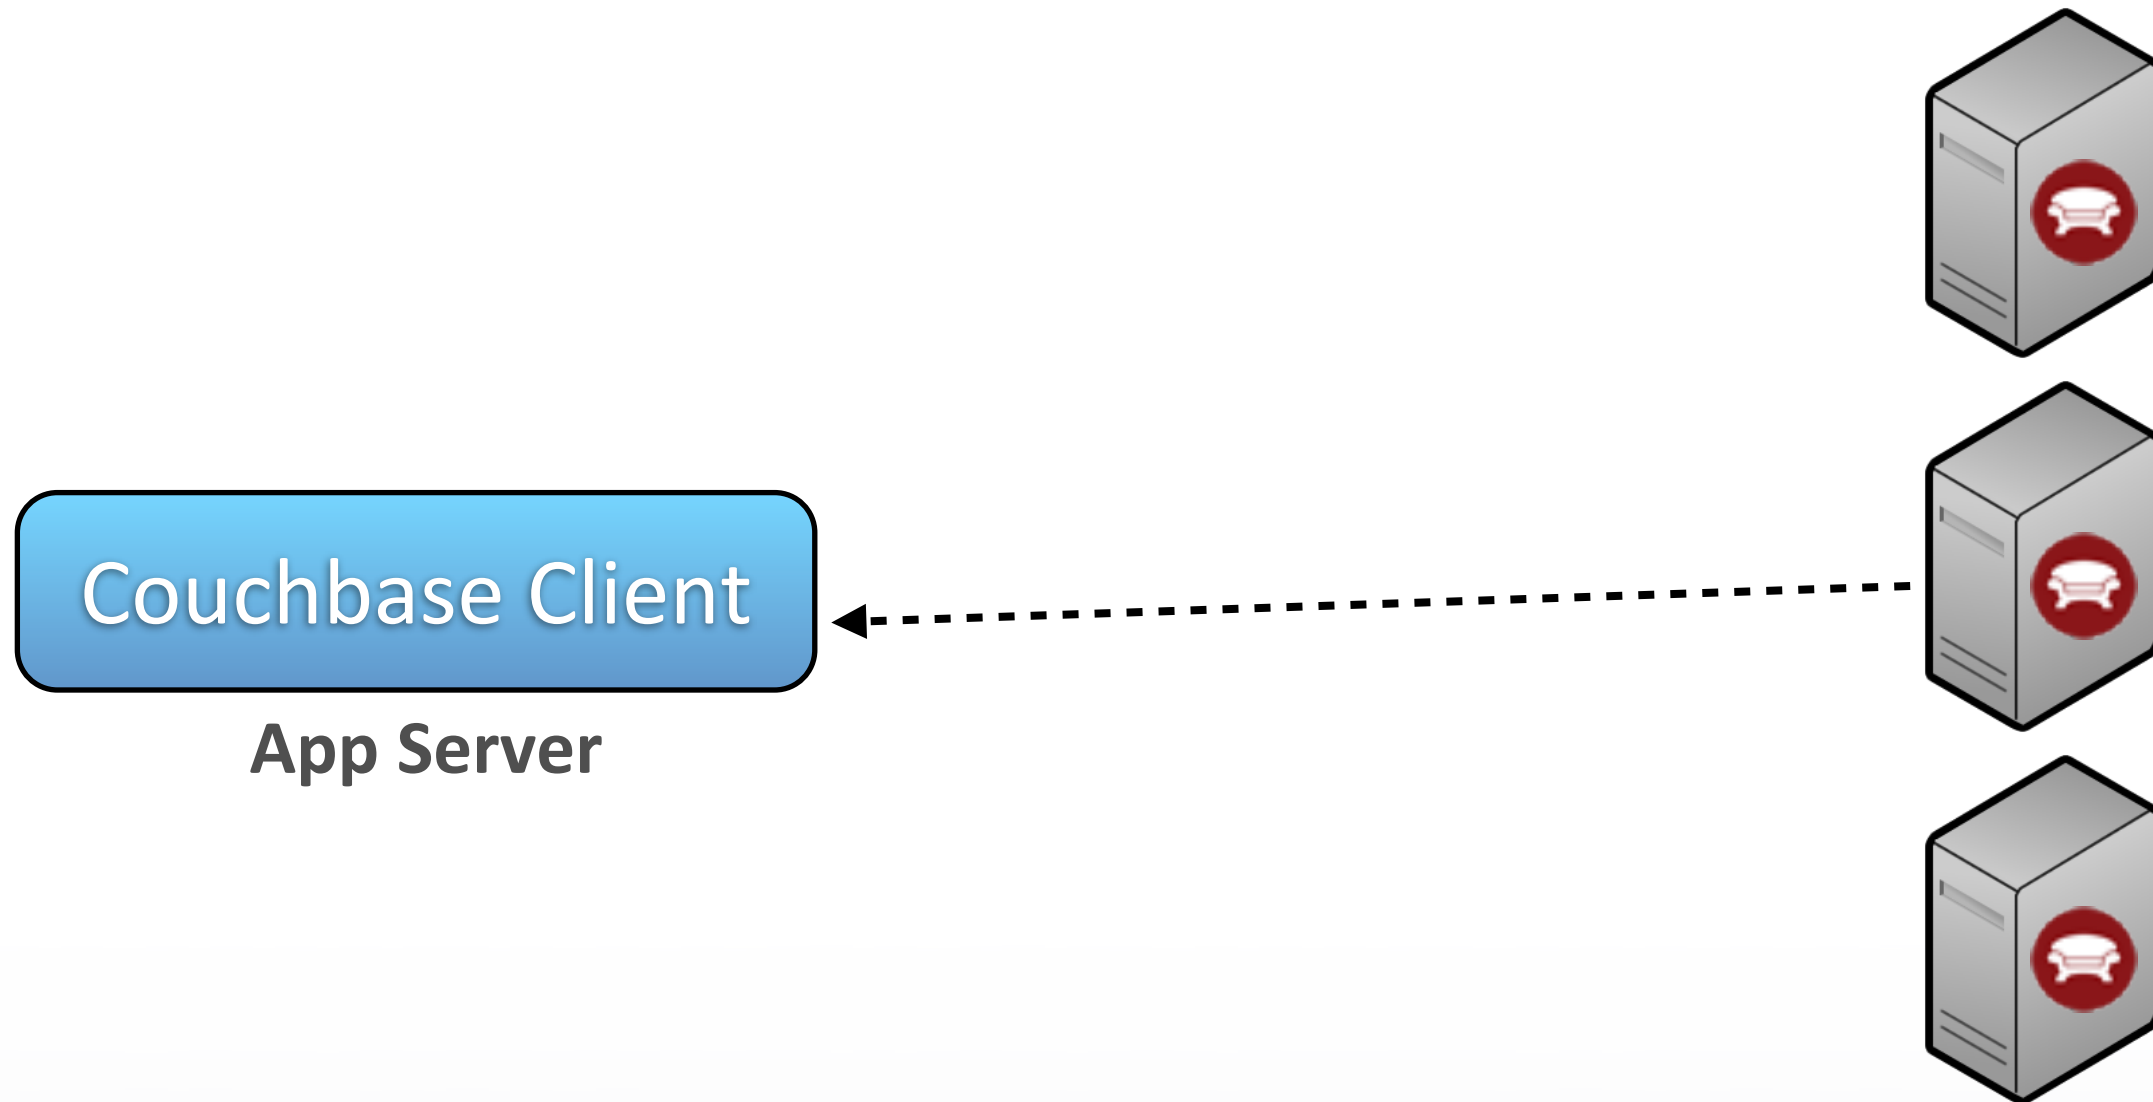
Couchbase

Client SDKs



Couchbase

Client SDKs



Couchbase

Client SDKs

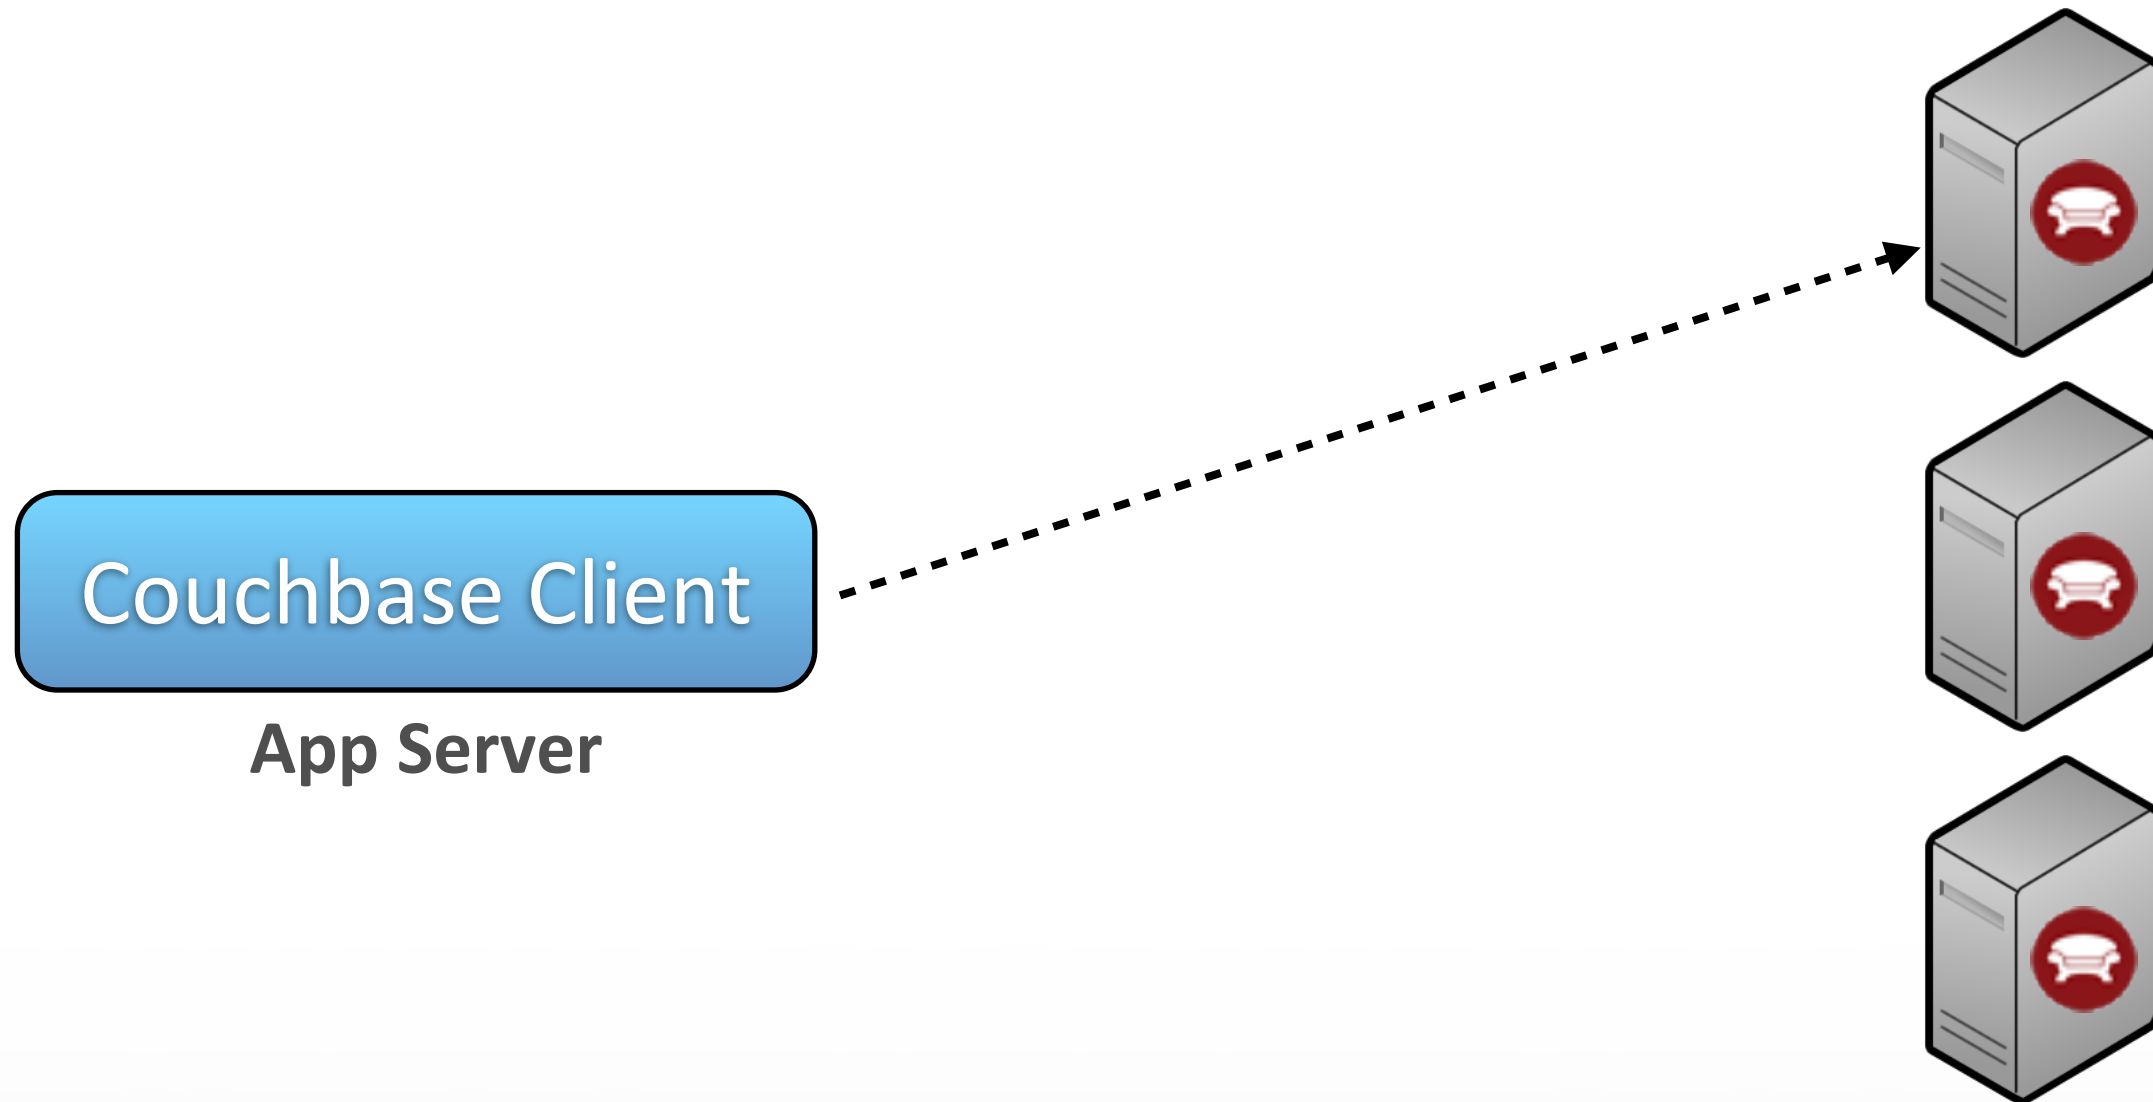
Couchbase Client

App Server



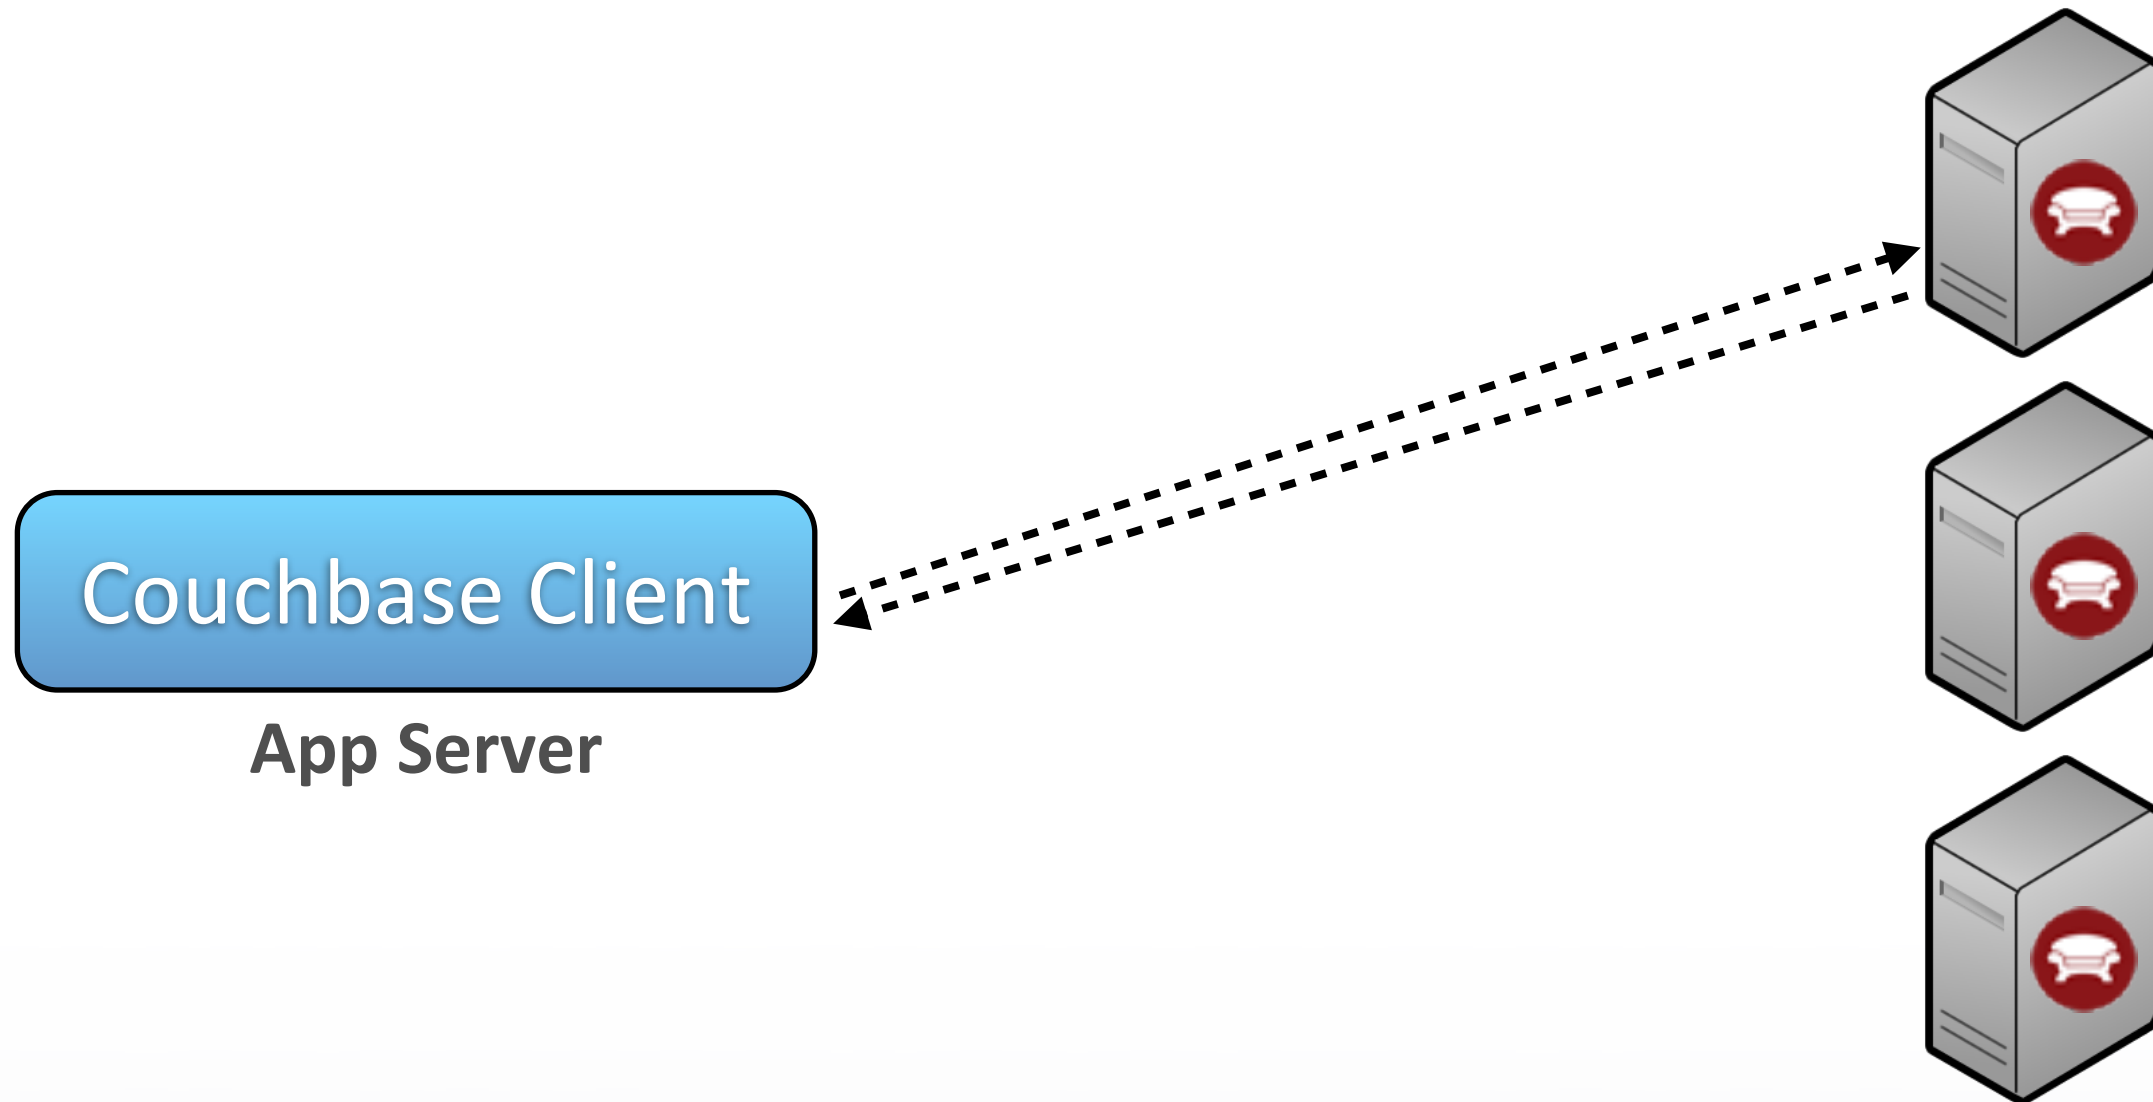
Couchbase

Client SDKs



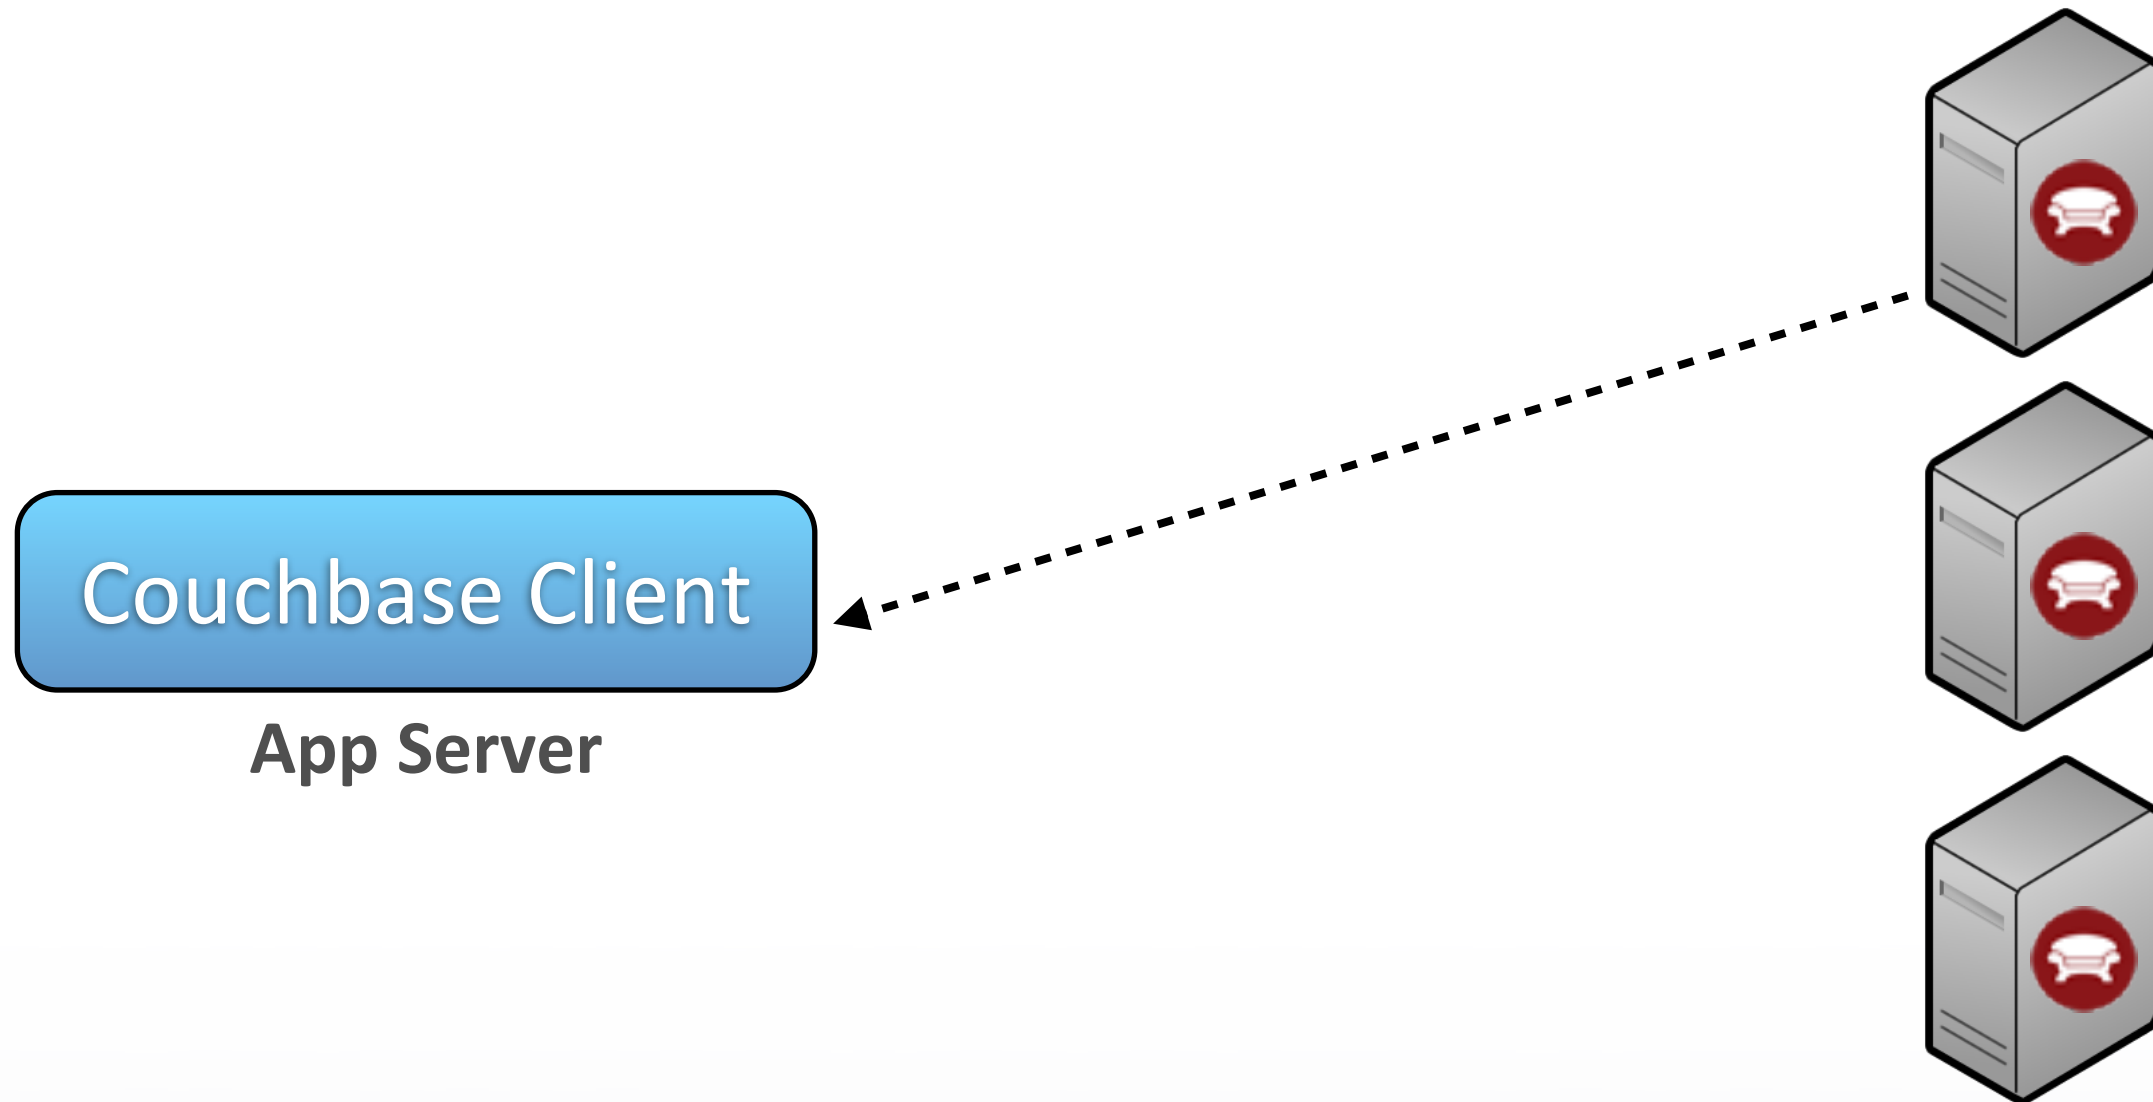
Couchbase

Client SDKs



Couchbase

Client SDKs



Couchbase

Client SDKs

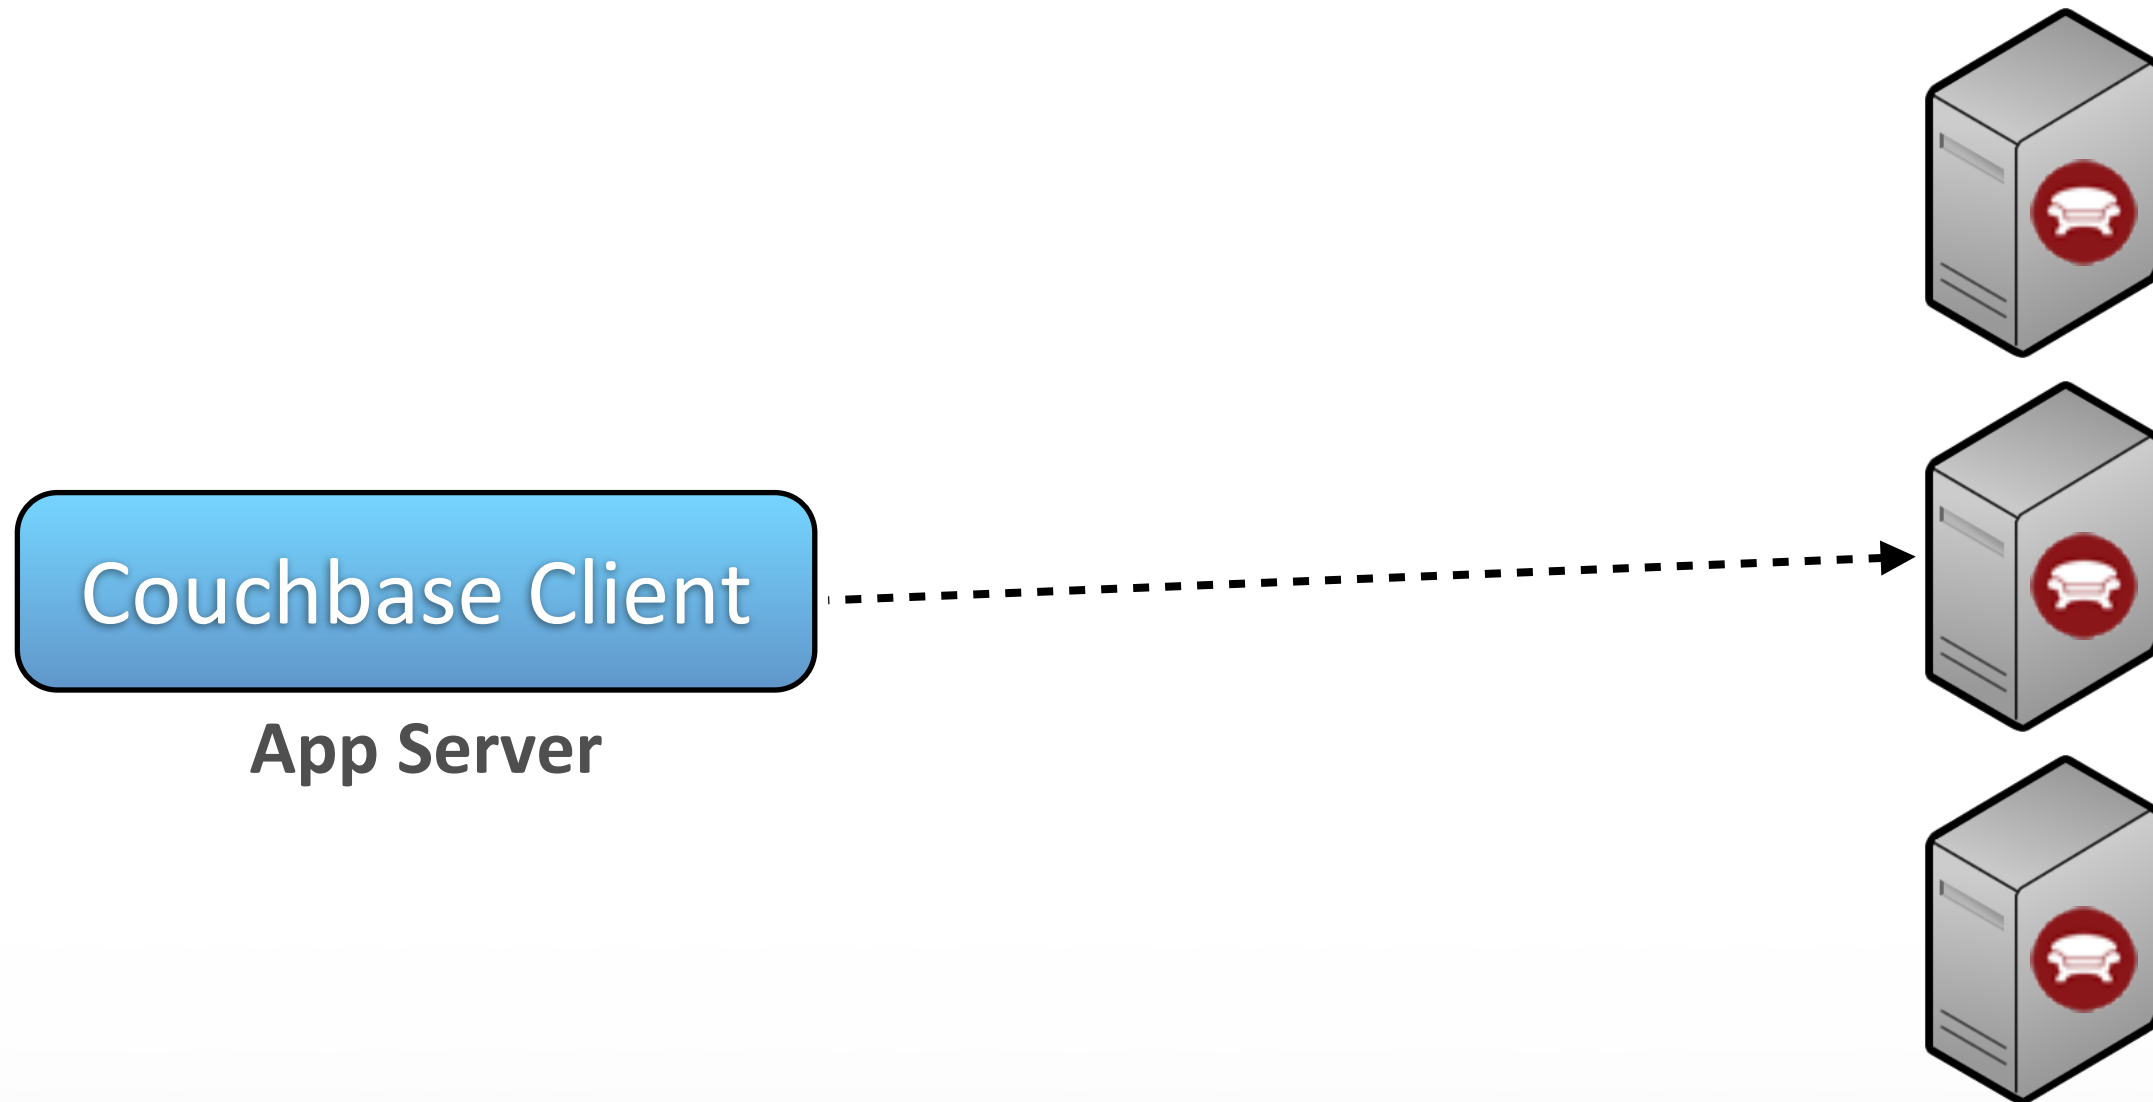
Couchbase Client

App Server



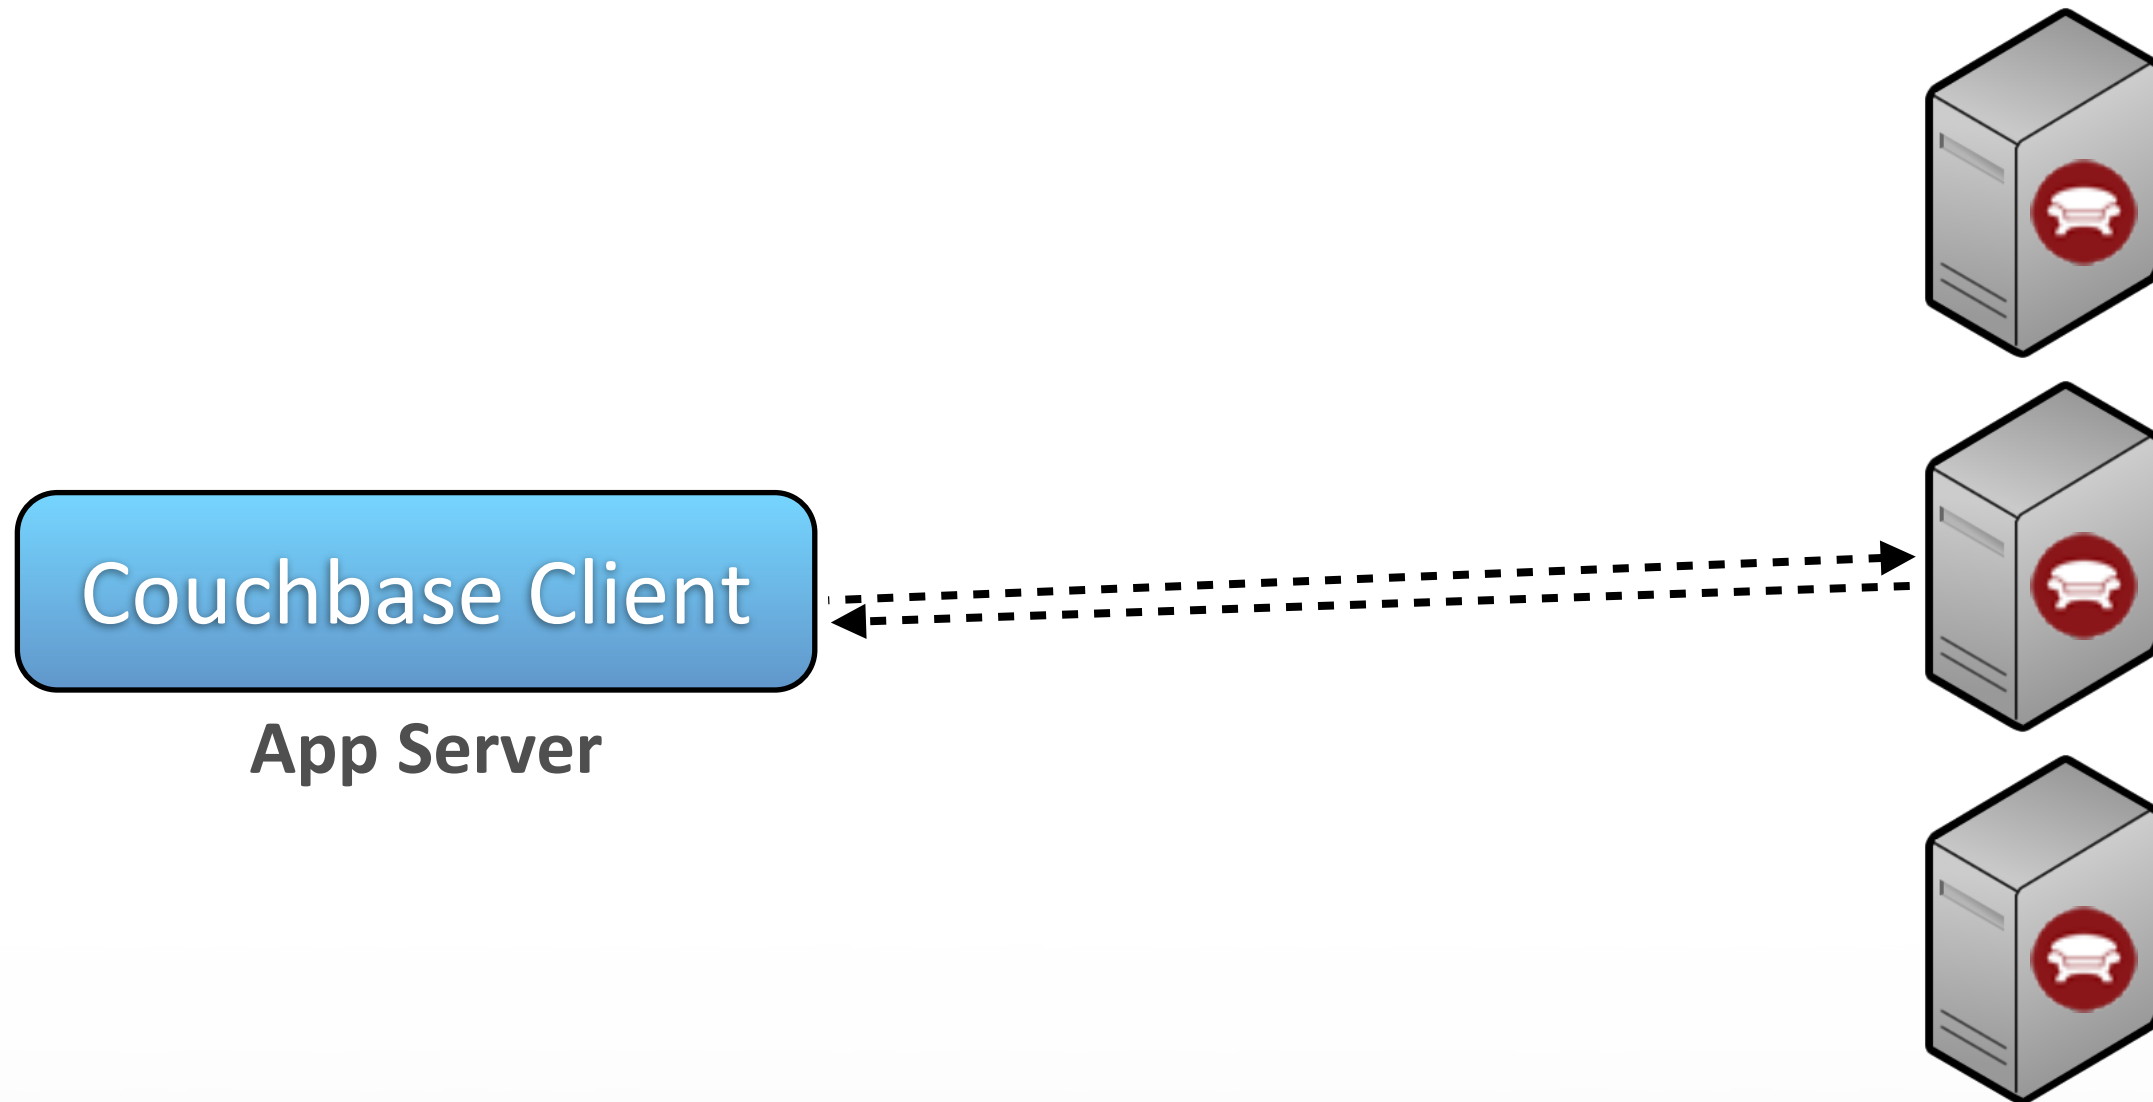
Couchbase

Client SDKs



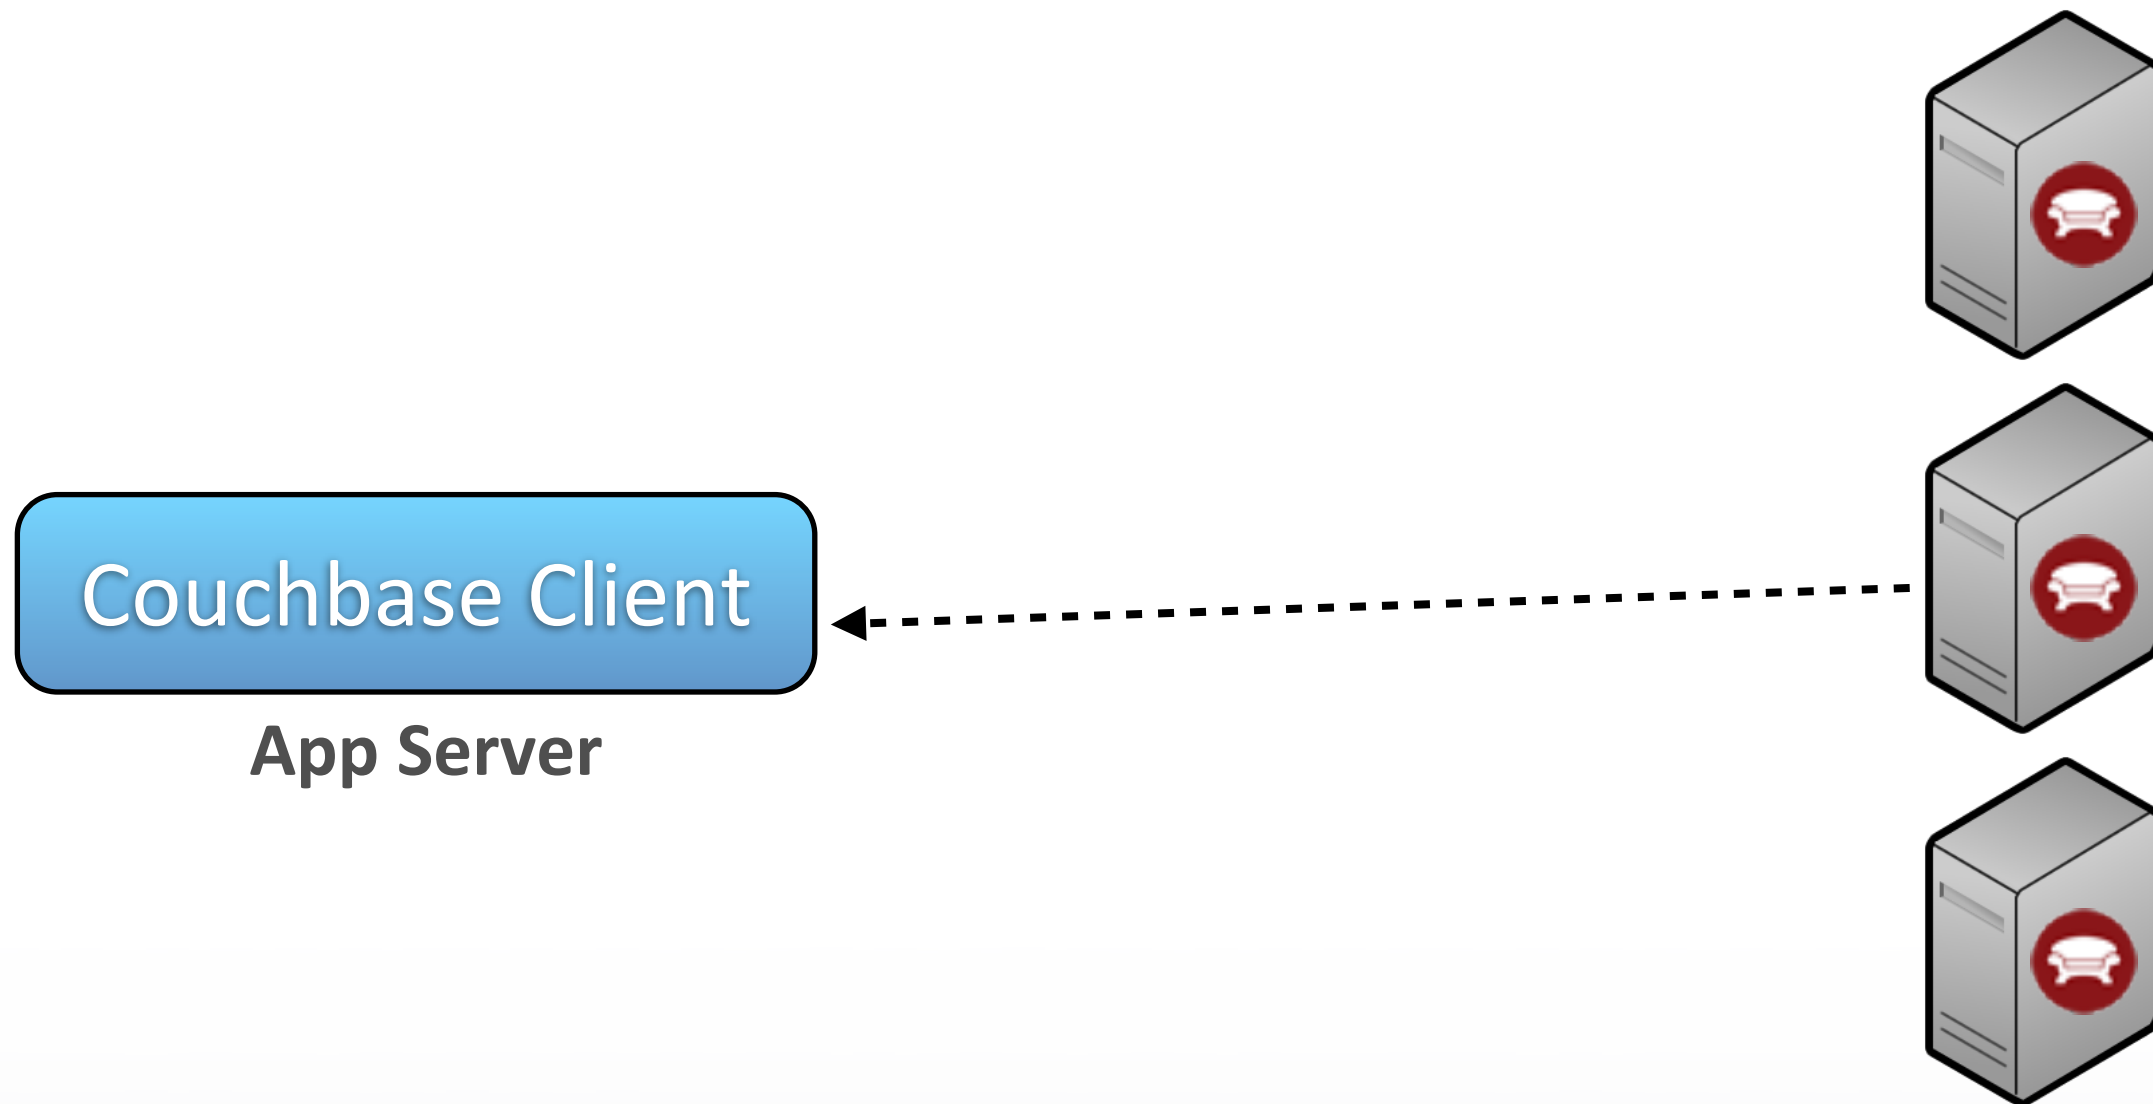
Couchbase

Client SDKs



Couchbase

Client SDKs



Couchbase

Client SDKs

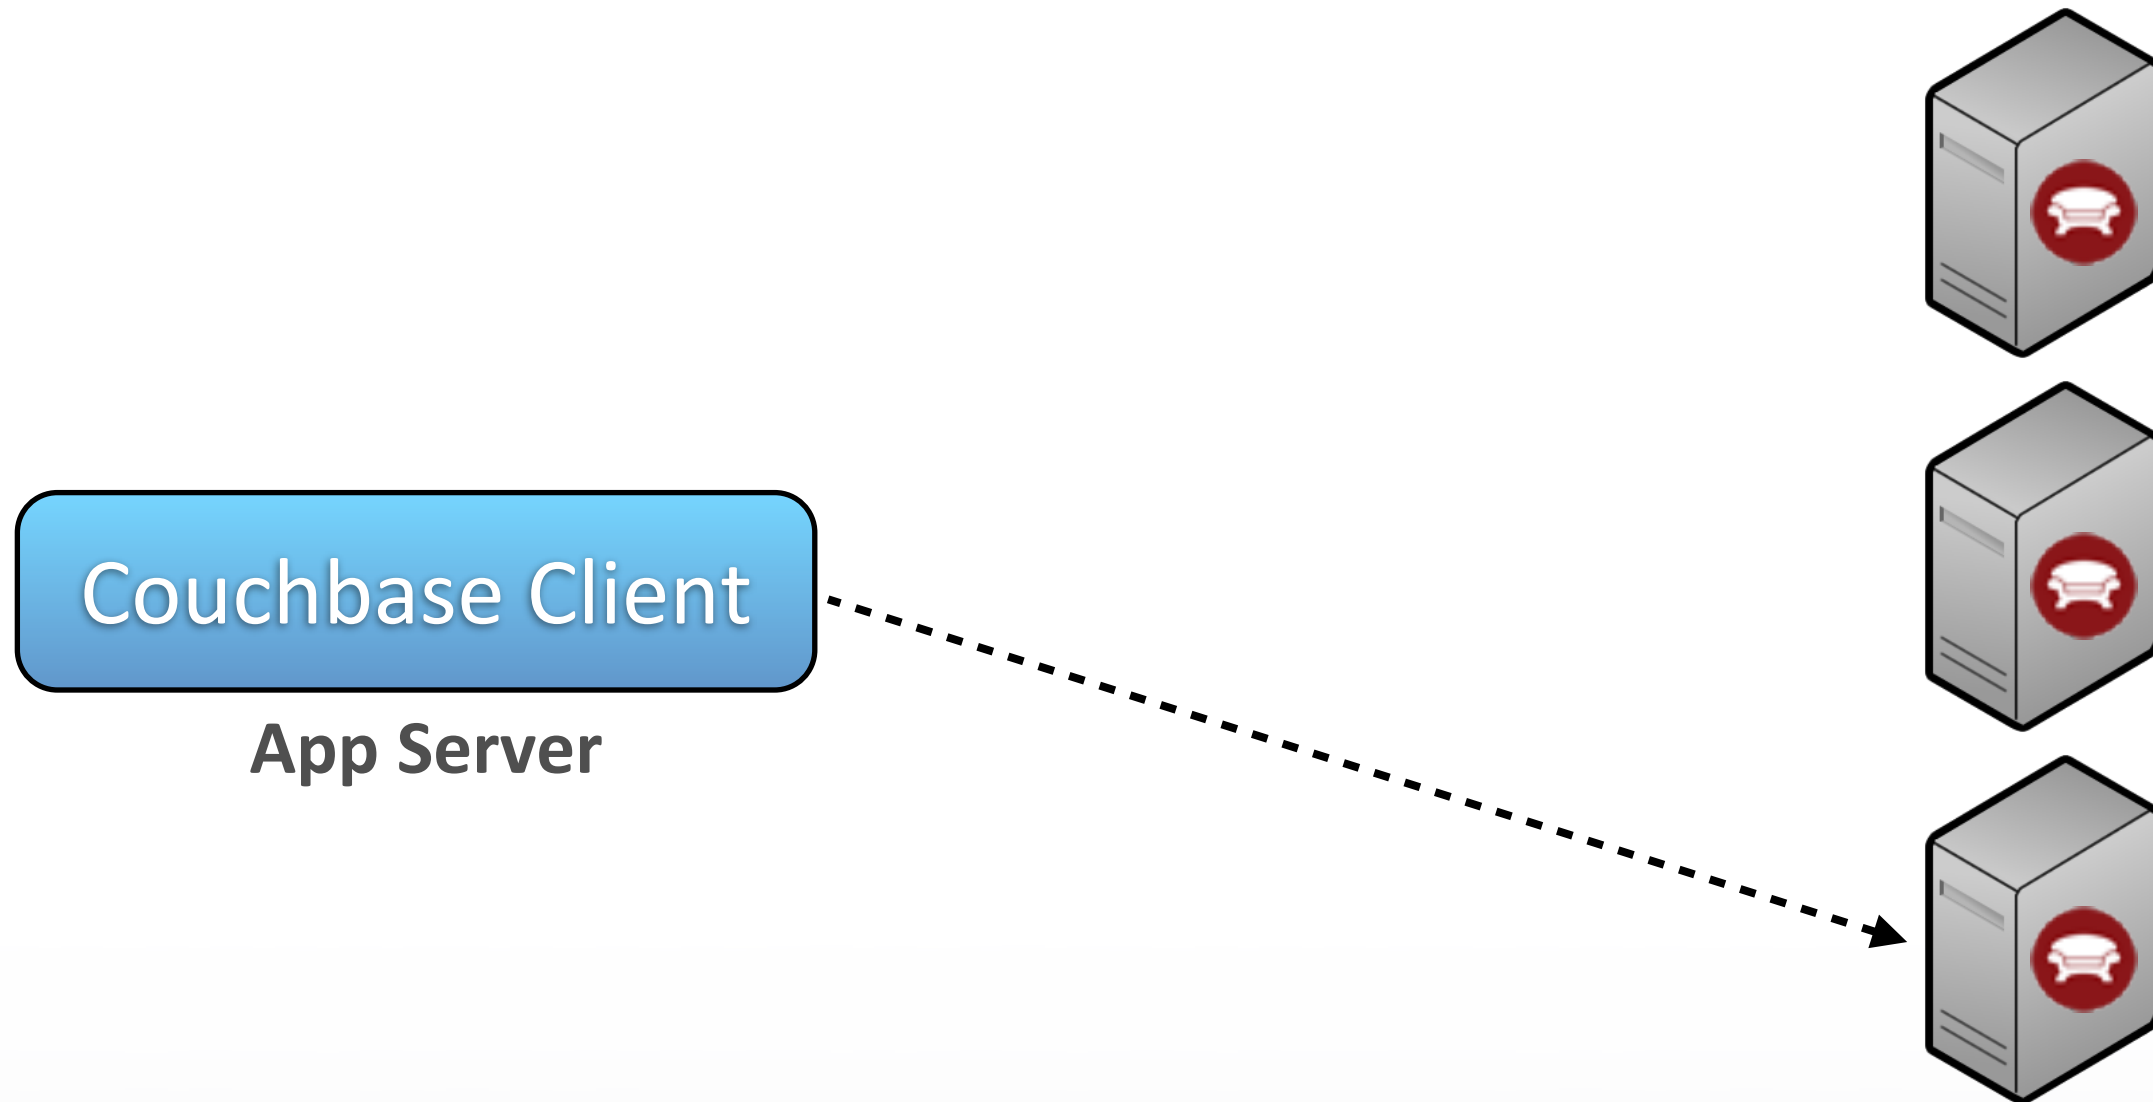
Couchbase Client

App Server



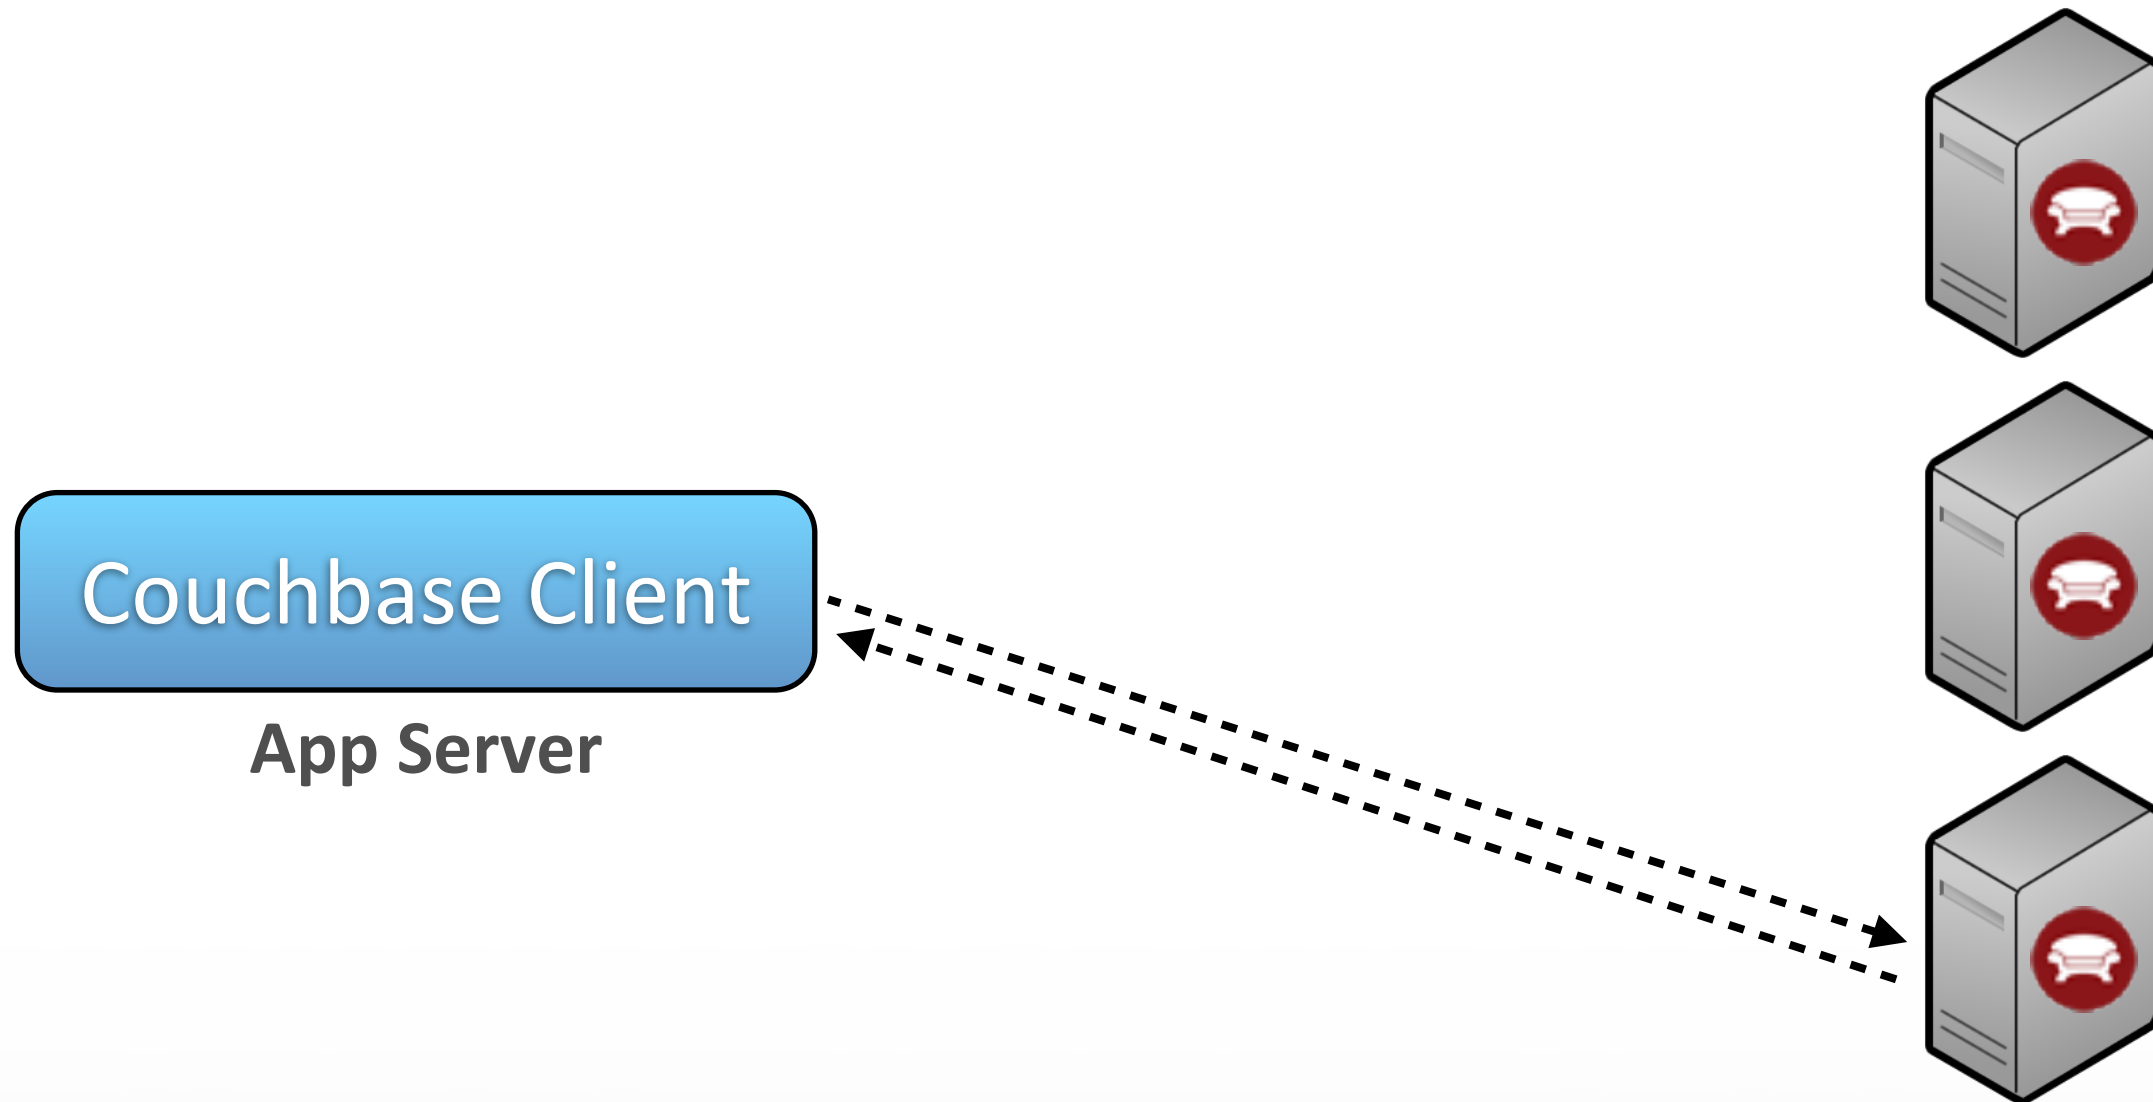
Couchbase

Client SDKs



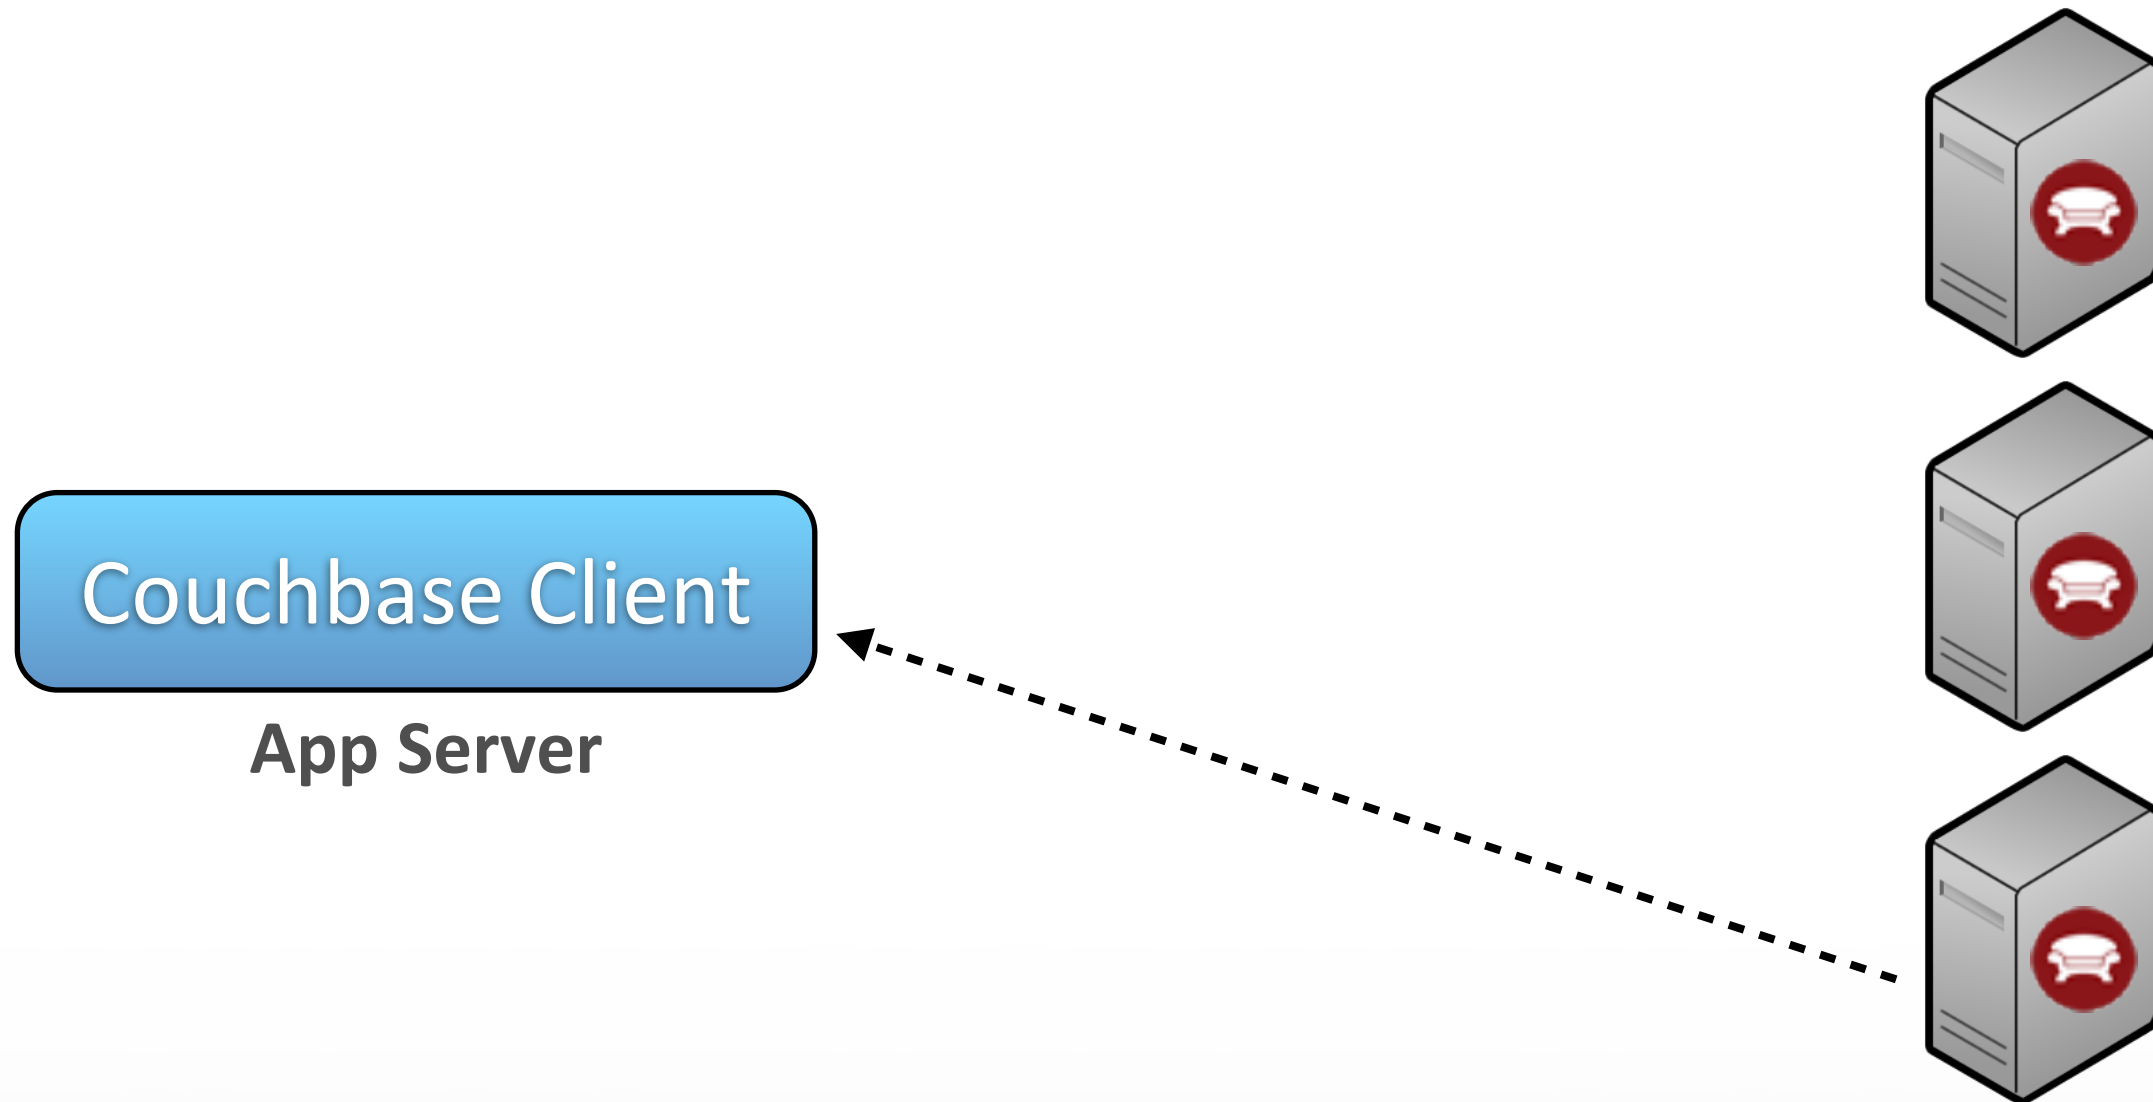
Couchbase

Client SDKs



Couchbase

Client SDKs



Couchbase

Client SDKs

Couchbase Client

App Server




Couchbase

Client SDKs

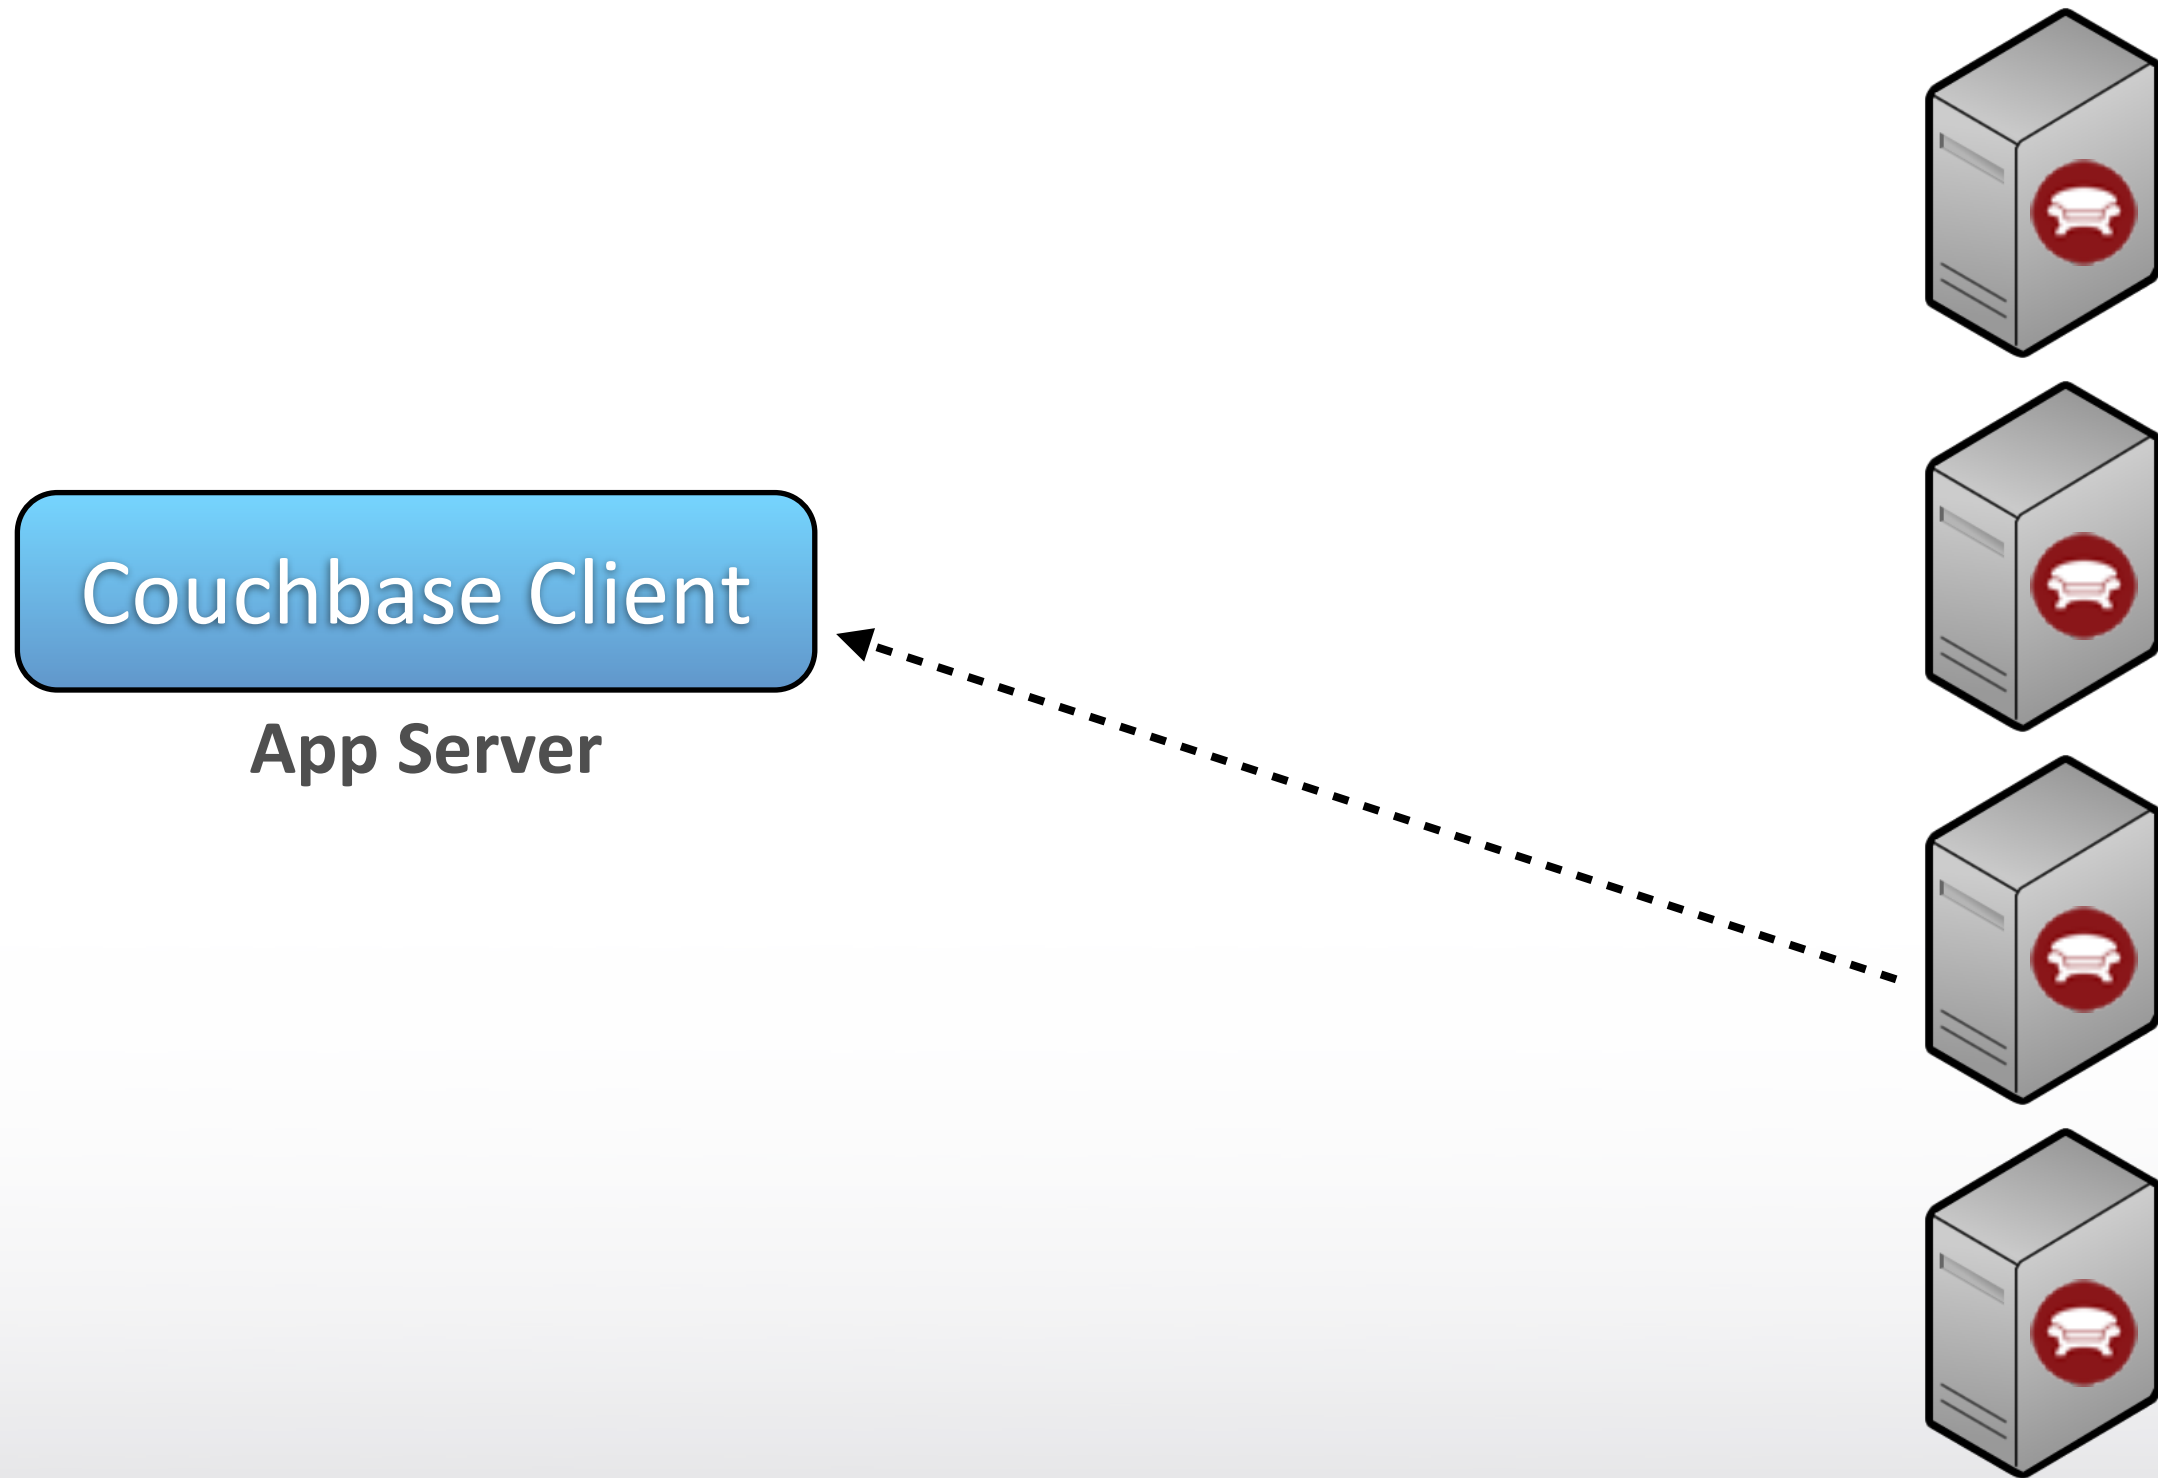
Couchbase Client

App Server




Couchbase

Client SDKs



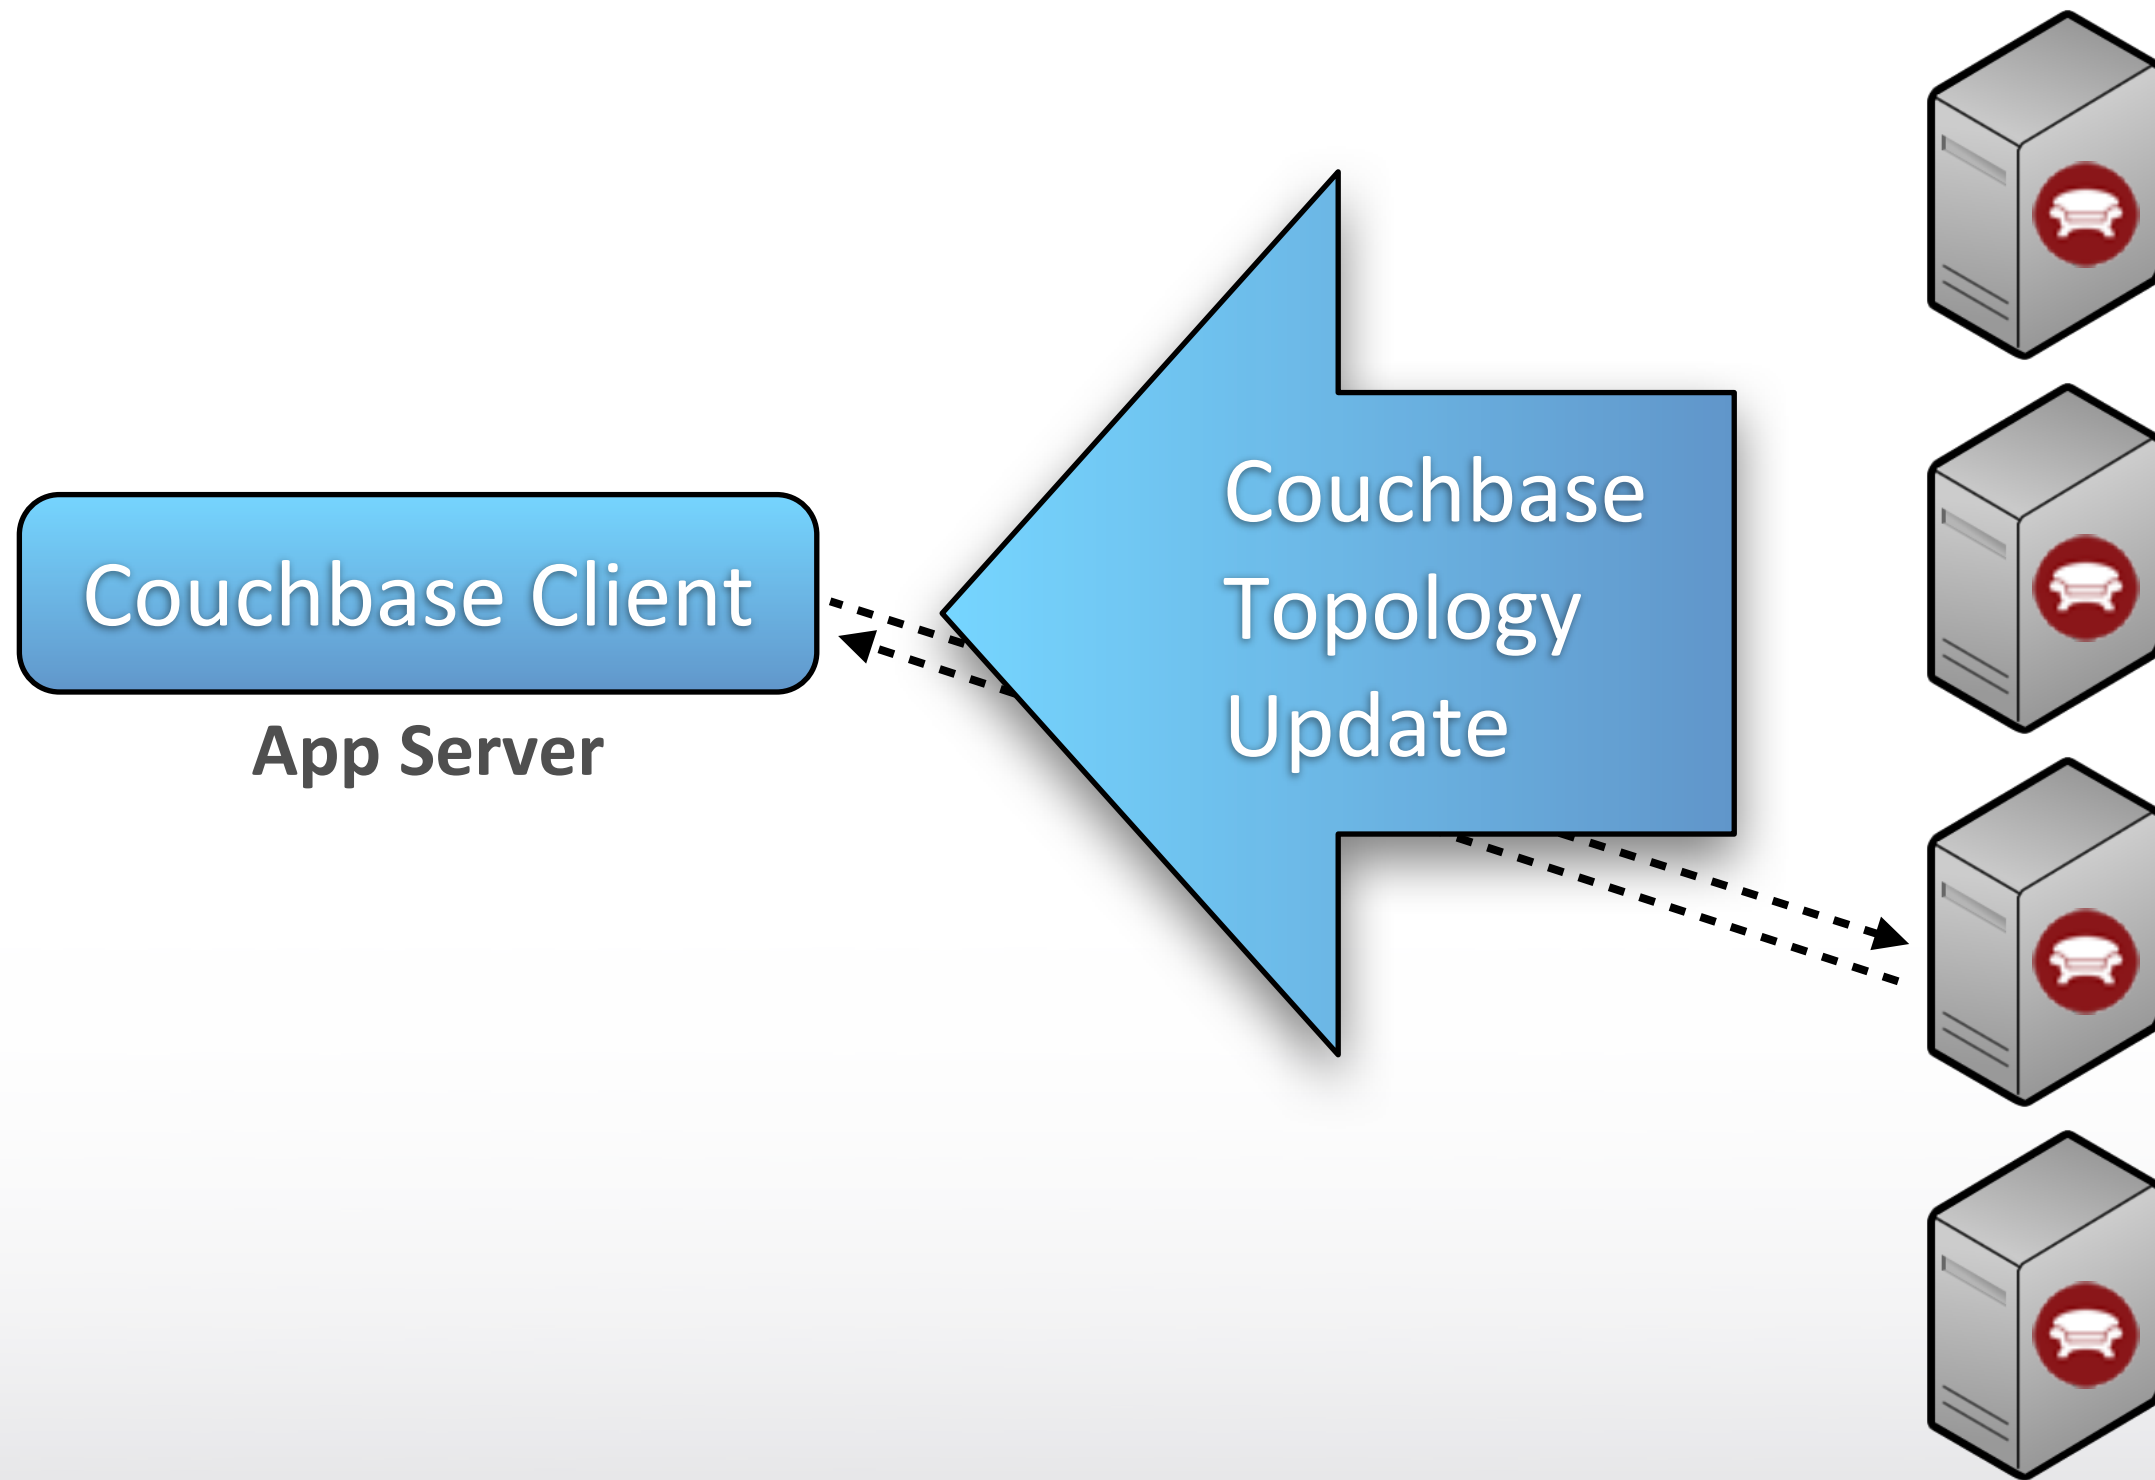
Couchbase

Client SDKs



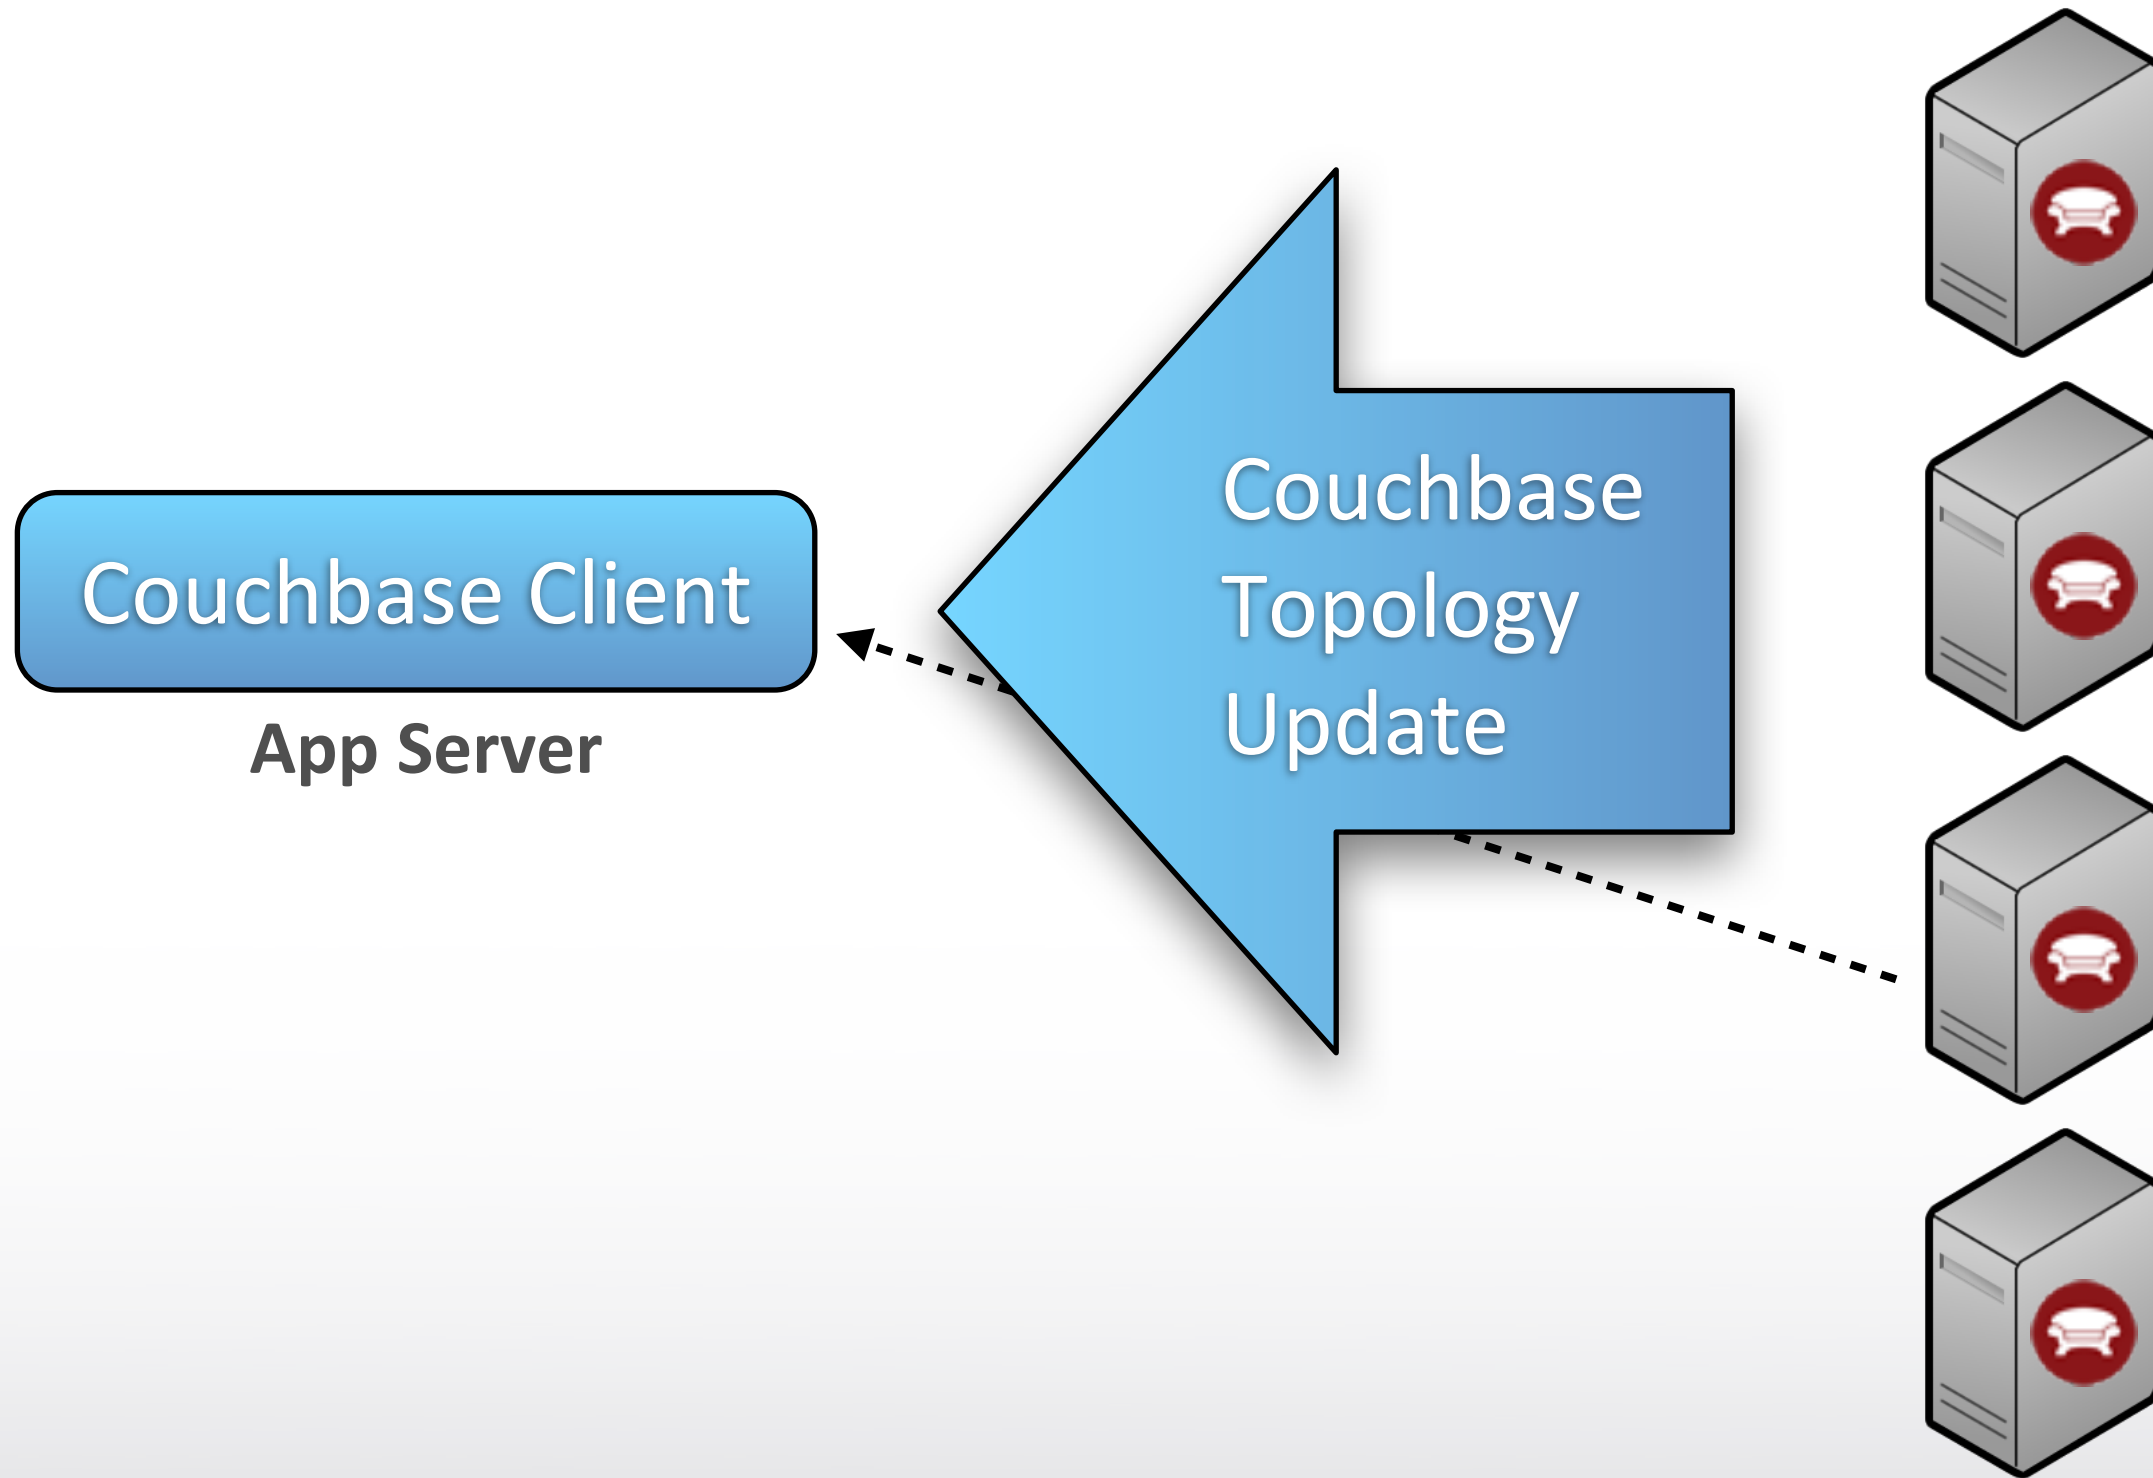
Couchbase

Client SDKs



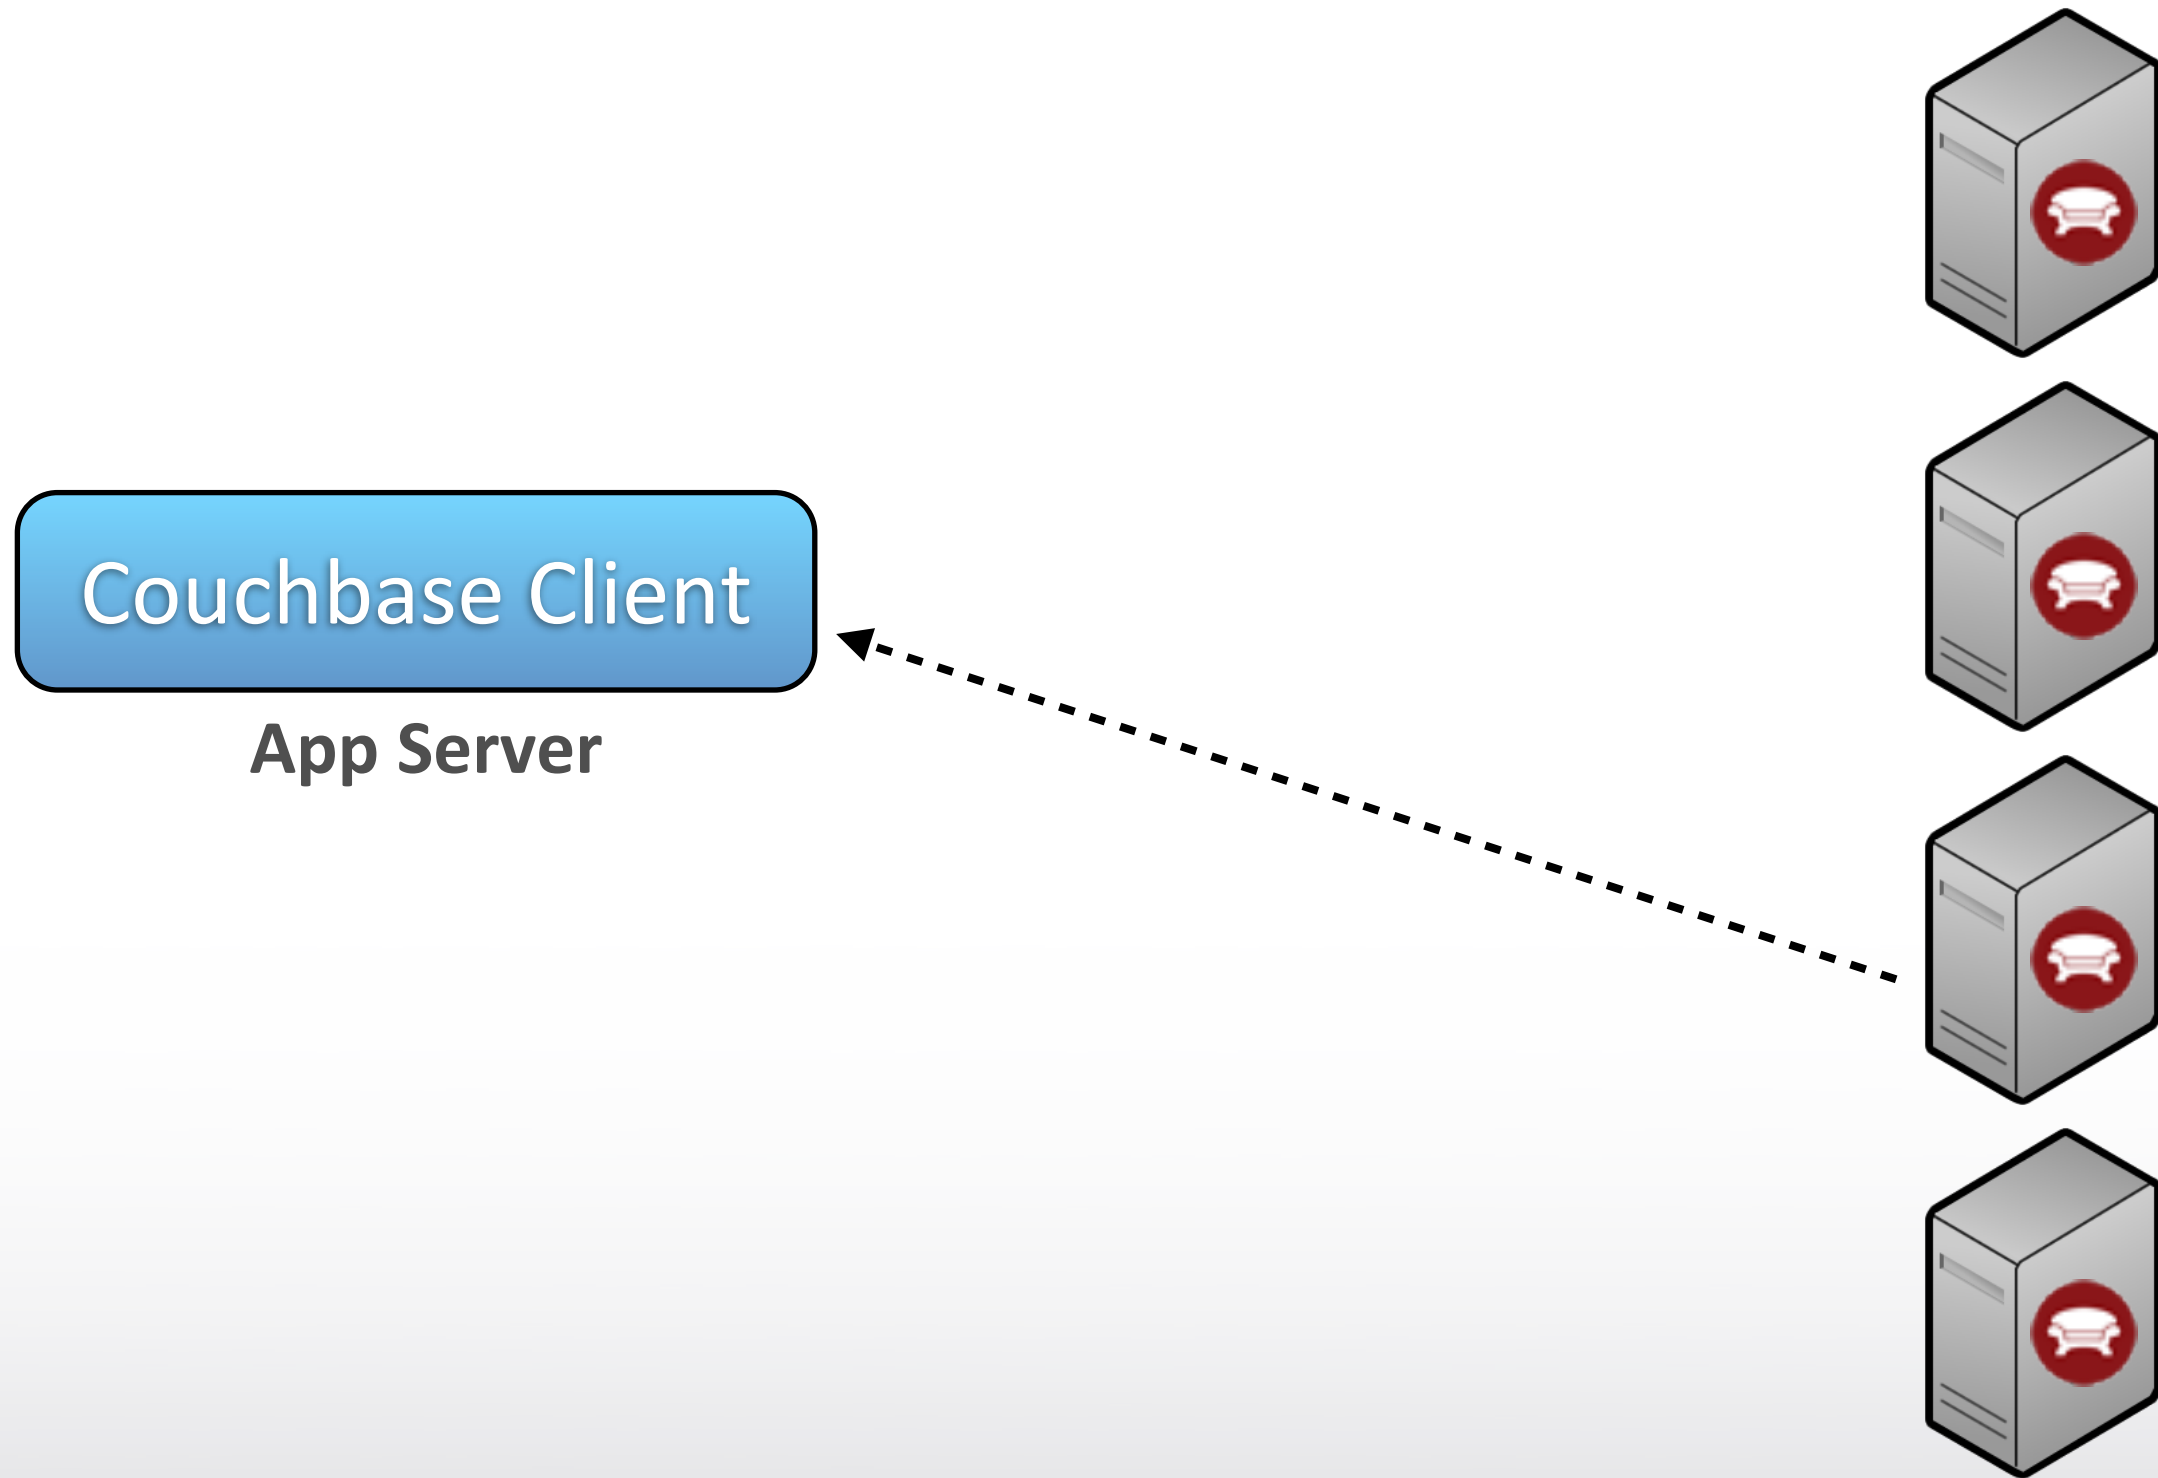
Couchbase

Client SDKs



Couchbase

Client SDKs




Couchbase

Client SDKs

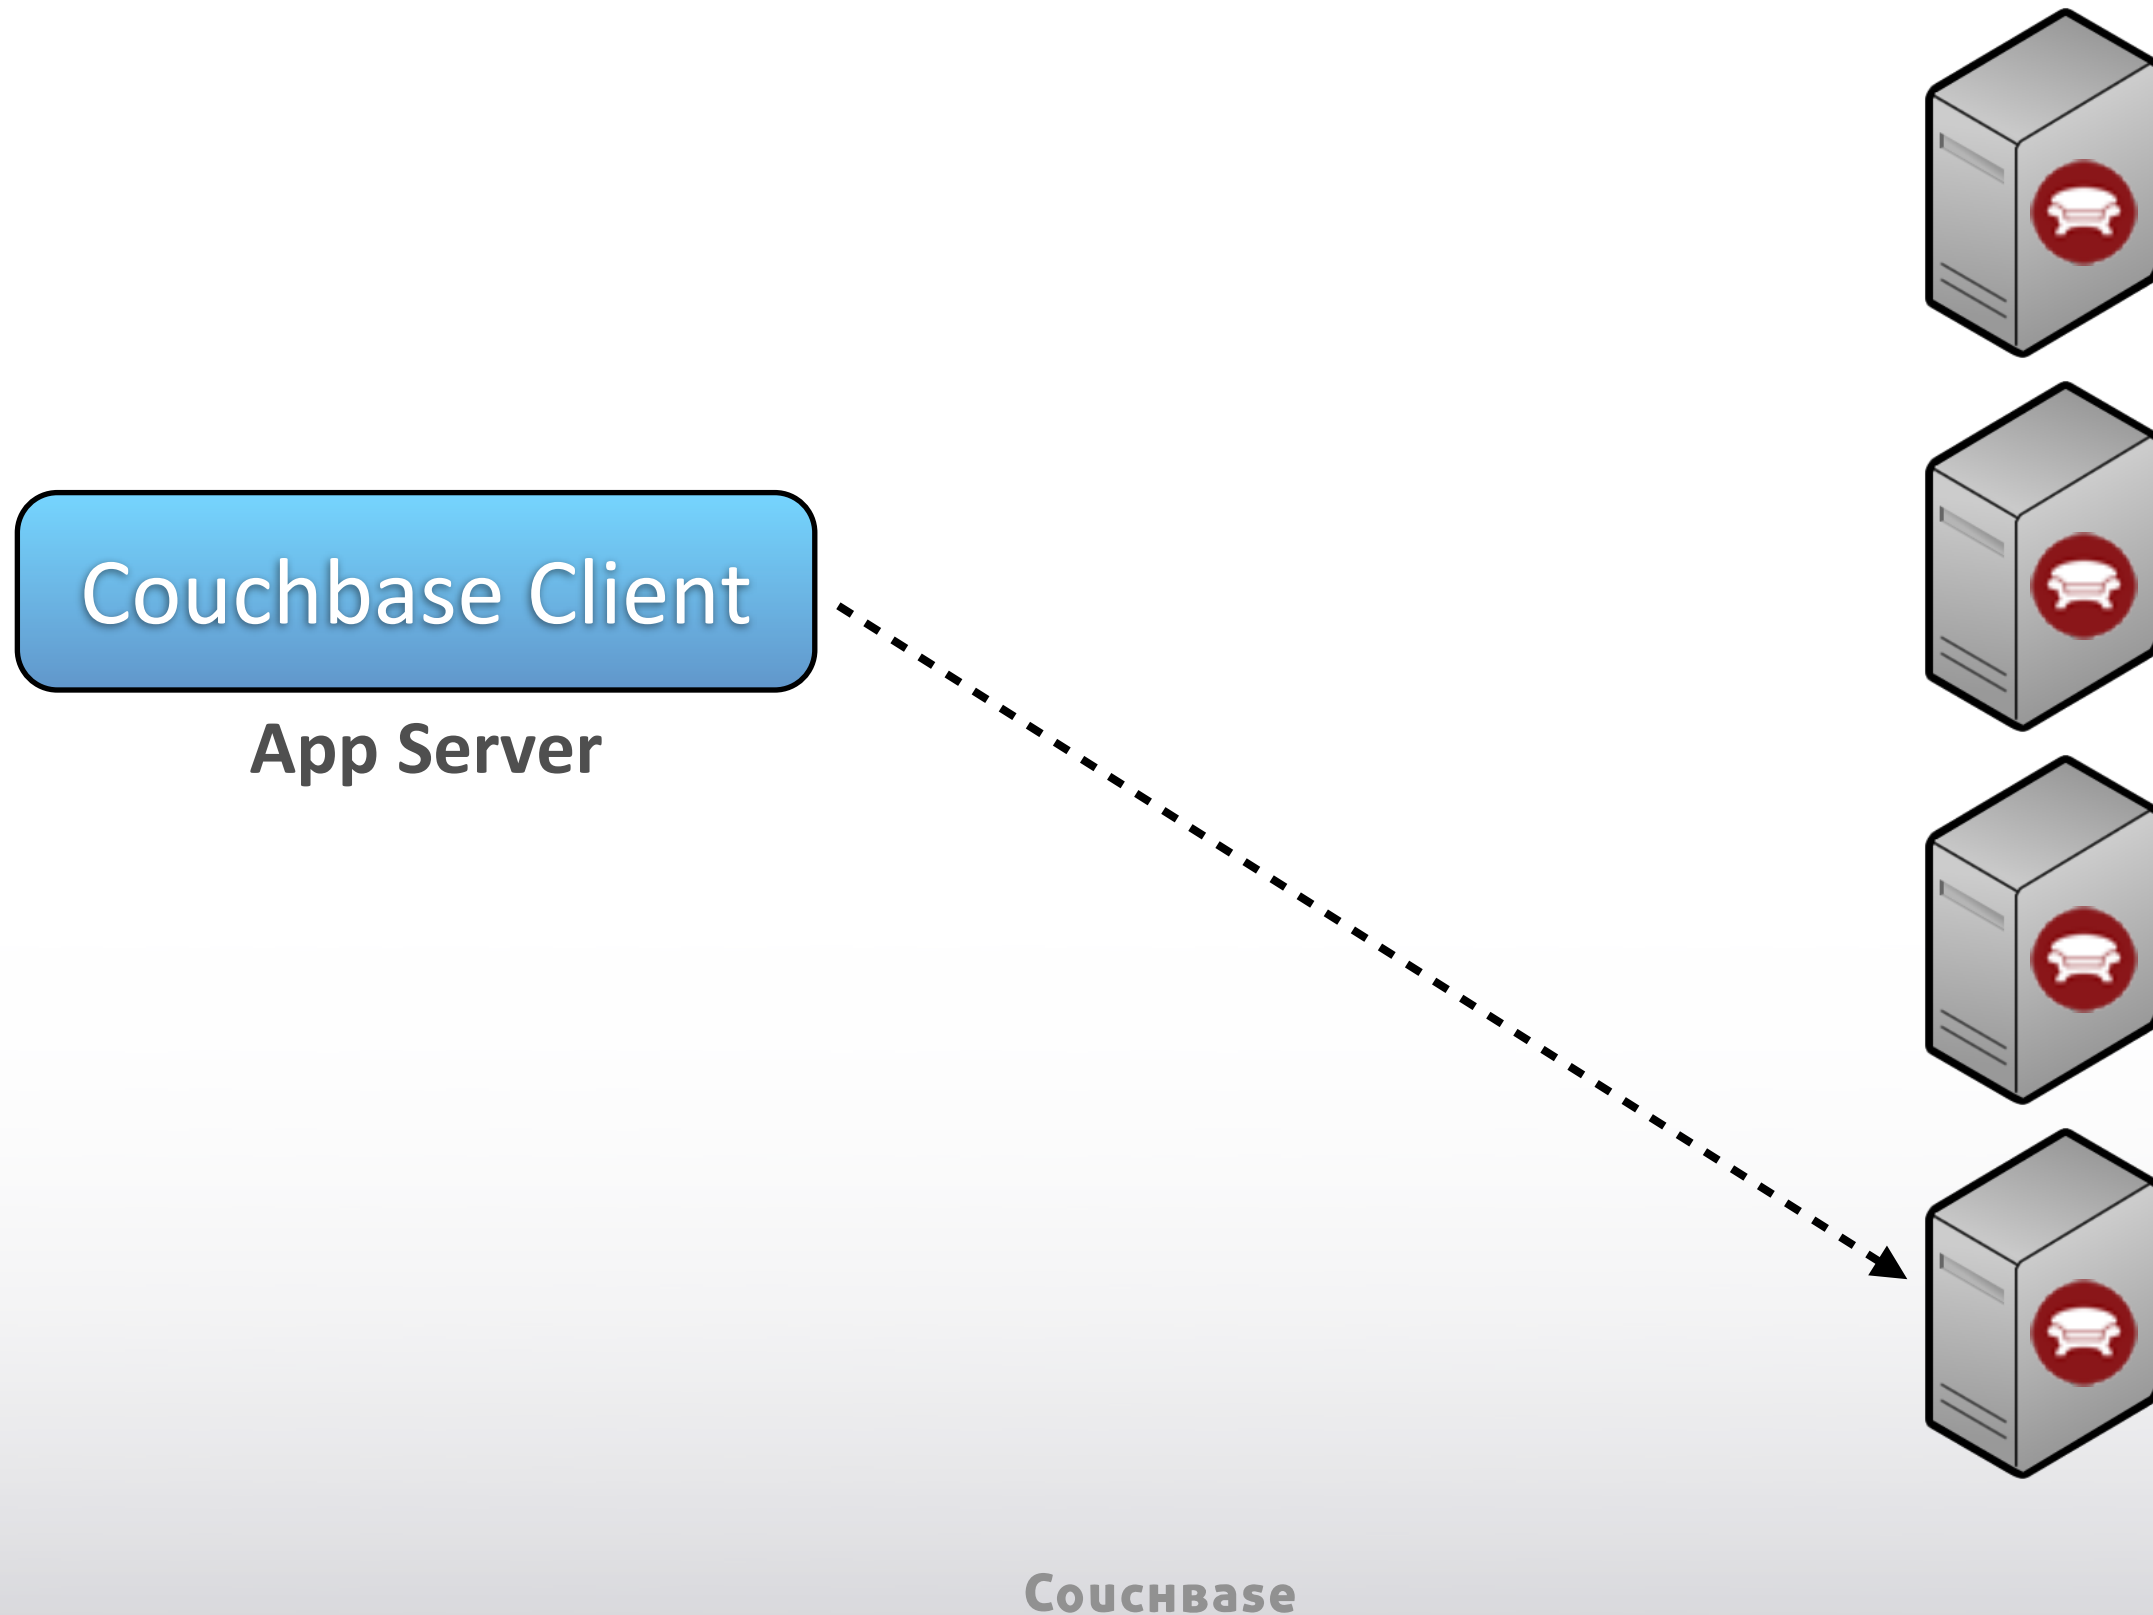
Couchbase Client

App Server

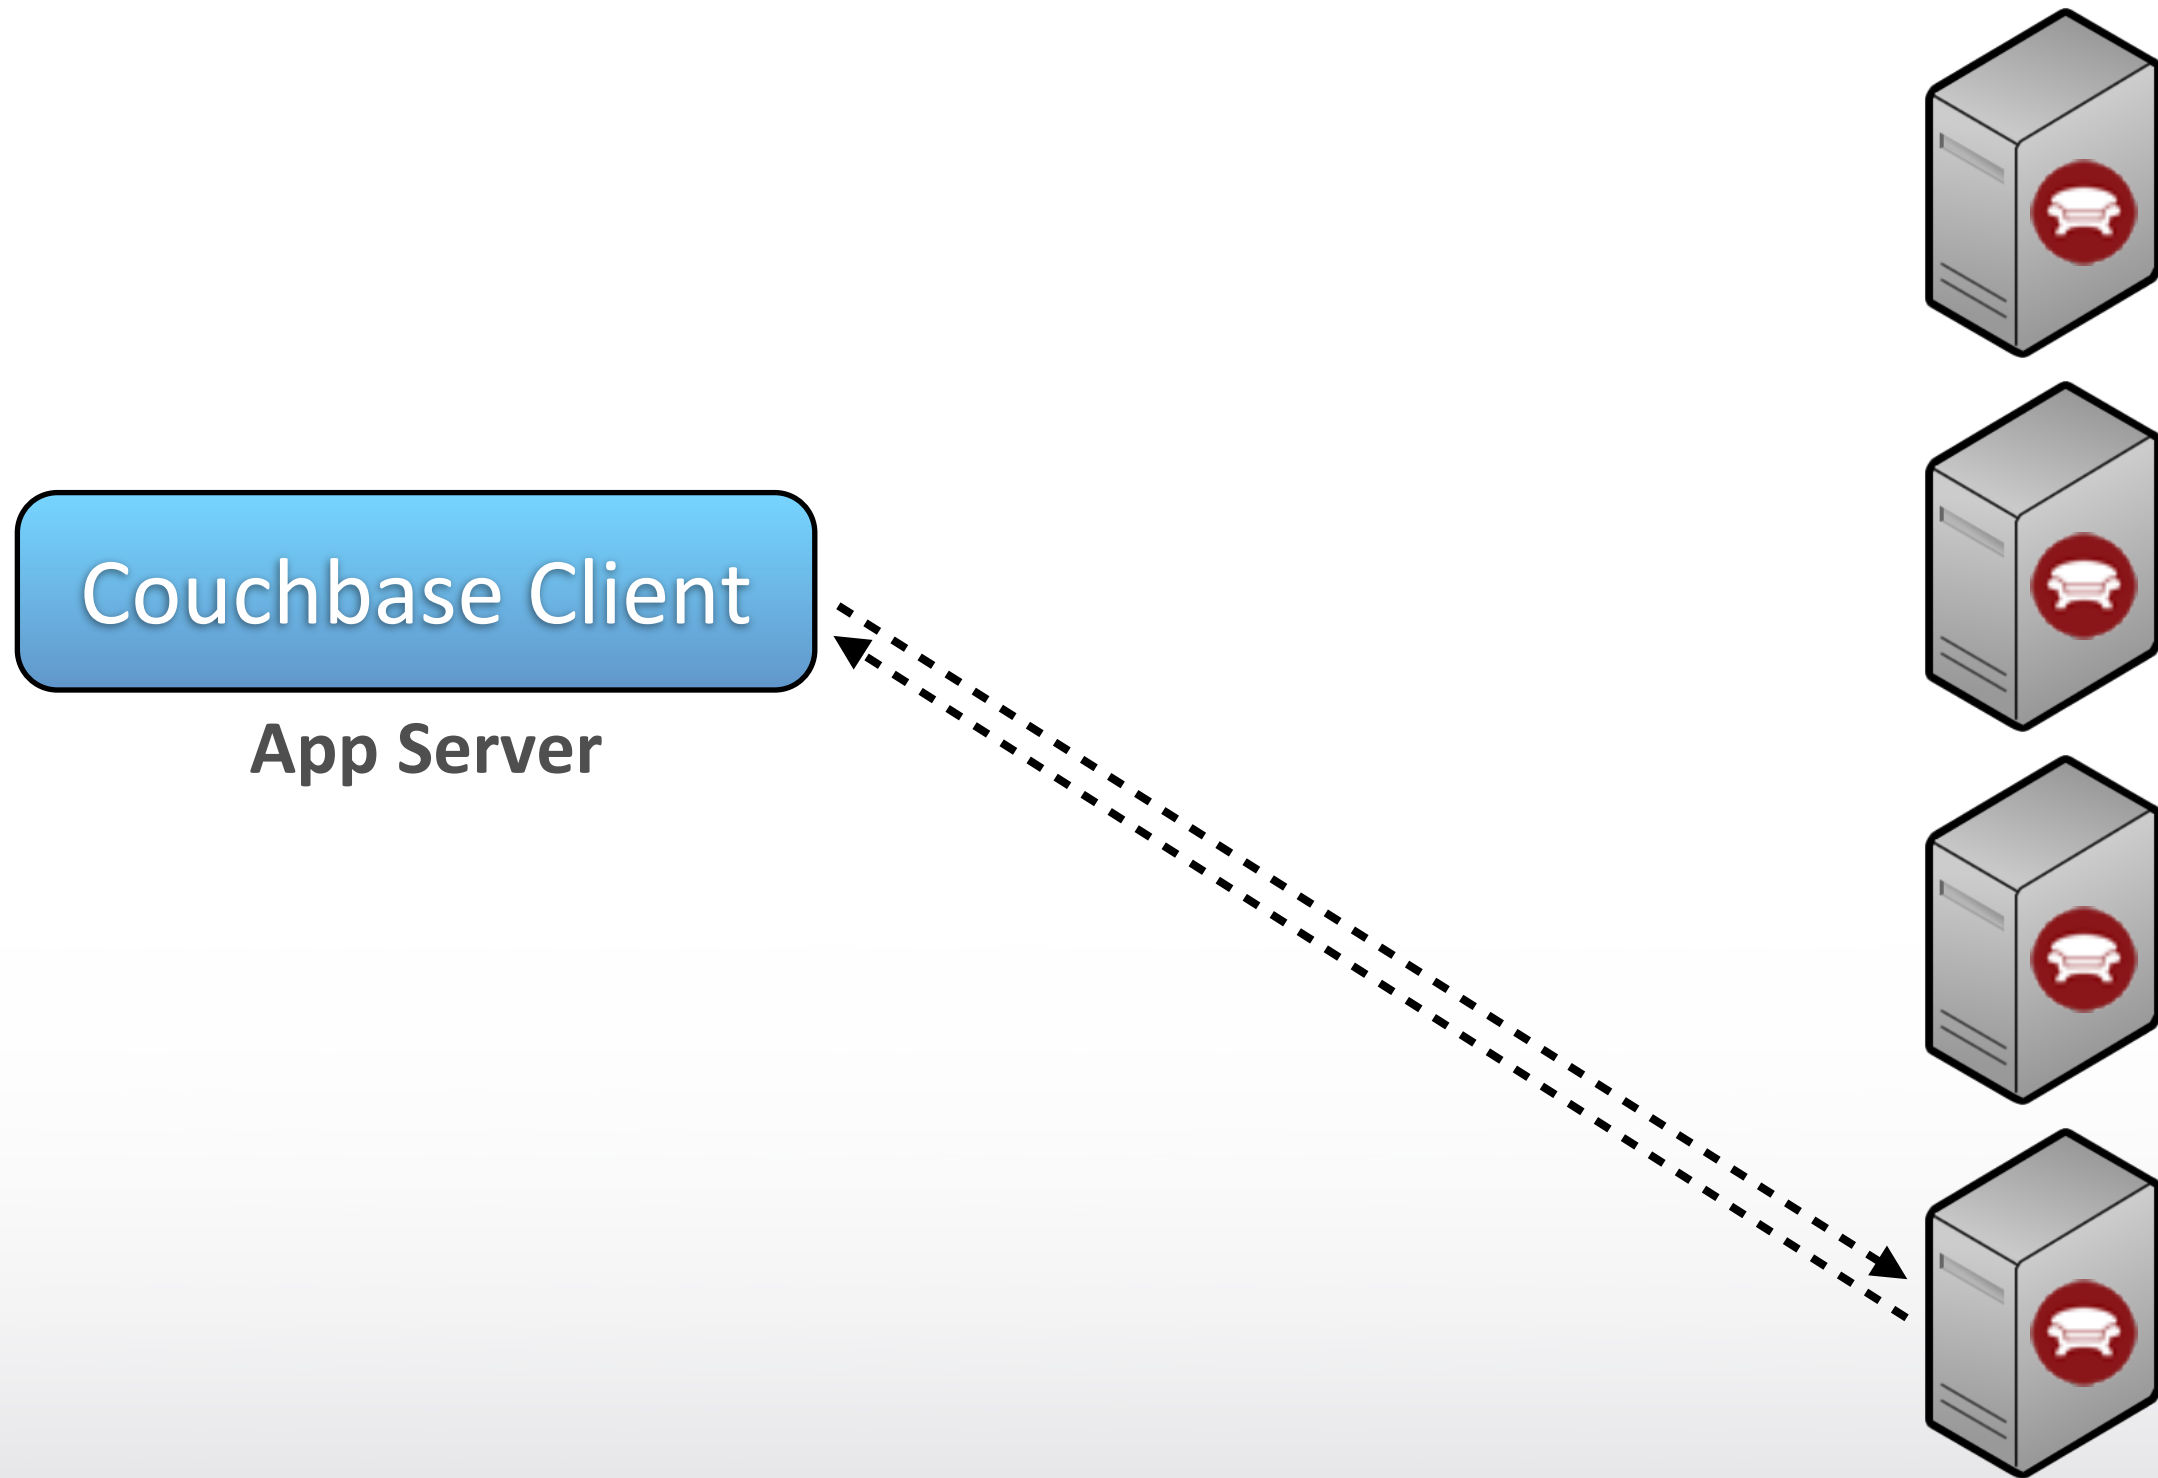


Couchbase

Client SDKs

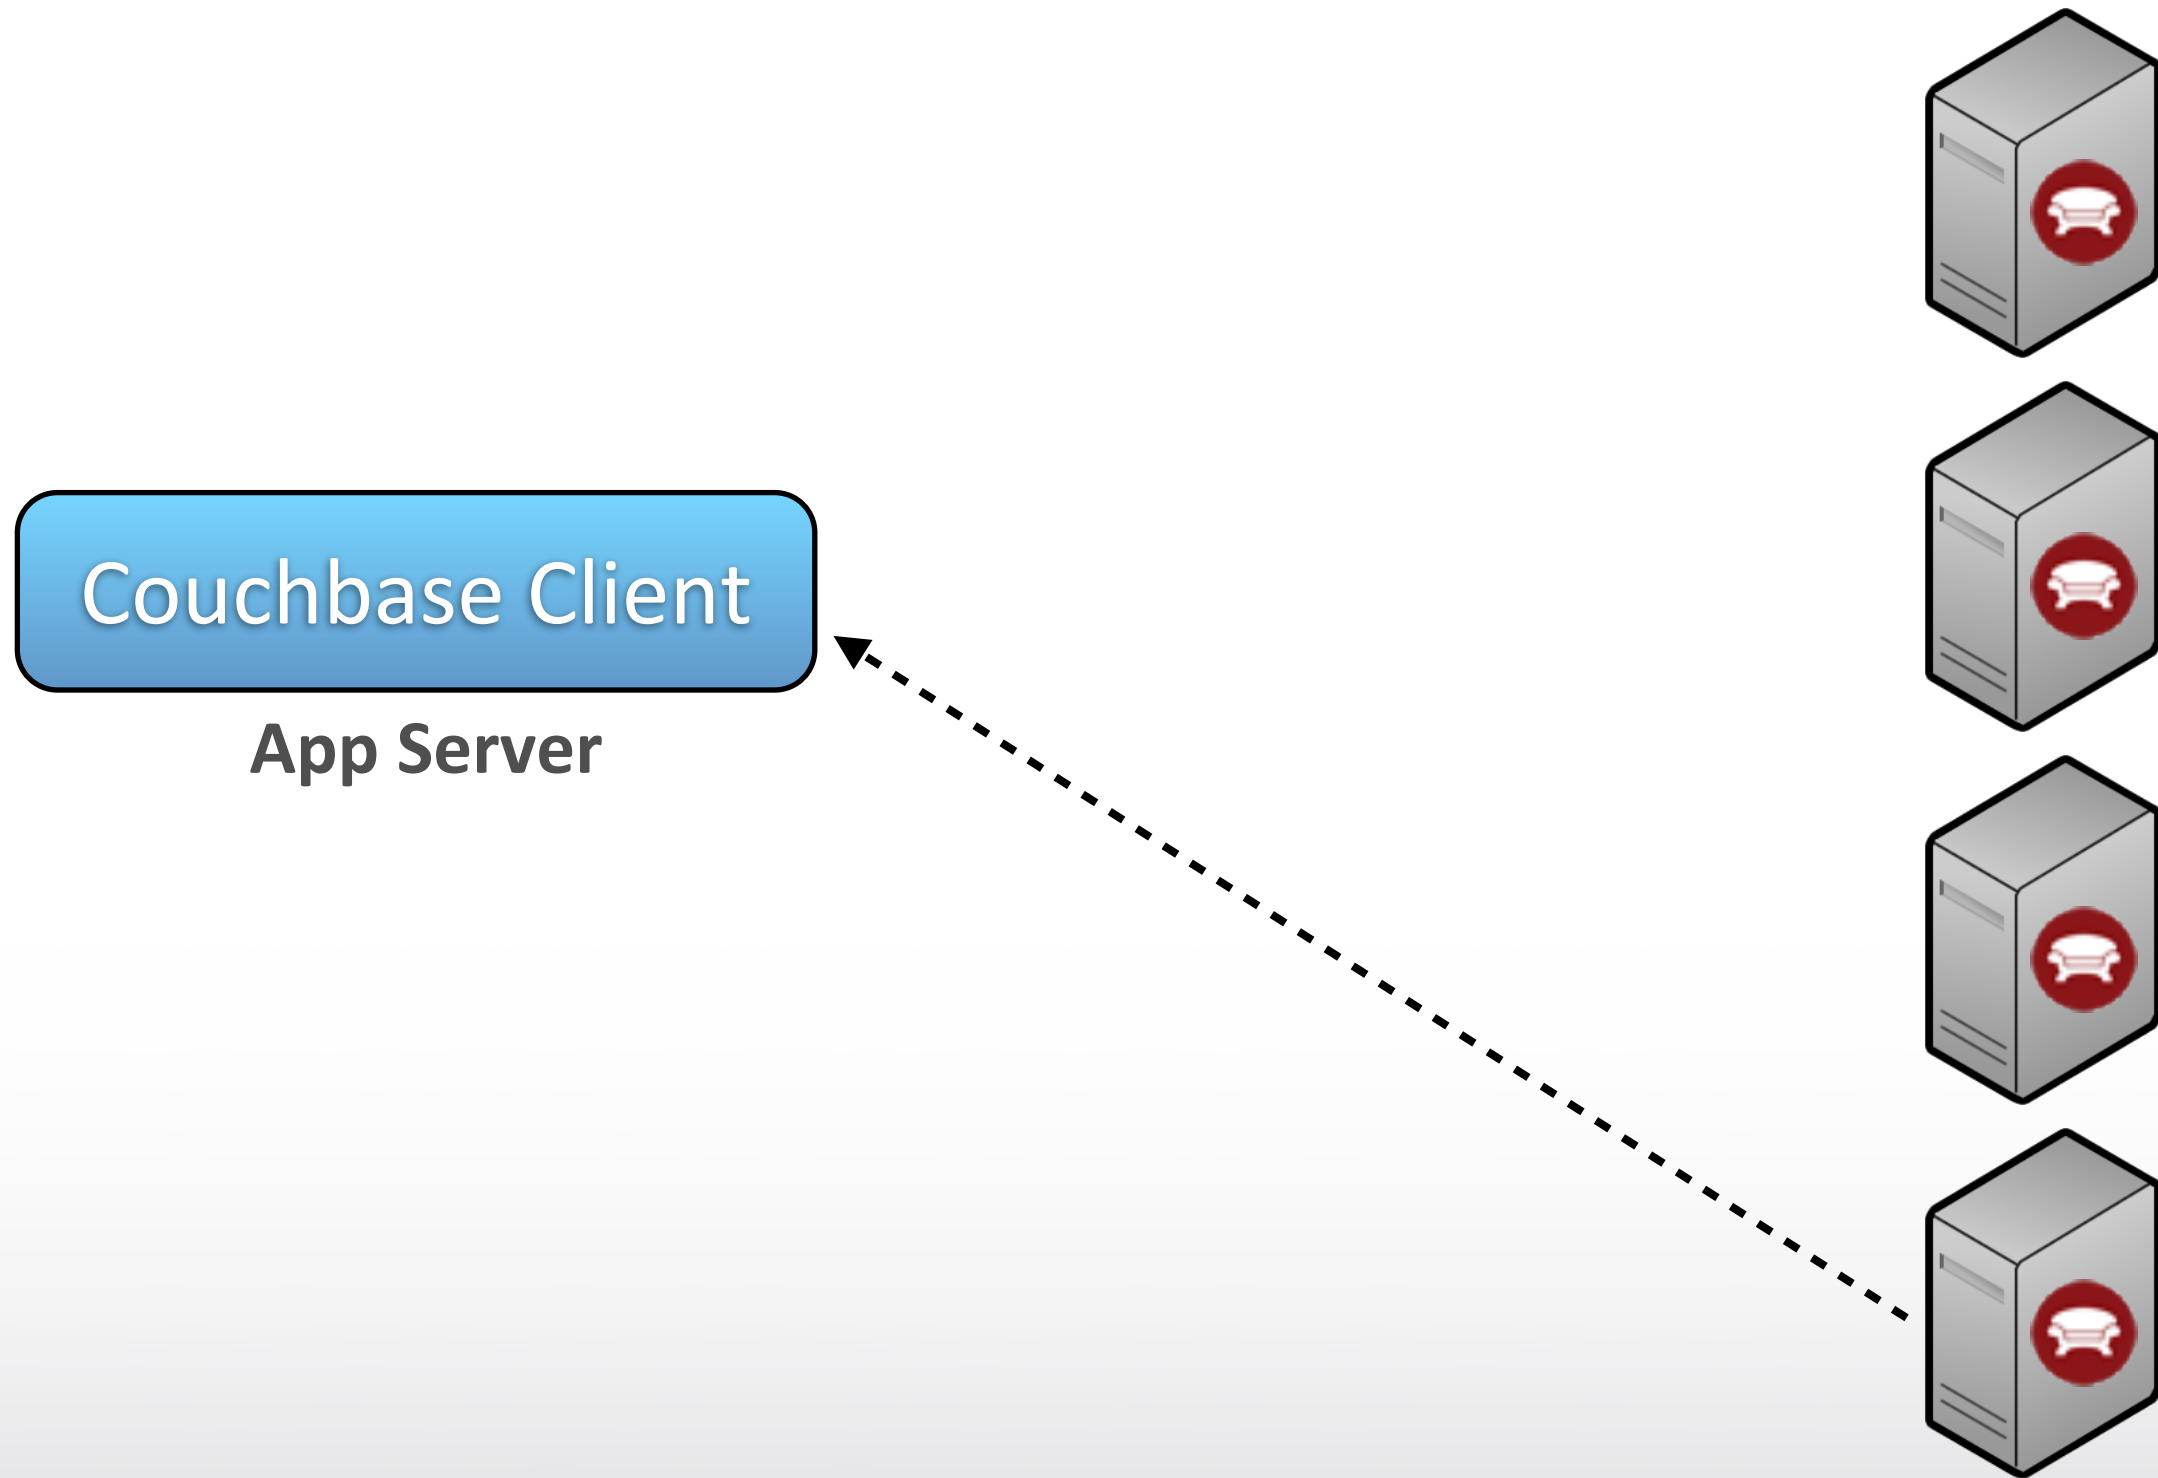


Client SDKs



Couchbase

Client SDKs




Couchbase

Client SDKs

Couchbase Client

App Server



Couchbase

Couchbase

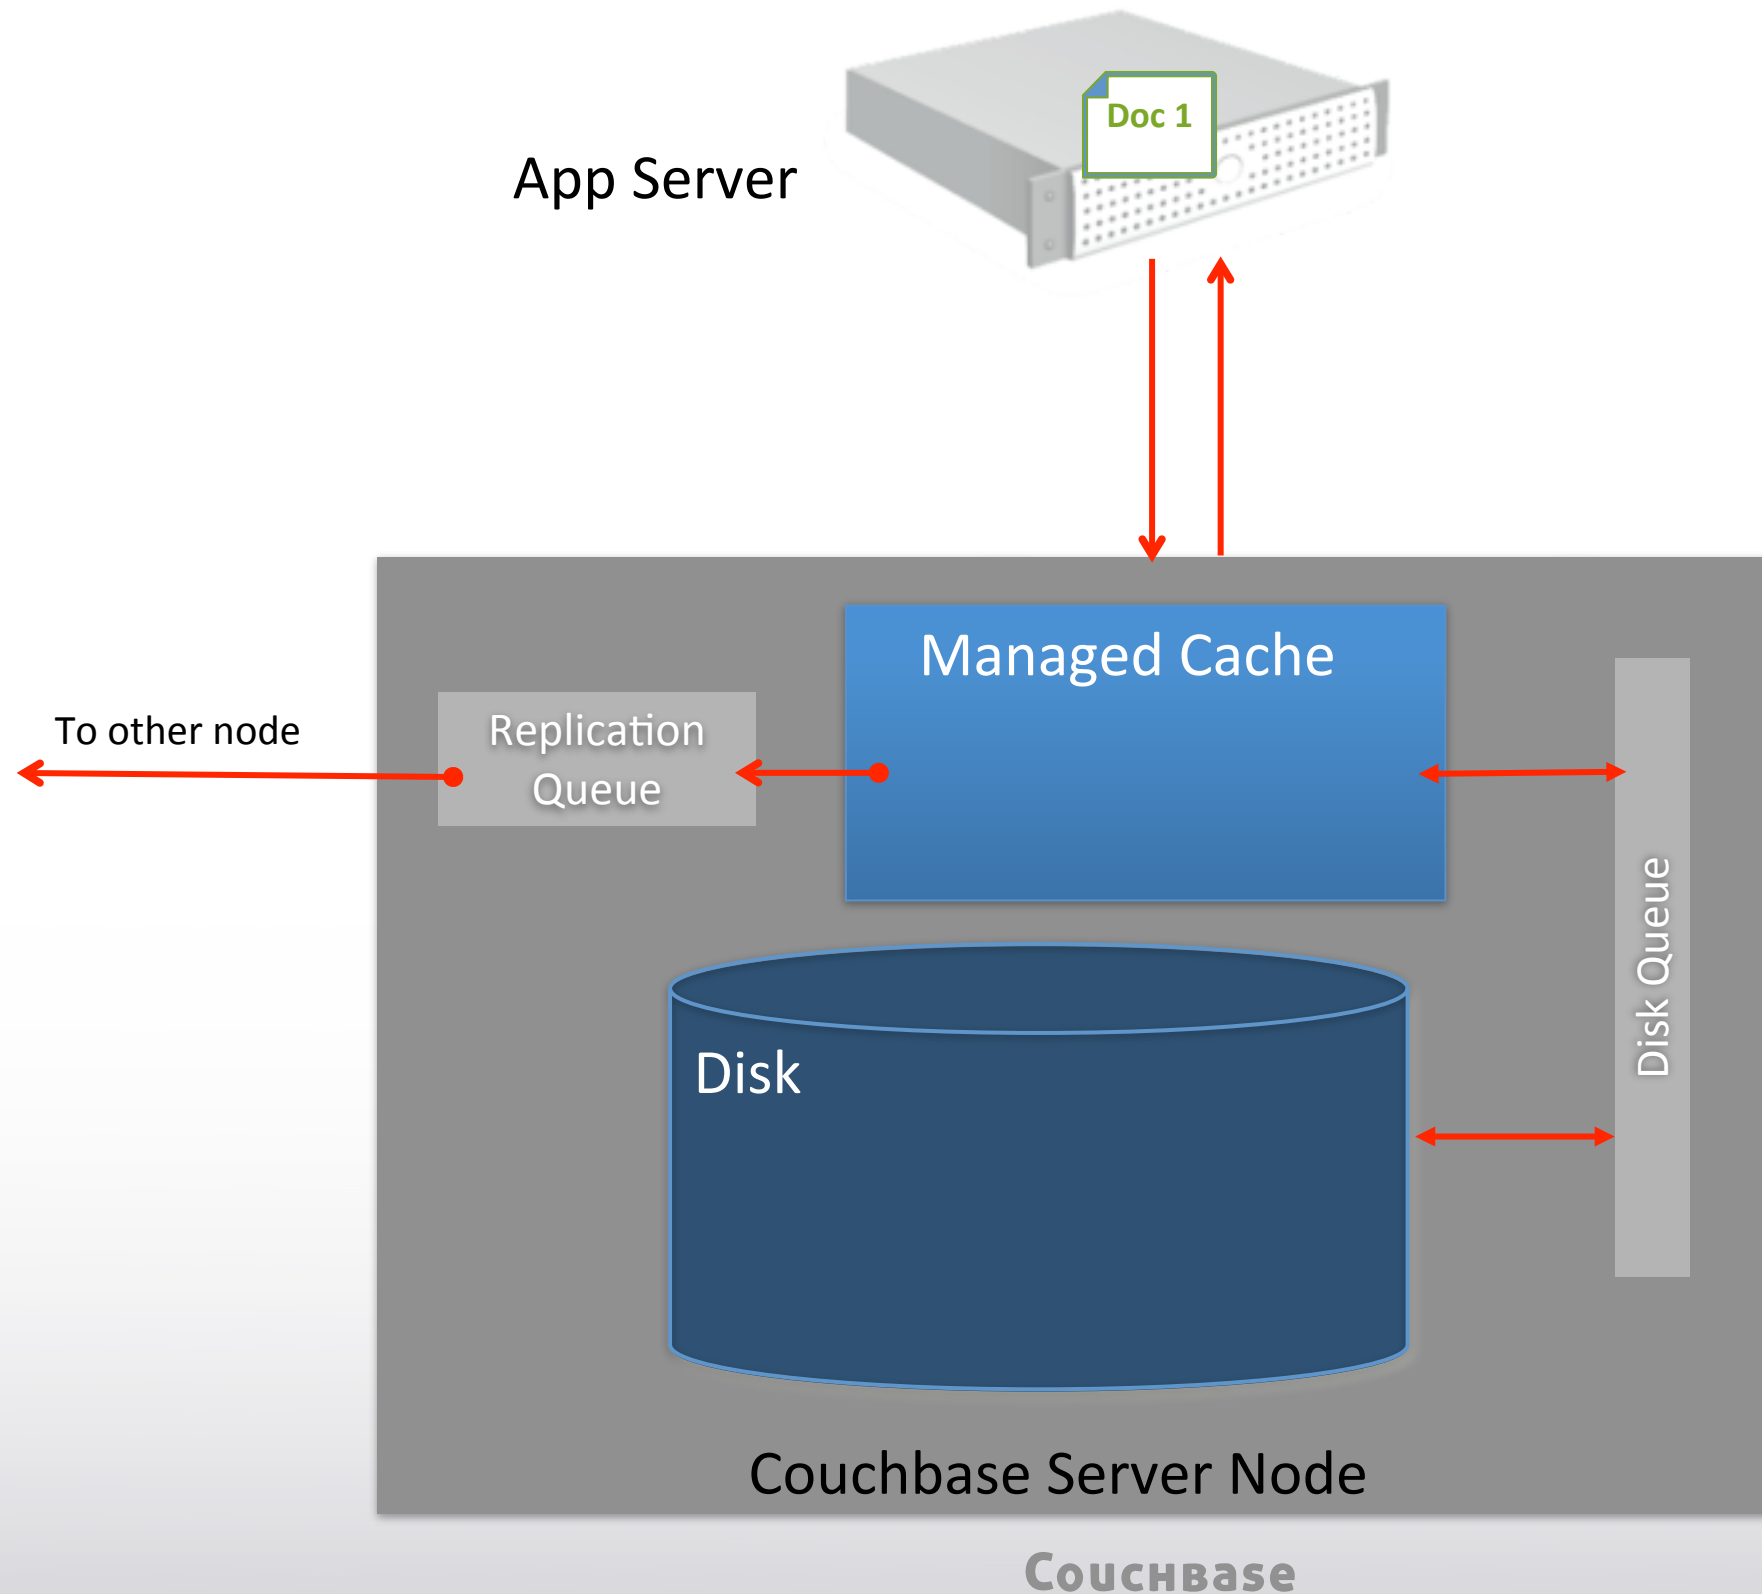


Couchbase

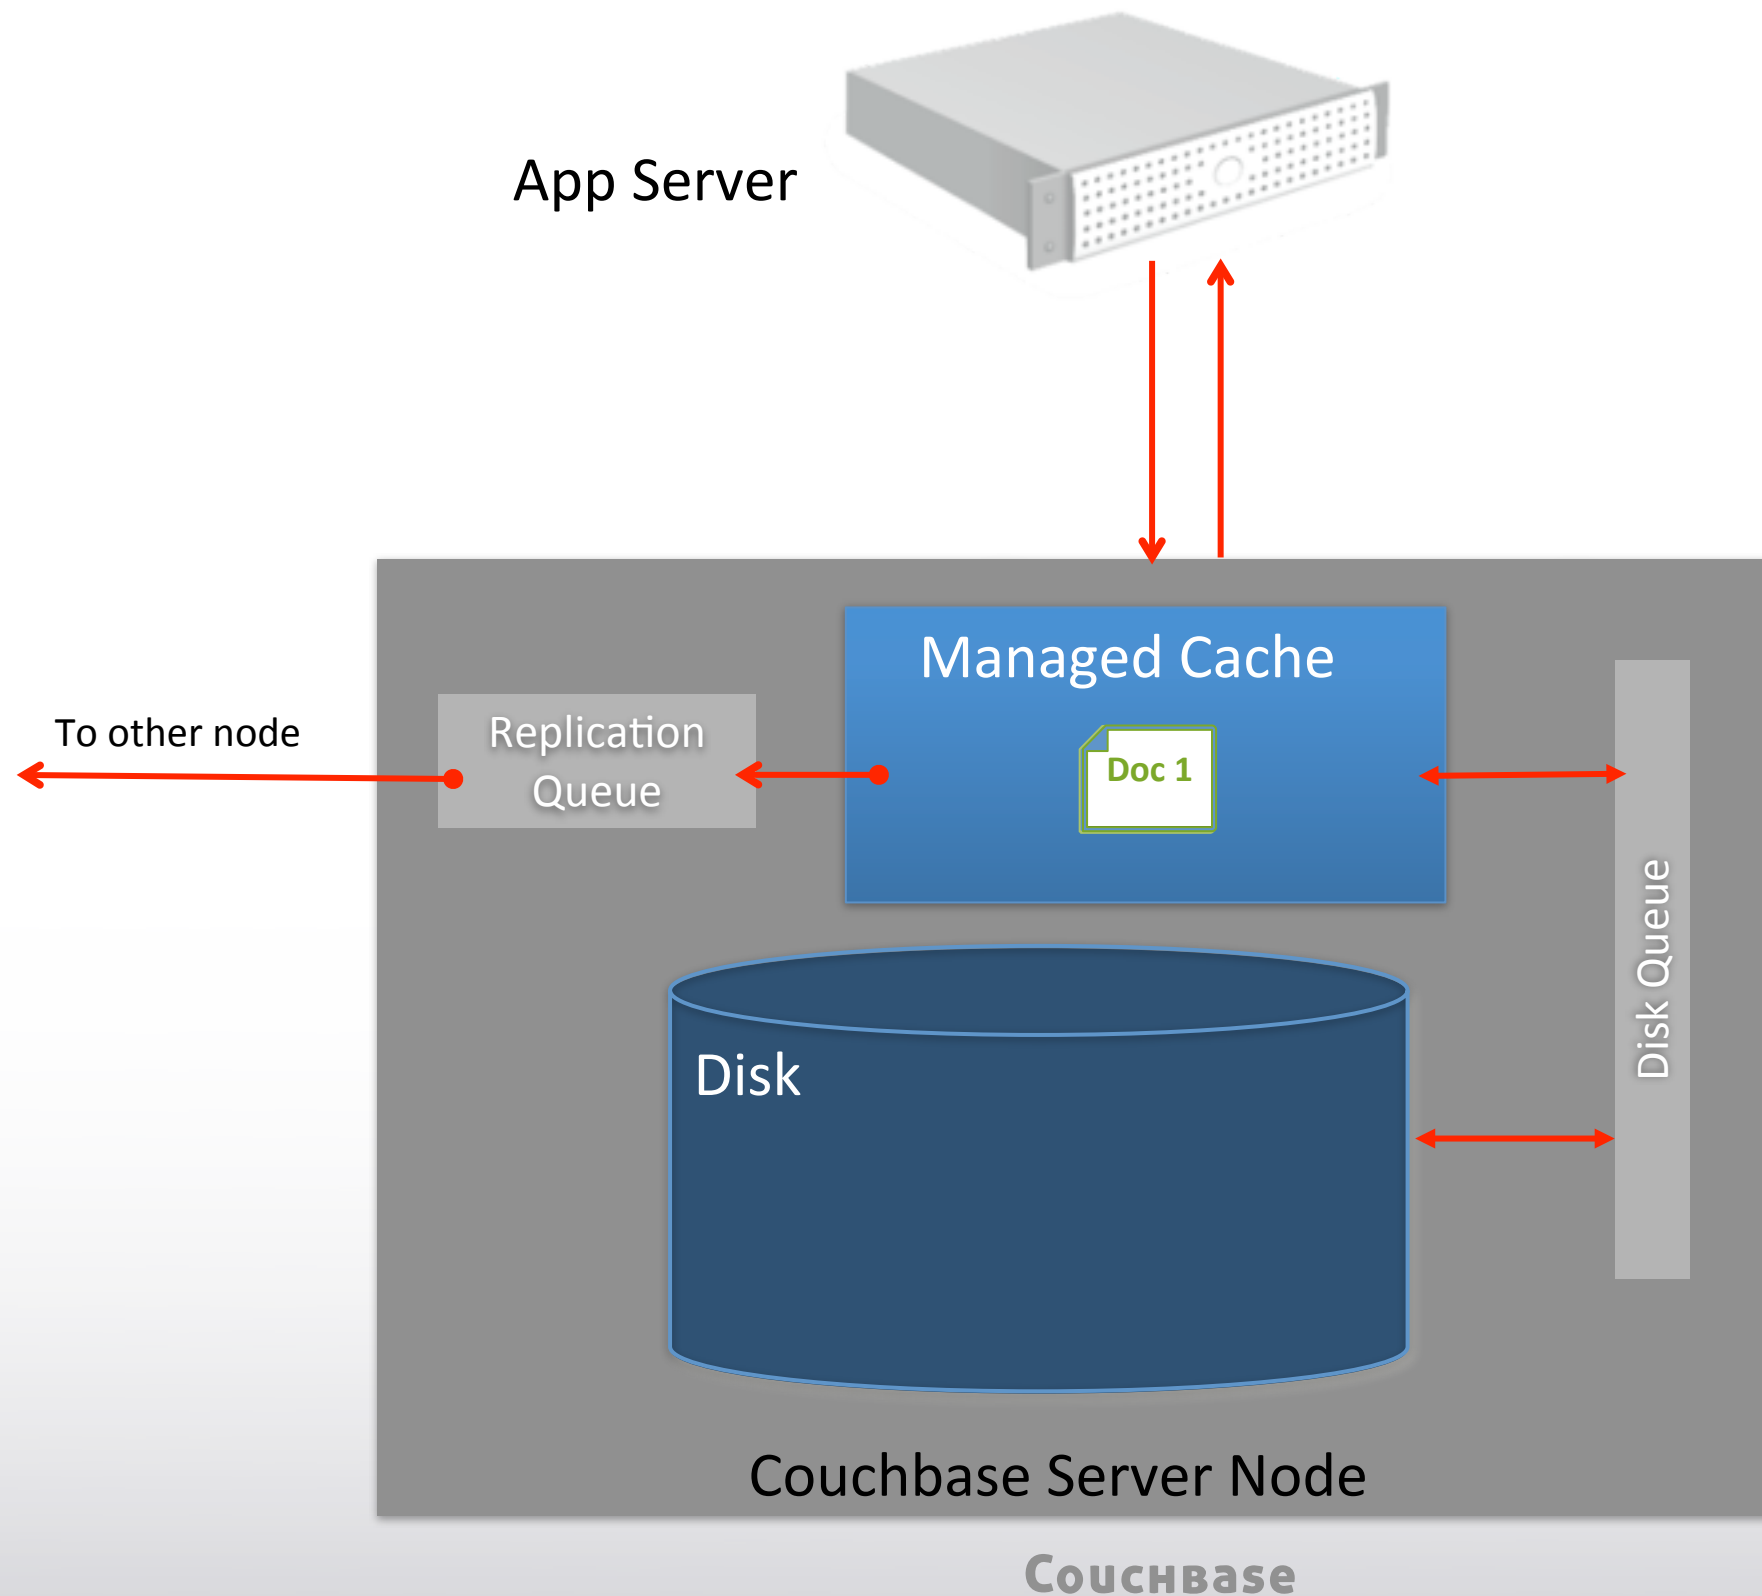
COUCHBASE OPERATIONS

Couchbase

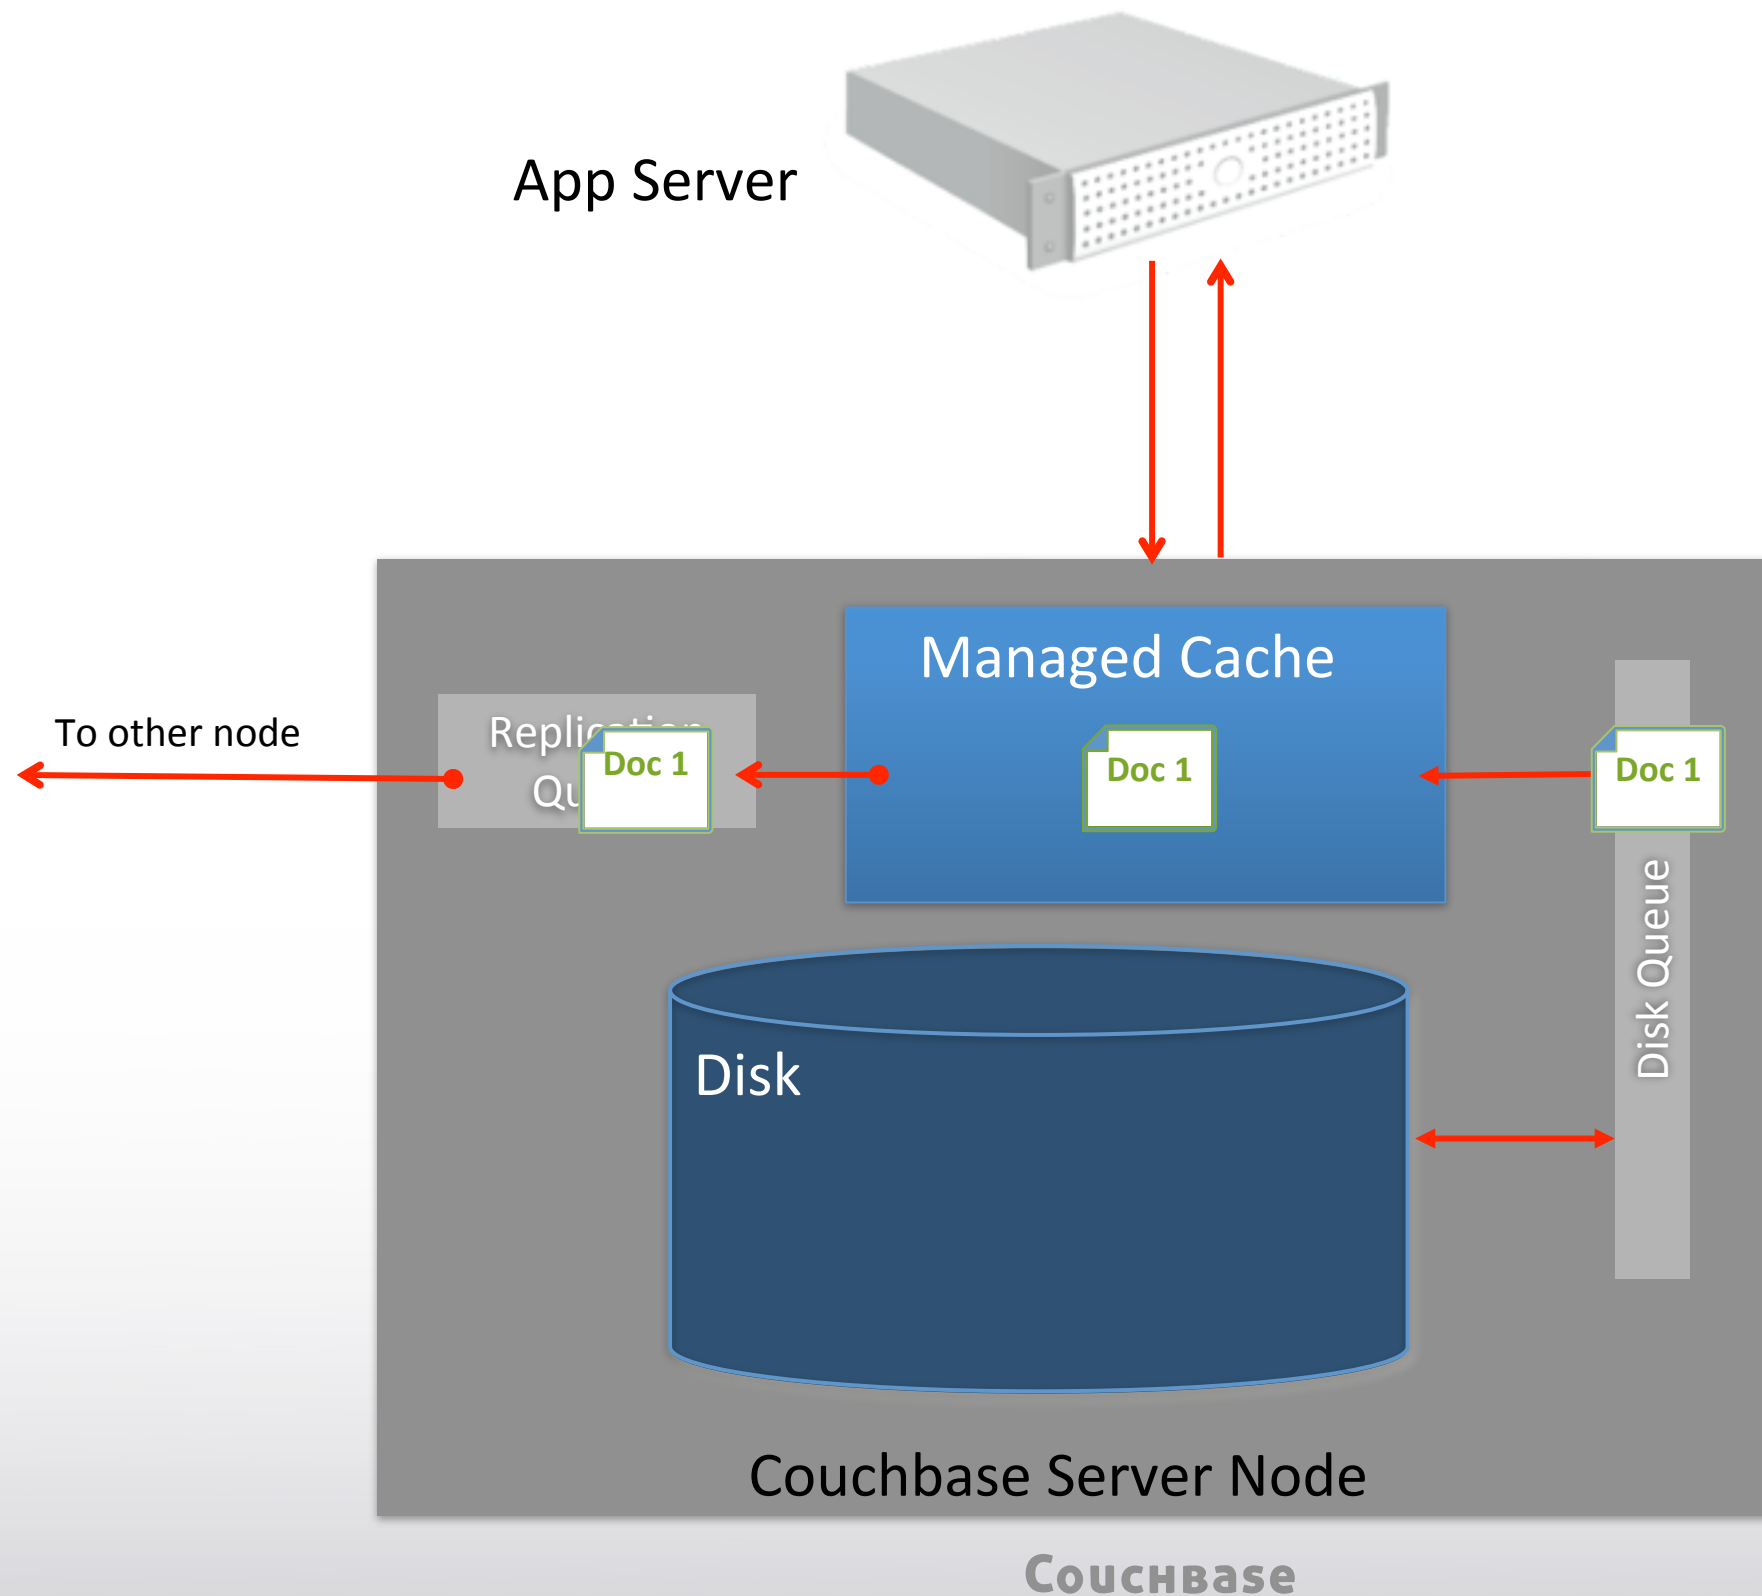
Write Operation



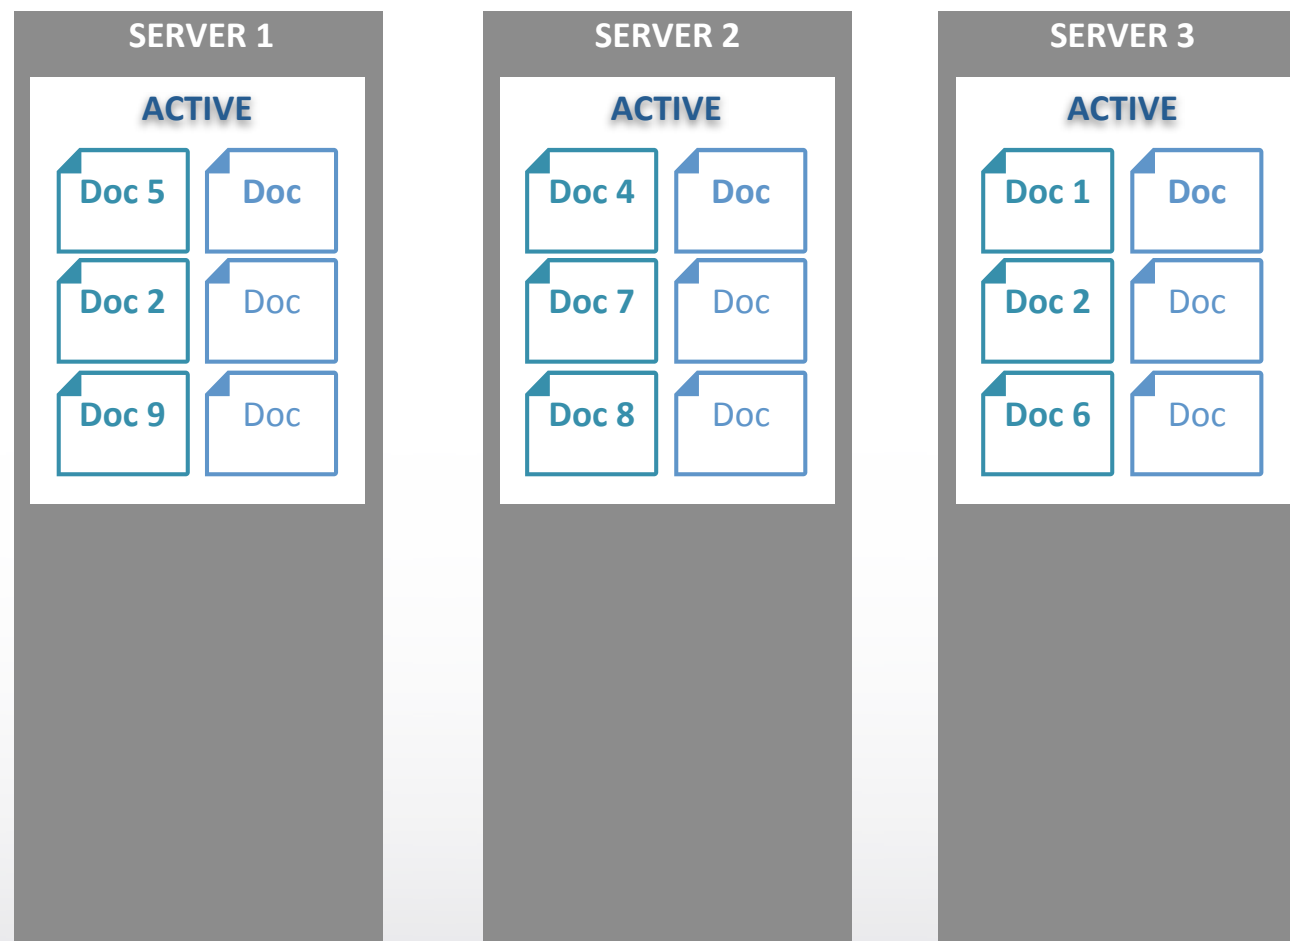
Write Operation



Write Operation



Basic Operations

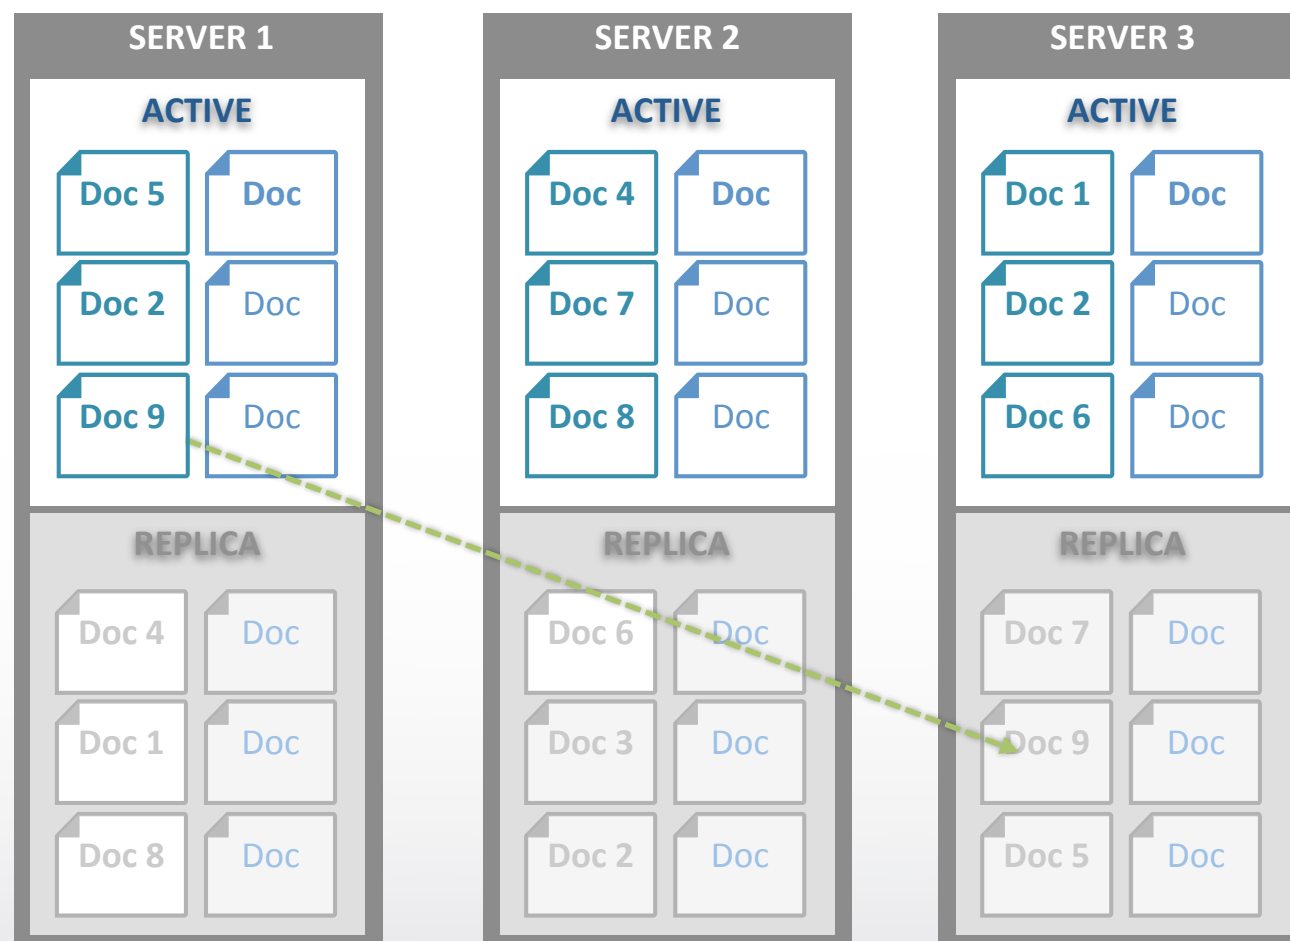


COUCHBASE SERVER CLUSTER

Couchbase

- Docs distributed evenly across servers
- Each server stores both active and replica docs
 - Only one doc 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

Basic Operations


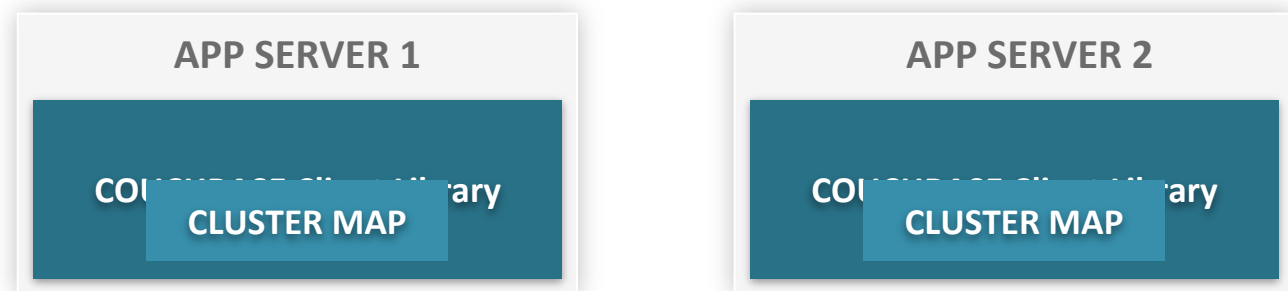


COUCHBASE SERVER CLUSTER

Couchbase

- Docs distributed evenly across servers
- Each server stores both active and replica docs
 - Only one doc 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

Basic Operations

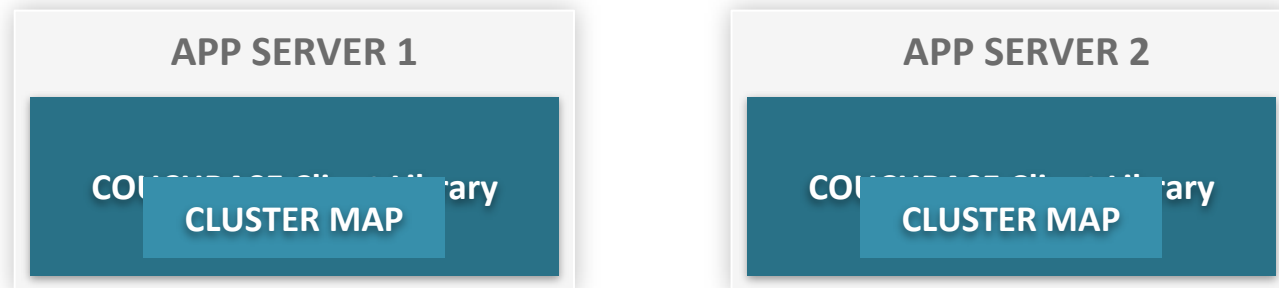


COUCHBASE SERVER CLUSTER


Couchbase

- Docs distributed evenly across servers
- Each server stores both active and replica docs
 - Only one doc 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

Basic Operations



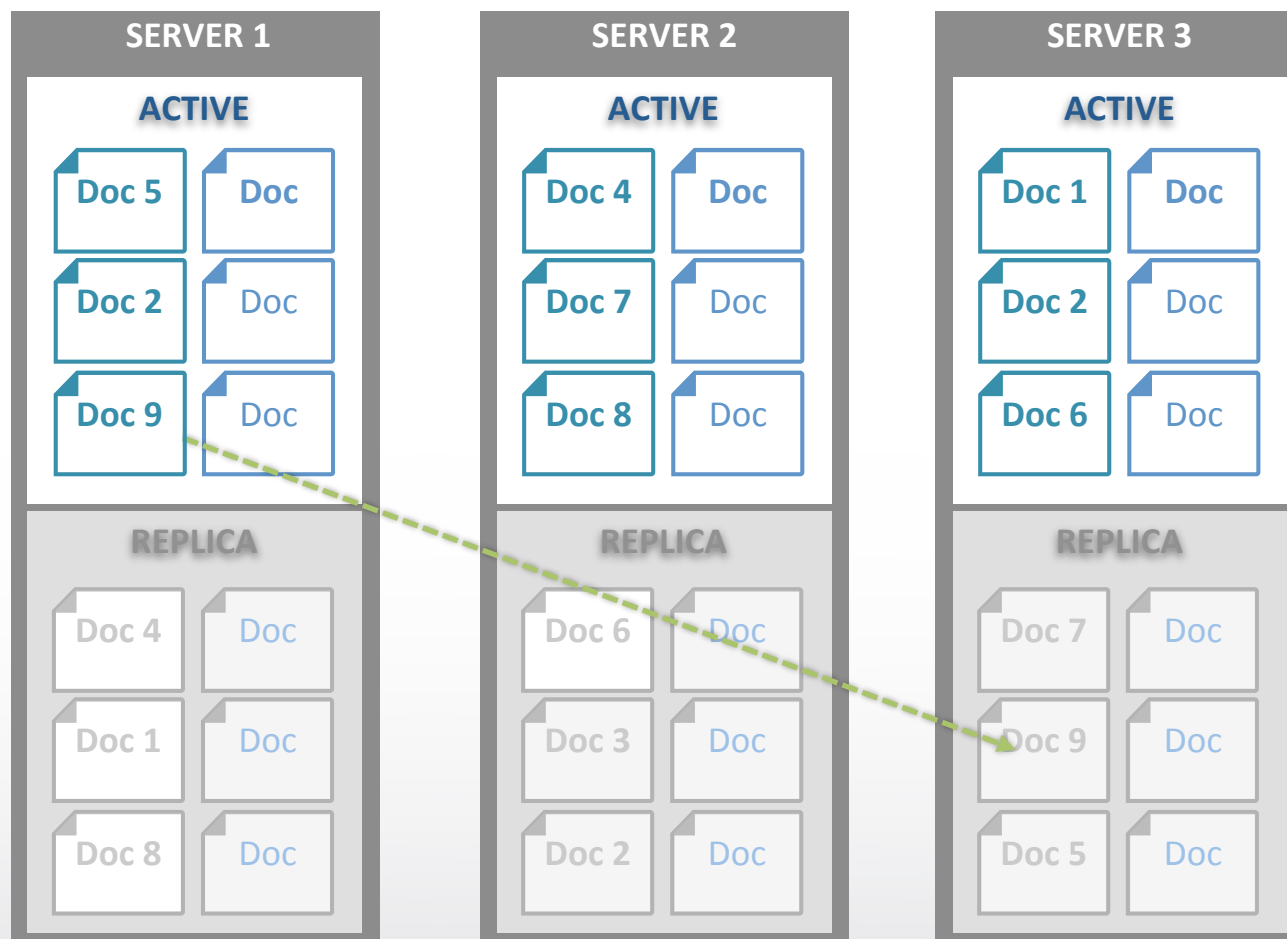
- Docs distributed evenly across servers
- Each server stores both active and replica docs



COUCHBASE SERVER CLUSTER

Couchbase

Basic Operations


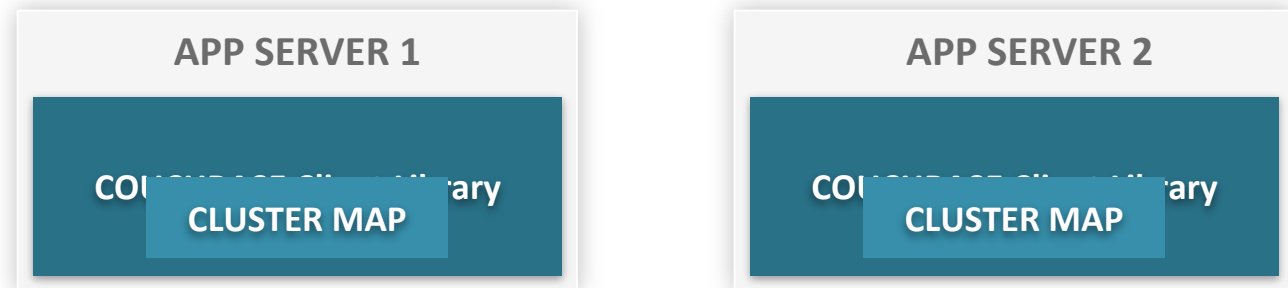


COUCHBASE SERVER CLUSTER

Couchbase

- Docs distributed evenly across servers
- Each server stores both active and replica docs
 - Only one doc active at a time

Basic Operations




COUCHBASE SERVER CLUSTER

Couchbase

- Docs distributed evenly across servers
- Each server stores both active and replica docs
Only one doc active at a time
- Client library provides app with simple interface to database

Basic Operations




COUCHBASE SERVER CLUSTER

Couchbase

- Docs distributed evenly across servers
- Each server stores both active and replica docs
 - Only one doc active at a time
- Client library provides app with simple interface to database
- Cluster map provides map to which server doc is on

Basic Operations

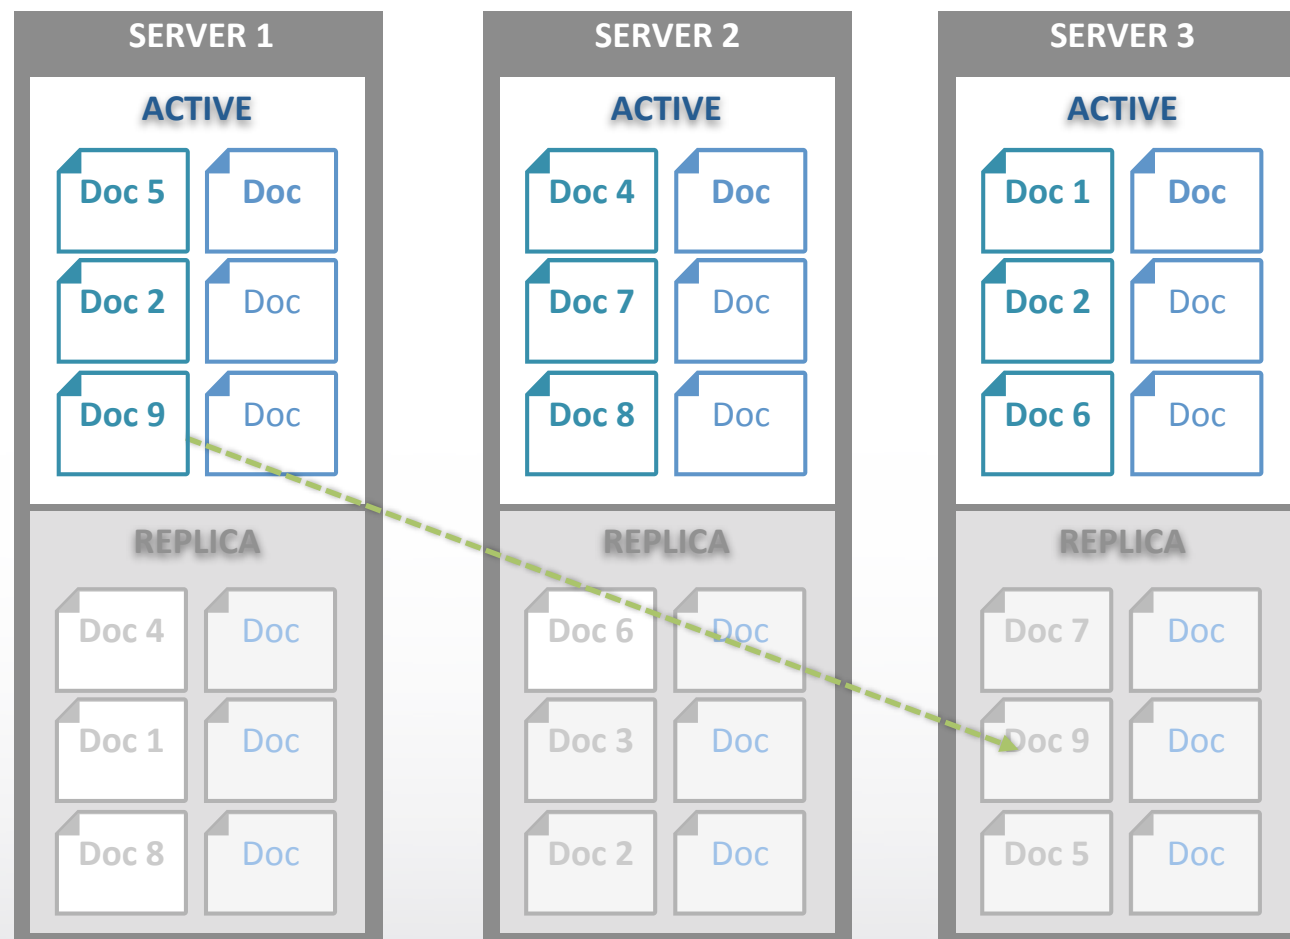



COUCHBASE SERVER CLUSTER

Couchbase

- Docs distributed evenly across servers
- Each server stores both active and replica docs
 - Only one doc 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

Basic Operations

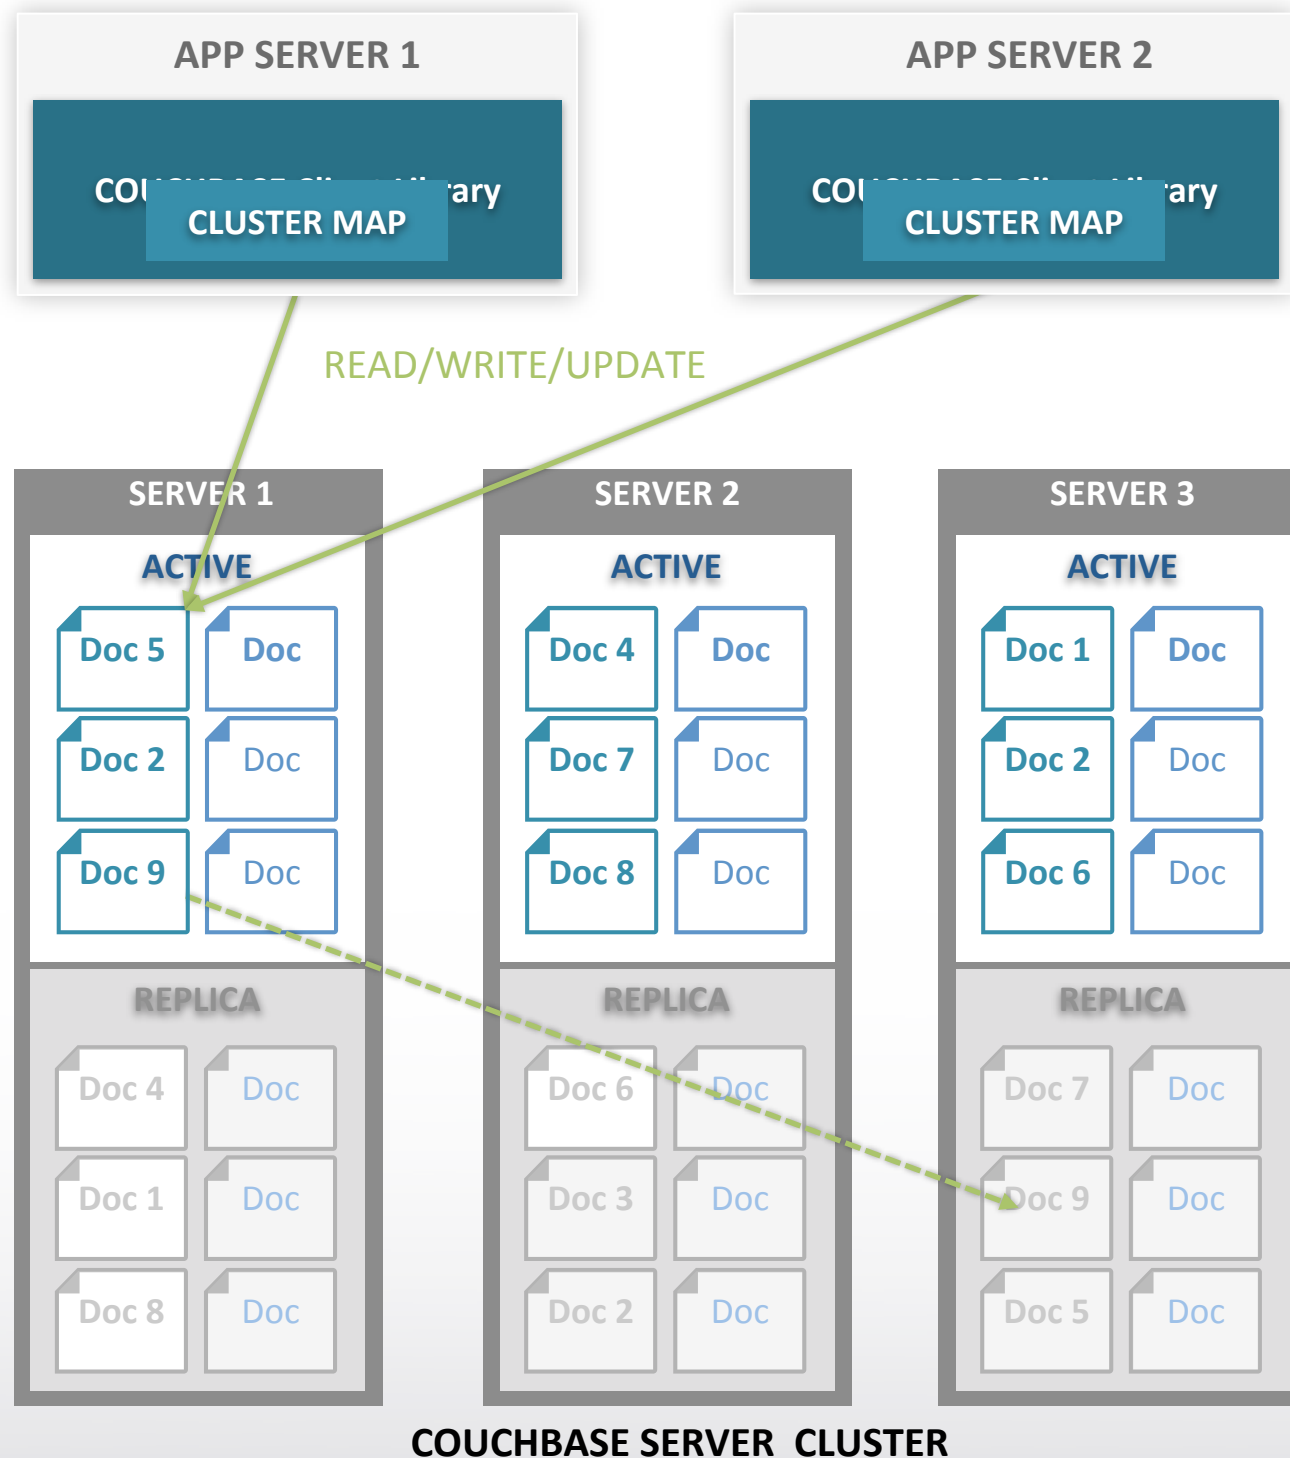


COUCHBASE SERVER CLUSTER

Couchbase

- Docs distributed evenly across servers
- Each server stores both active and replica docs
 - Only one doc 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

Basic Operations



- Docs distributed evenly across servers
- Each server stores both active and replica docs
 - Only one doc 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

Couchbase

Store & Retrieve Operations

- **get (key)**
 - Retrieve a document
- **set (key, value)**
 - Store a document, overwrites if exists
- **add (key, value)**
 - Store a document, error/exception if exists
- **replace (key, value)**
 - Store a document, error/exception if doesn't exist
- **cas (key, value, cas)**
 - Compare and swap, mutate document only if it hasn't changed while executing this operation

Atomic Counter Operations

These operations are always executed in order atomically.

- **set (key, value)**

- Use set to initialize the counter

- `cb.set("my_counter", 1)`

- **incr (key)**

- Increase an atomic counter value, default by 1

- `cb.incr("my_counter")` # now it's 2

- **decr (key)**

- Decrease an atomic counter value, default by 1

- `cb.decr("my_counter")` # now it's 1

Mental Adjustments


- In SQL we tend to want to avoid hitting the database as much as possible
 - Even with caching and indexing tricks, and massive improvements over the years, SQL still gets bogged down by complex joins and huge indexes, so we avoid making database calls
- In Couchbase, get's and set's are so fast *they are trivial*, not bottlenecks, this is hard for many people to accept at first; Multiple get statements are commonplace, don't avoid it!

JSON Document Structure

Meta Information

Including Key

All Keys Unique and
Kept in RAM



meta

```
{
  "id": "u::jasdeep@couchbase.com",
  "rev": "1-0002bce00000000000",
  "flags": 0,
  "expiration": 0,
  "type": "json"
}
```

Document Value

Most Recent In Ram
And Persisted To Disk



document

```
{
  "uid": 123456,
  "firstname": "jasdeep",
  "lastname": "Jaitla",
  "age": 22,
  "favorite_colors": ["blue", "black"],
  "email": "jasdeep@couchbase.com"
}
```

COUCHBASE

Couchbase

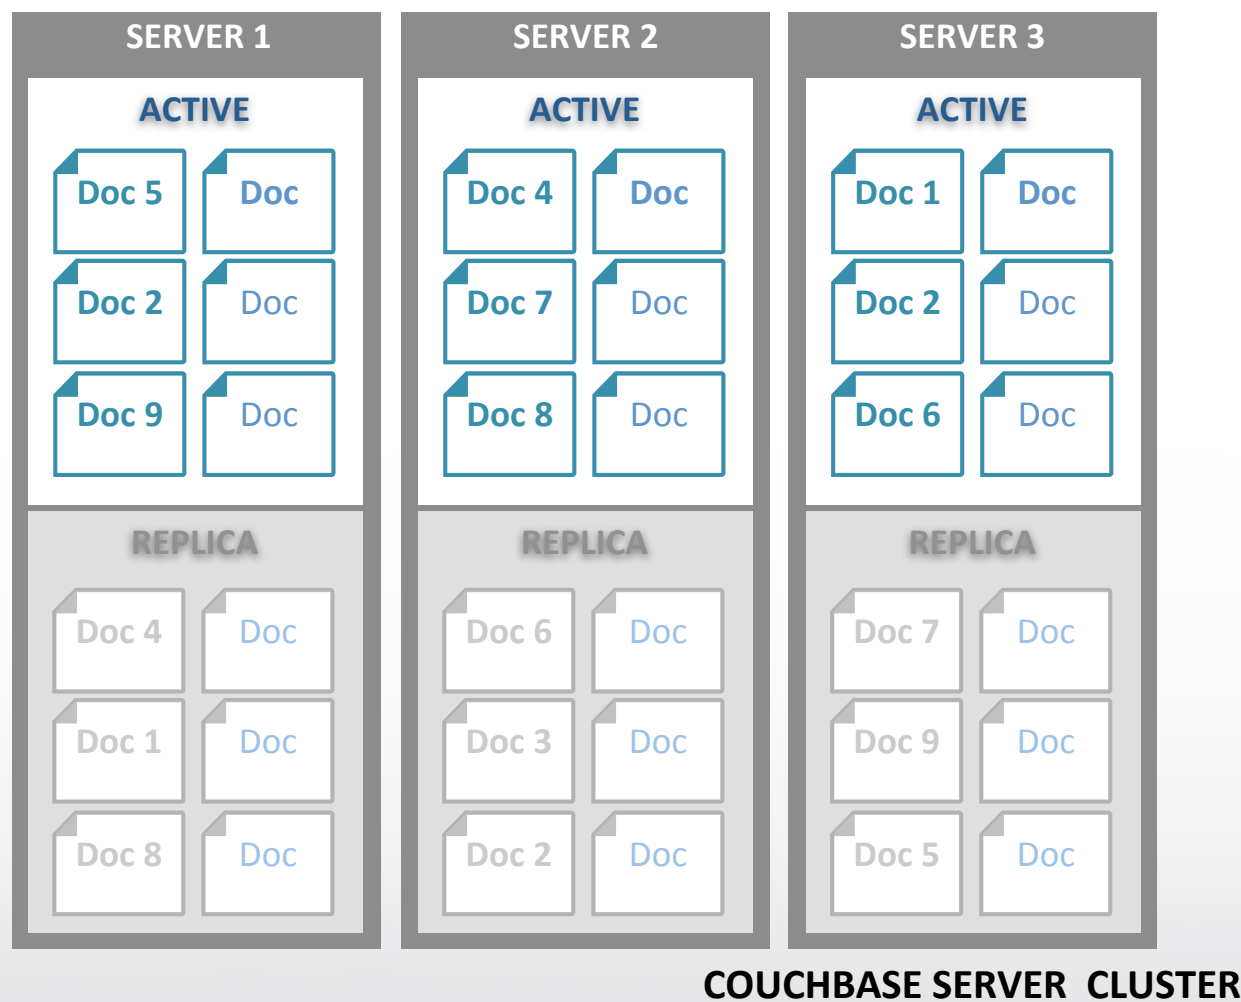




Couchbase

DEMONSTRATION




Couchbase

Add Nodes to Cluster



Couchbase




Add Nodes to Cluster



- Two servers added
One-click operation

Couchbase

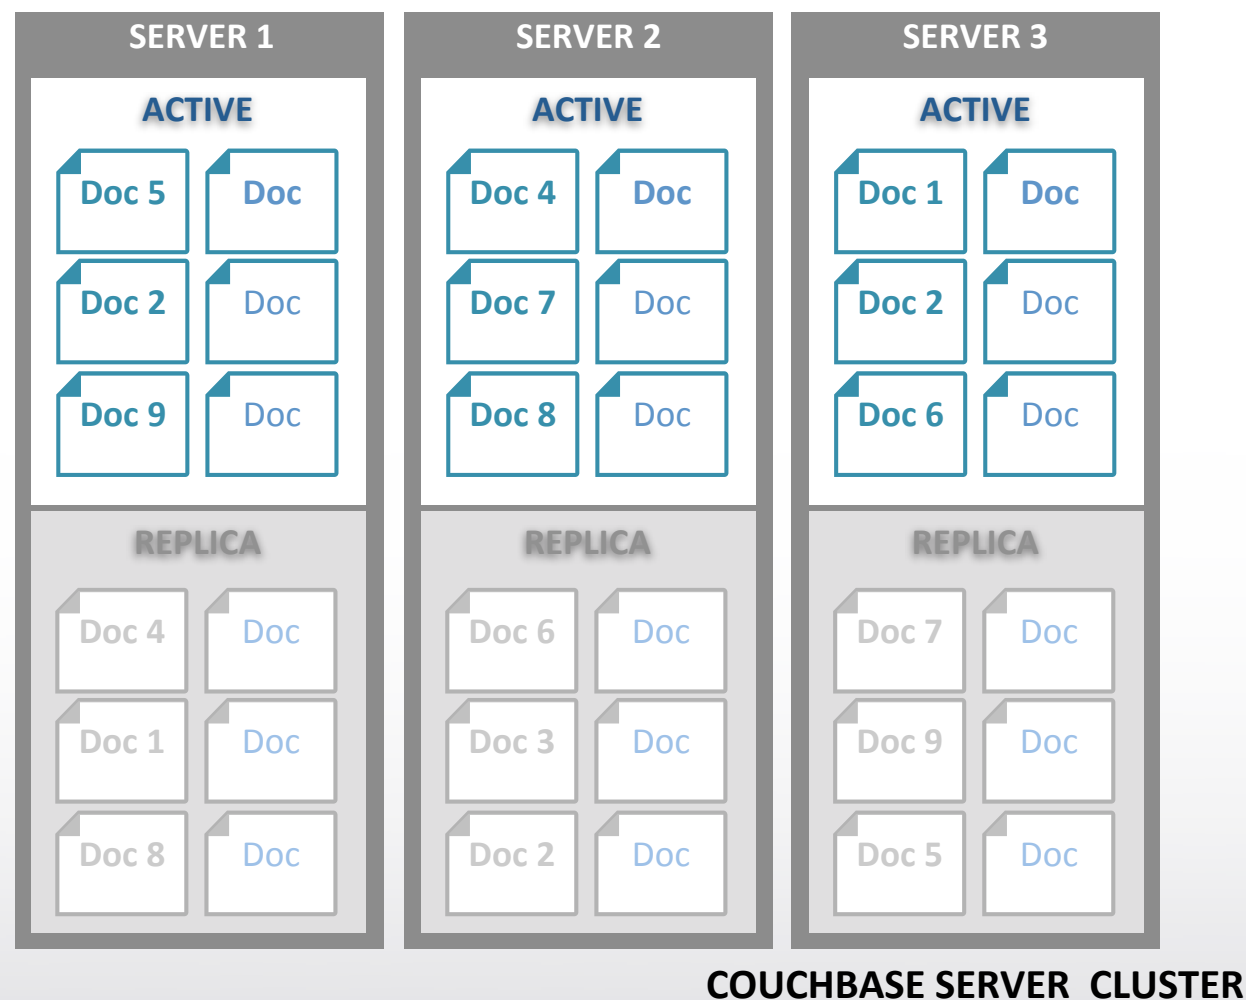

Add Nodes to Cluster



- Two servers added
One-click operation
- Docs automatically rebalanced across cluster
Even distribution of docs
Minimum doc movement

Couchbase

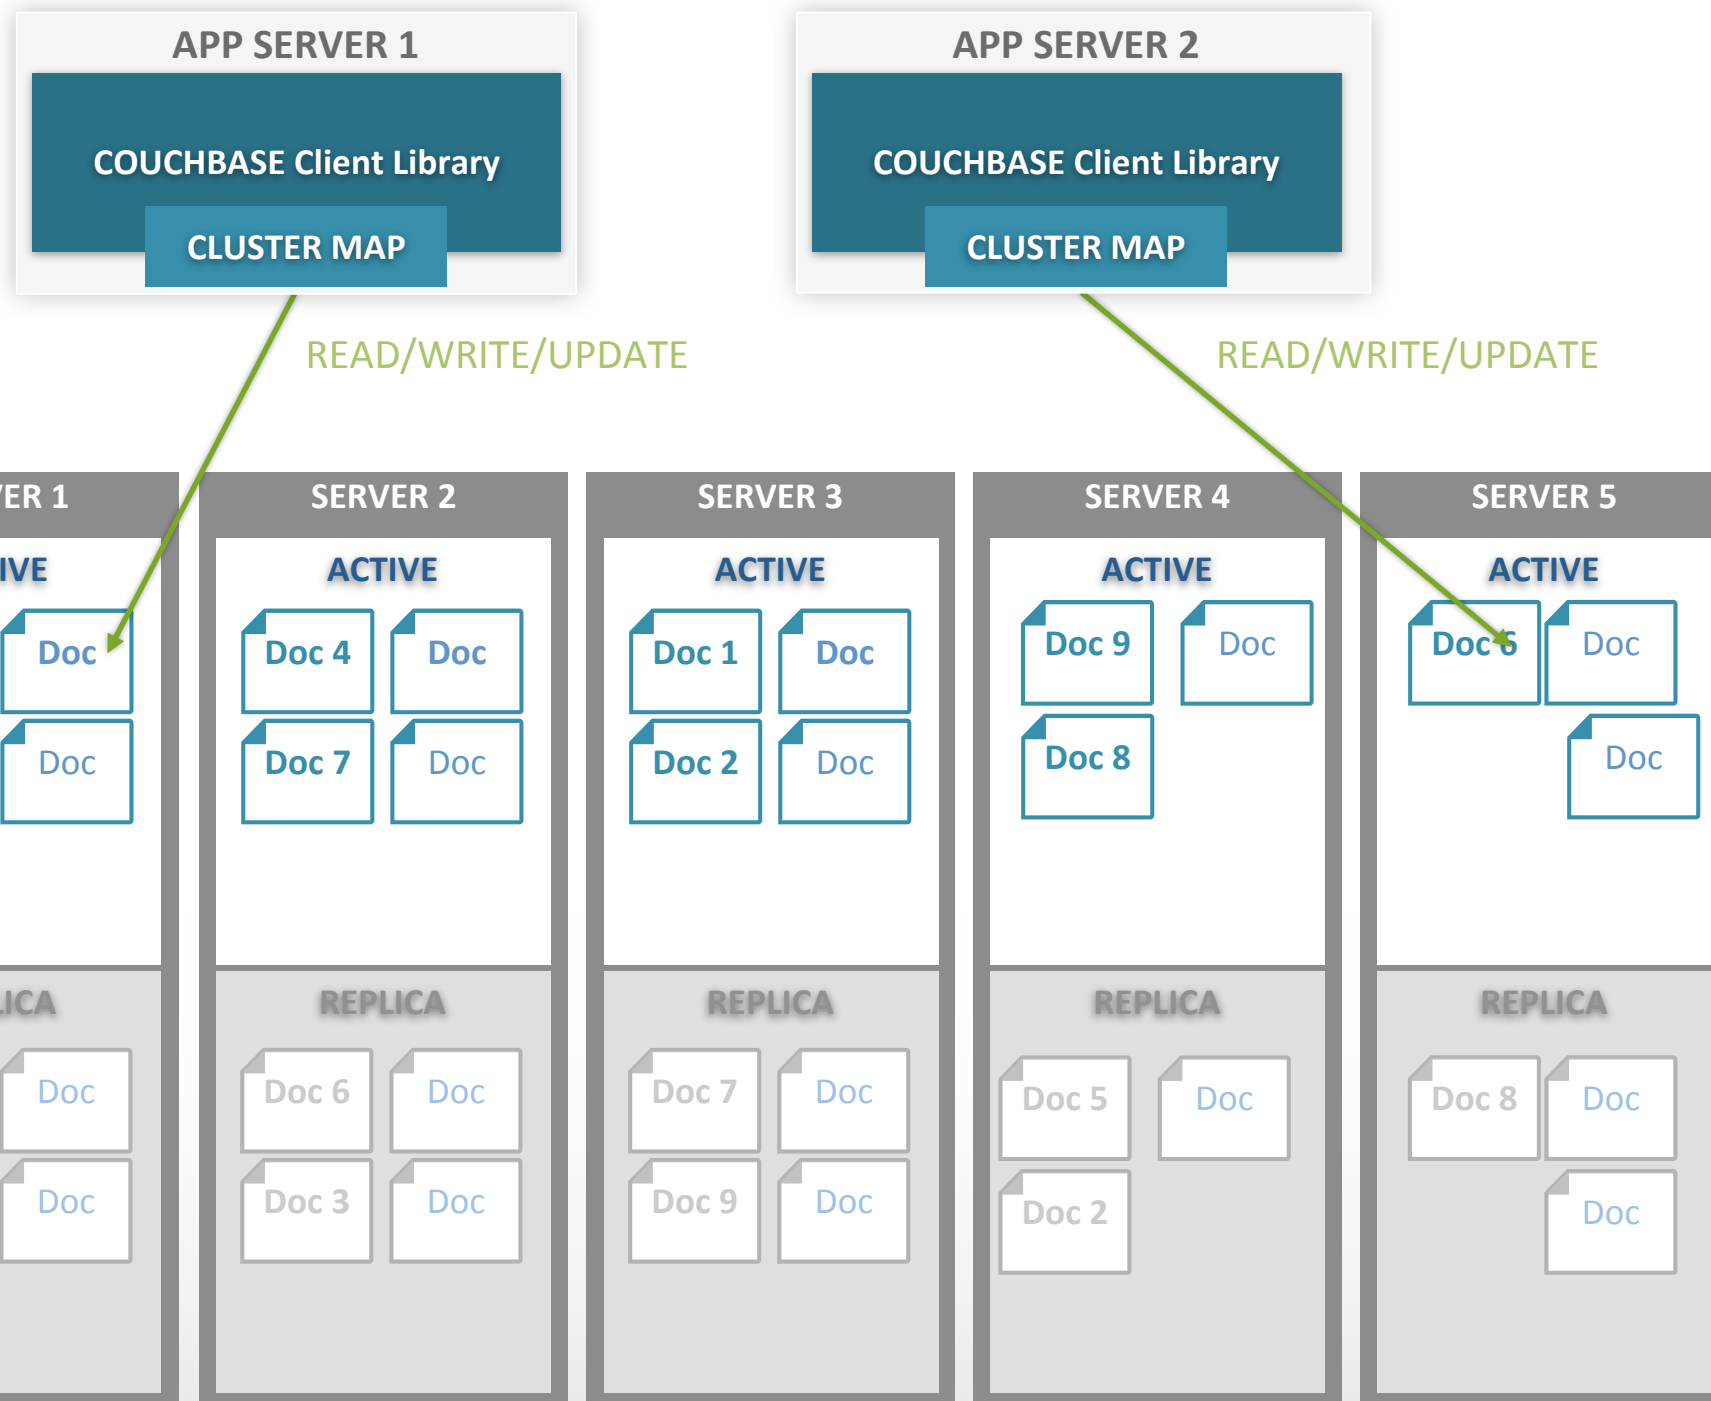
Add Nodes to Cluster



- Two servers added
One-click operation
- Docs automatically rebalanced across cluster
Even distribution of docs
Minimum doc movement
- Cluster map updated

Couchbase

Add Nodes to Cluster

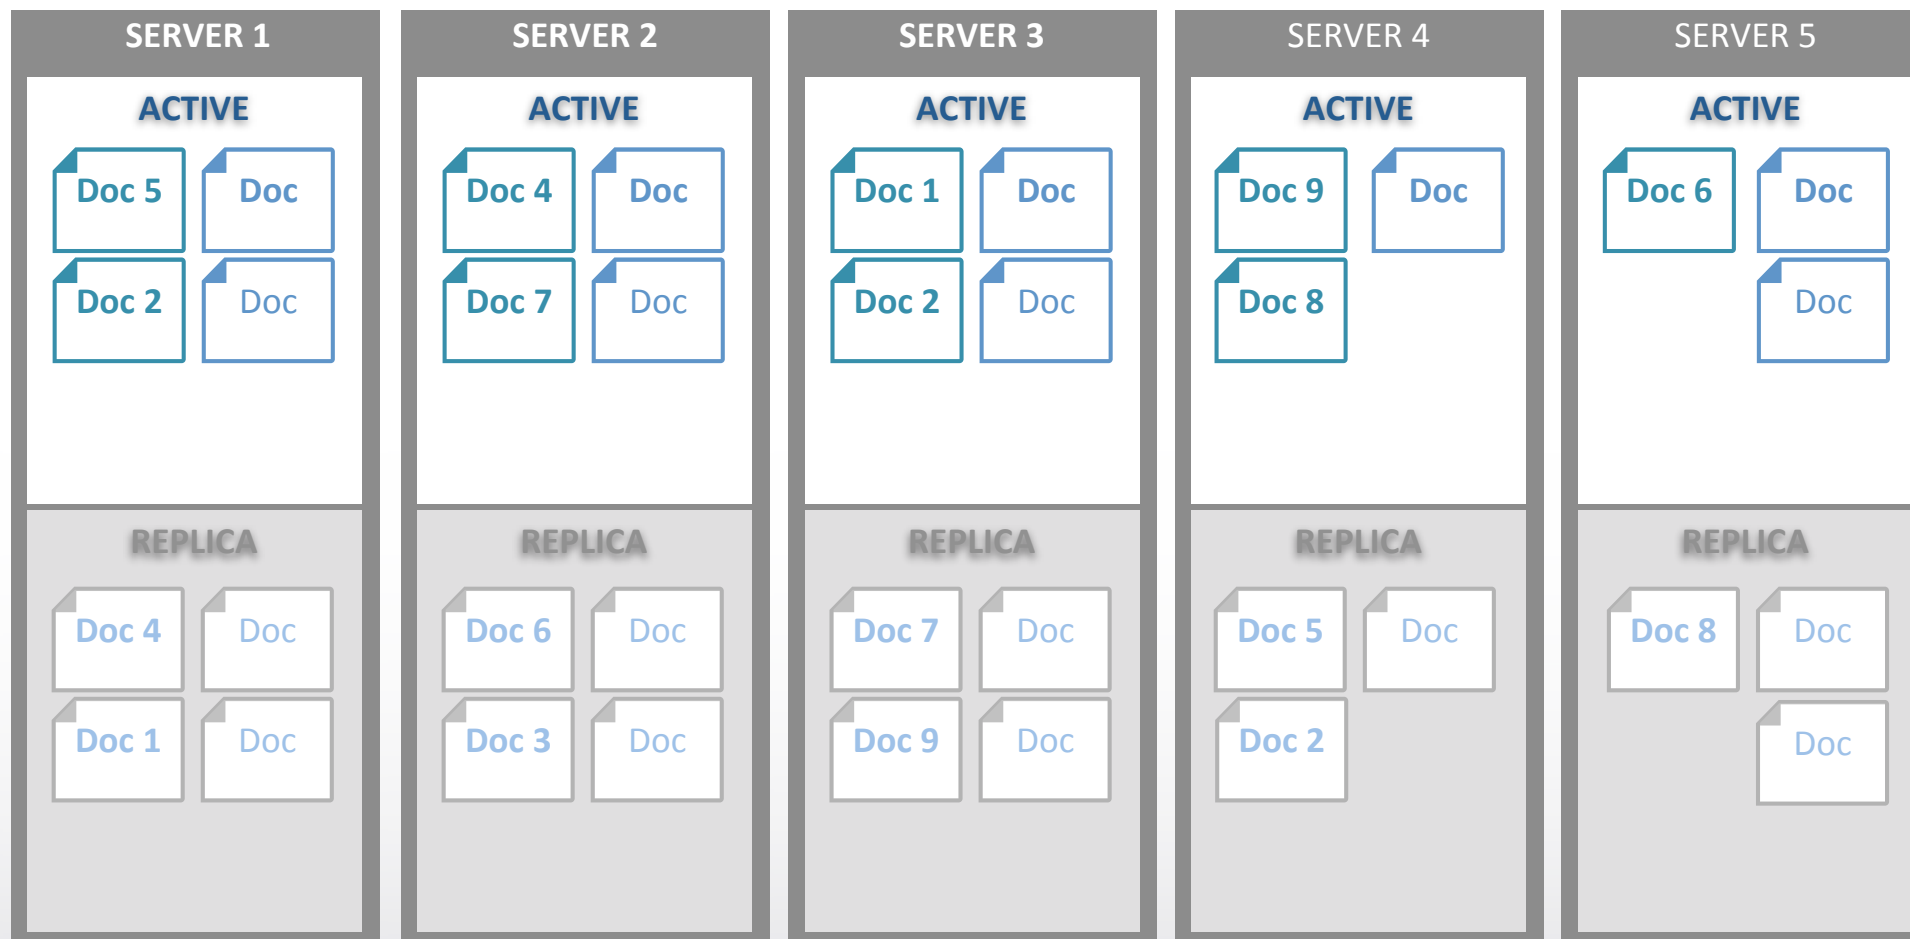
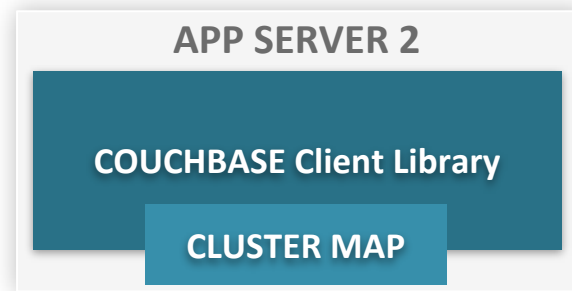
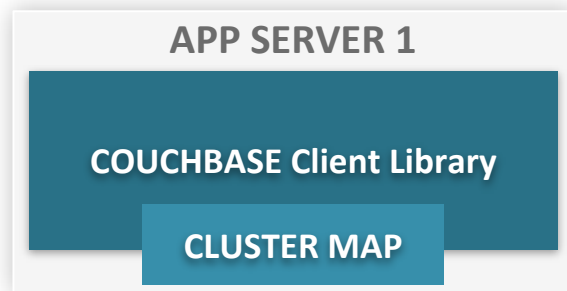


- Two servers added
One-click operation
- Docs automatically rebalanced across cluster
Even distribution of docs
Minimum doc movement
- Cluster map updated
- App database calls now distributed over larger number of servers

COUCHBASE SERVER CLUSTER

Couchbase


Fail Over Node



COUCHBASE SERVER CLUSTER

Couchbase

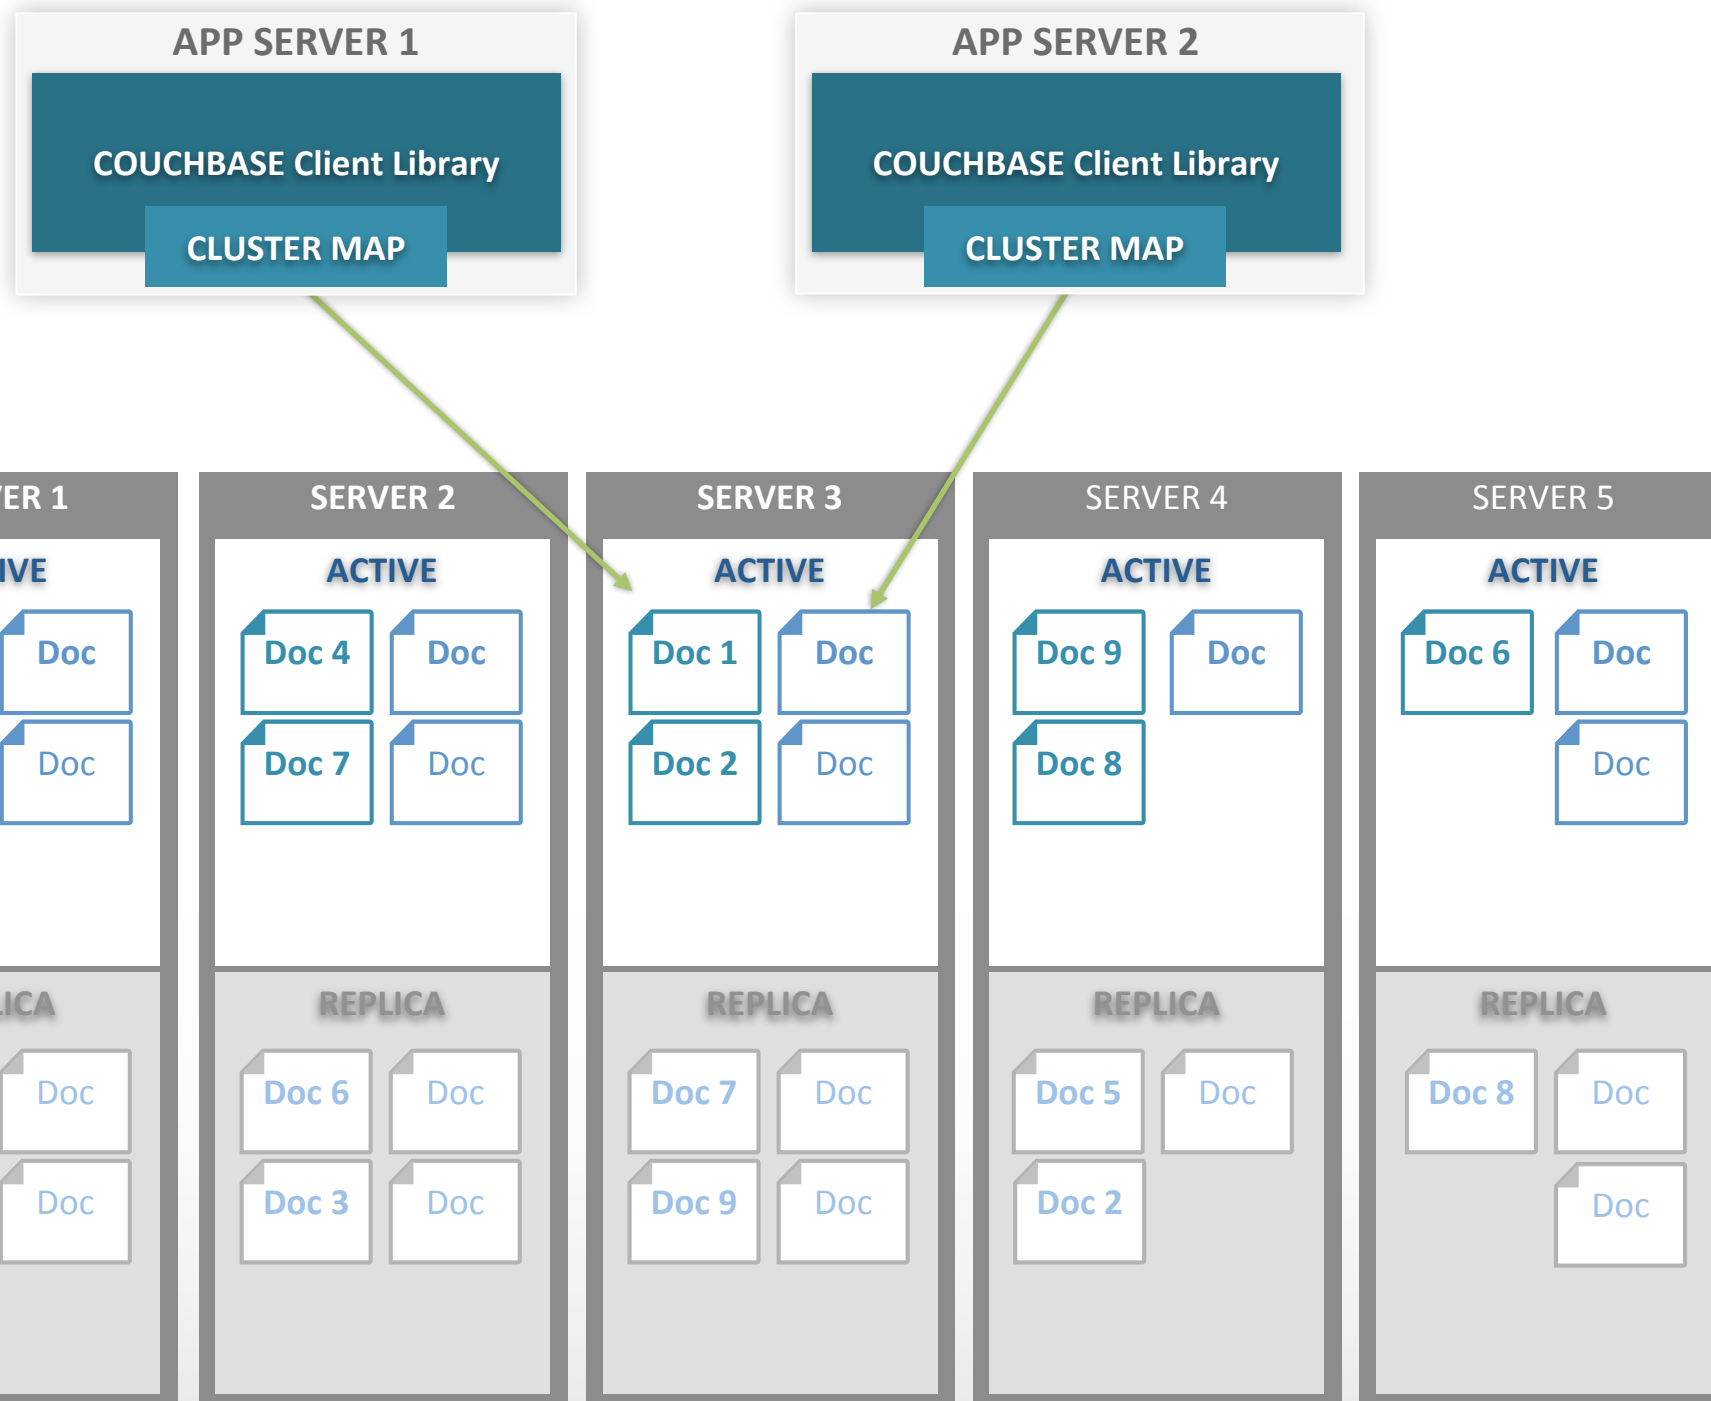
Fail Over Node



COUCHBASE SERVER CLUSTER

Couchbase

Fail Over Node

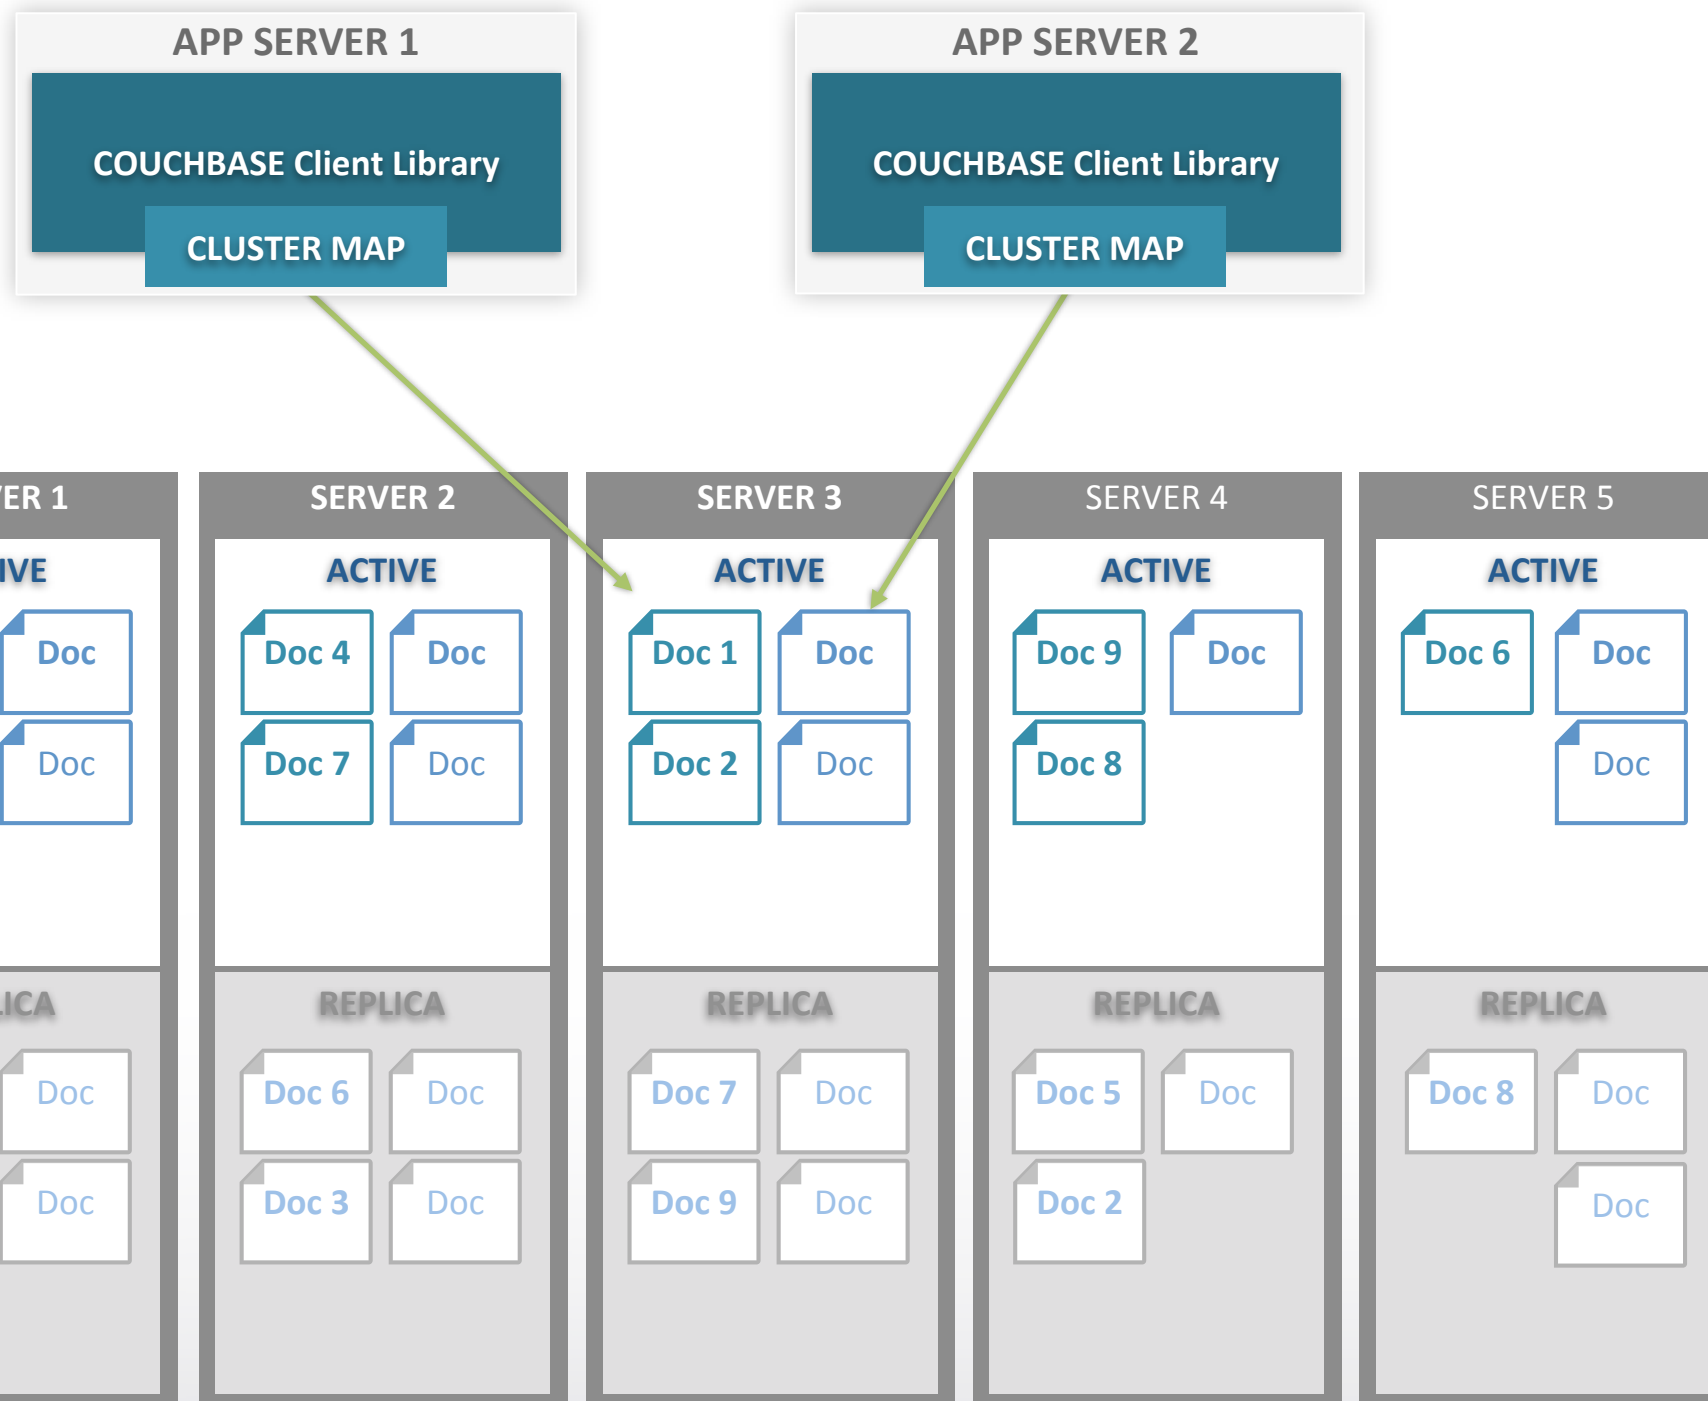


- App servers accessing docs

COUCHBASE SERVER CLUSTER

Couchbase

Fail Over Node




- App servers accessing docs
- Requests to Server 3 fail

COUCHBASE SERVER CLUSTER

Couchbase

Fail Over Node




- App servers accessing docs
- Requests to Server 3 fail
- Cluster detects server failed

COUCHBASE SERVER CLUSTER

Couchbase

Fail Over Node




- App servers accessing docs
- Requests to Server 3 fail
- Cluster detects server failed
Promotes replicas of docs to active

COUCHBASE SERVER CLUSTER

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

COUCHBASE SERVER CLUSTER

Couchbase

Fail Over Node

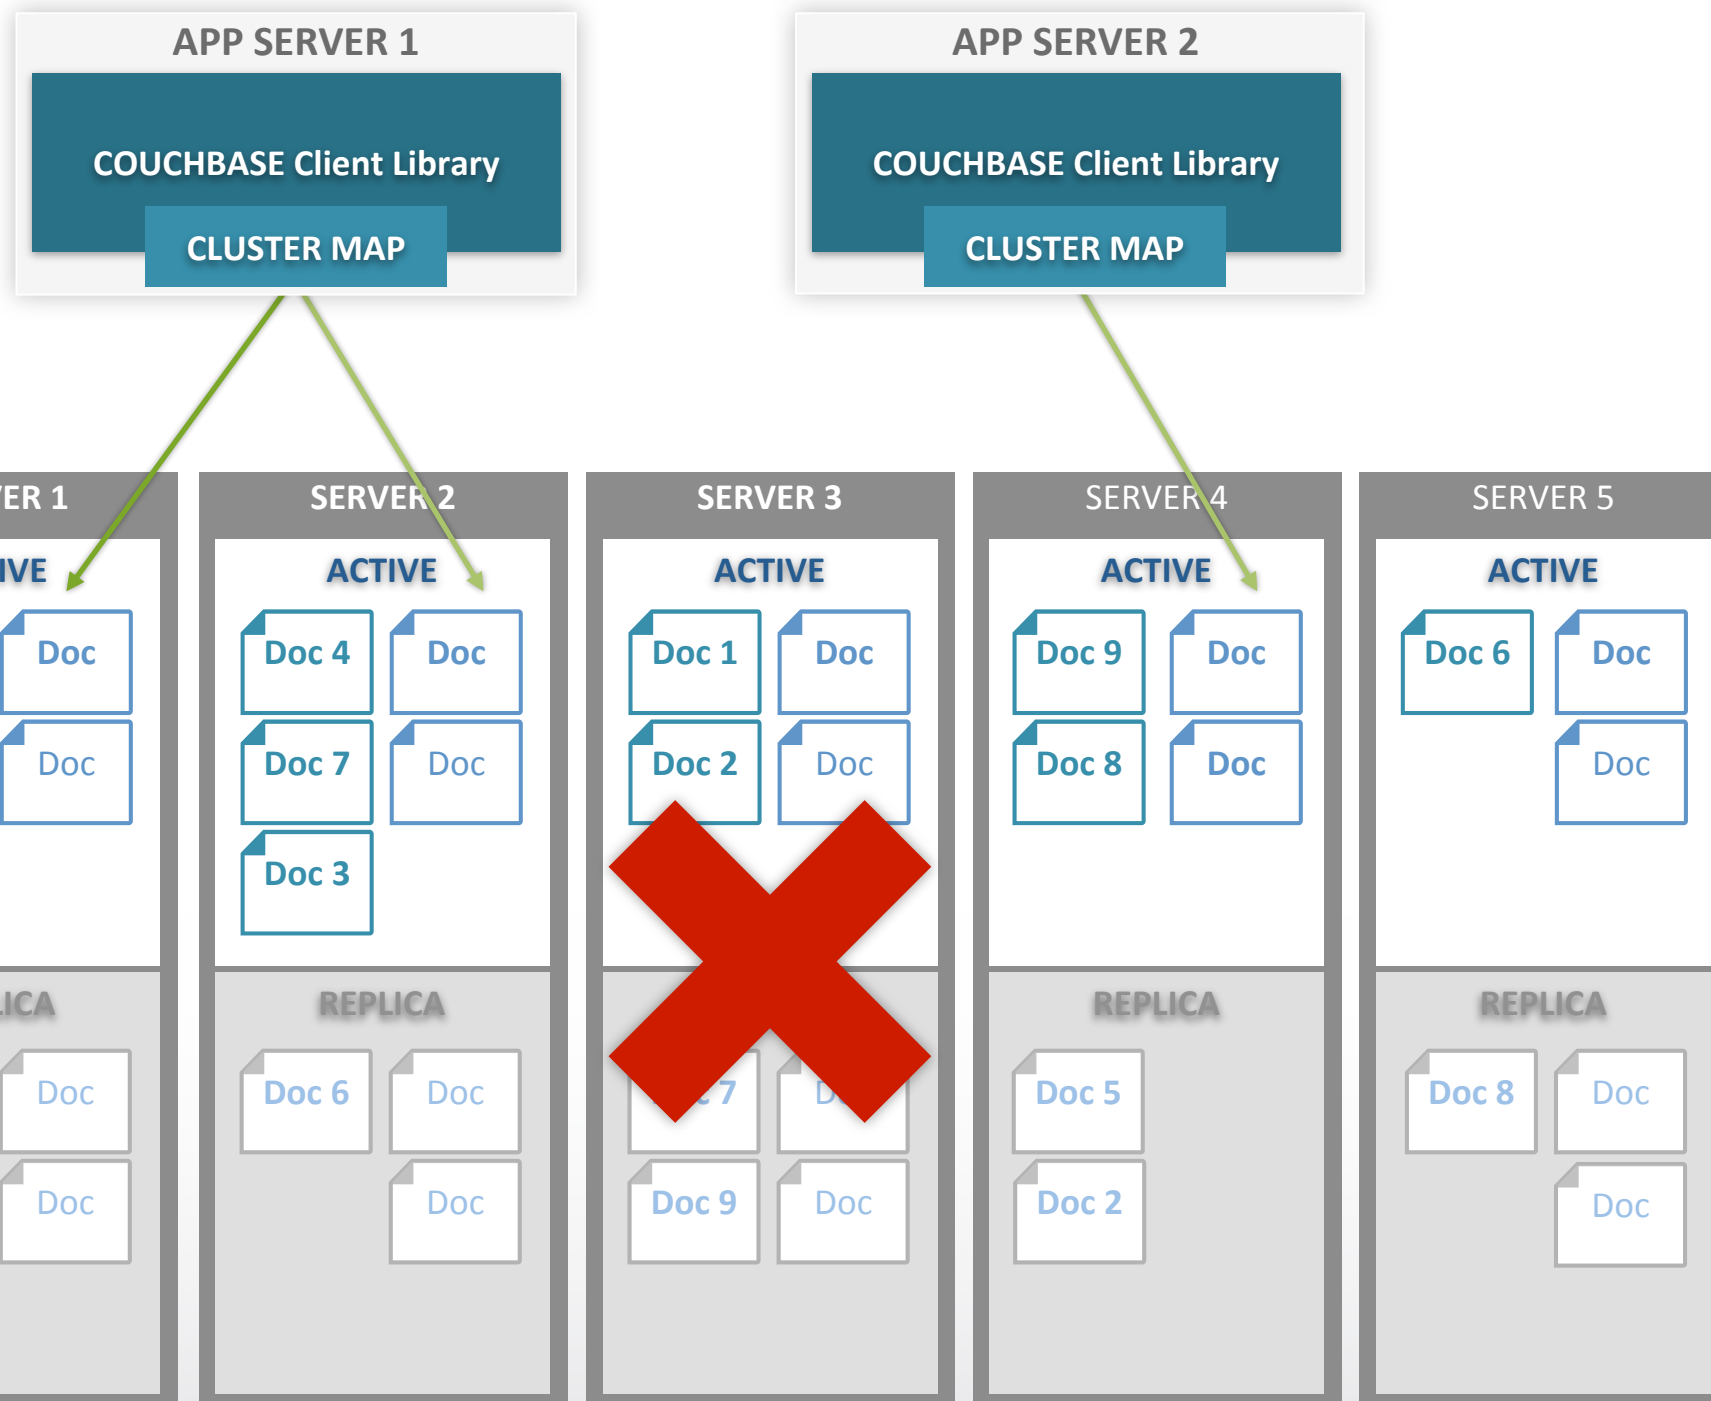


- App servers accessing docs
- Requests to Server 3 fail
- Cluster detects server failed
Promotes replicas of docs to active
Updates cluster map
- Requests for docs now go to appropriate server

COUCHBASE SERVER CLUSTER

Couchbase

Fail Over Node

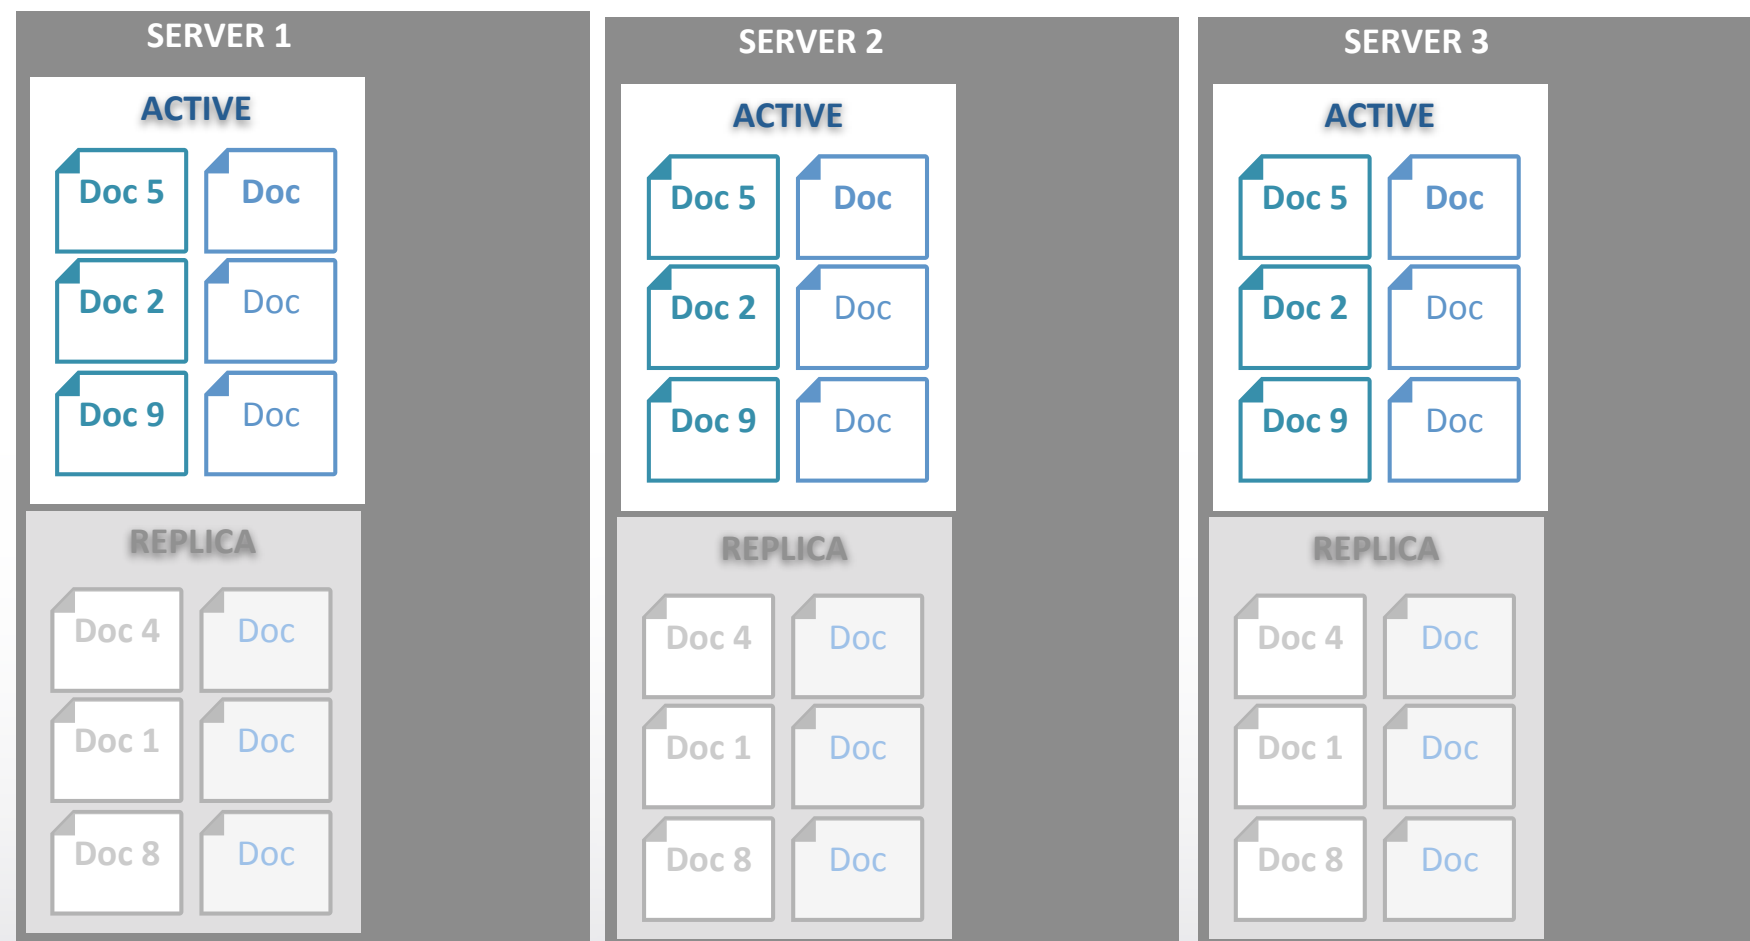



- App servers accessing docs
- Requests to Server 3 fail
- Cluster detects server failed
Promotes replicas of docs to active
Updates cluster map
- Requests for docs now go to appropriate server
- Typically rebalance would follow

COUCHBASE SERVER CLUSTER

Couchbase

Indexing and Querying



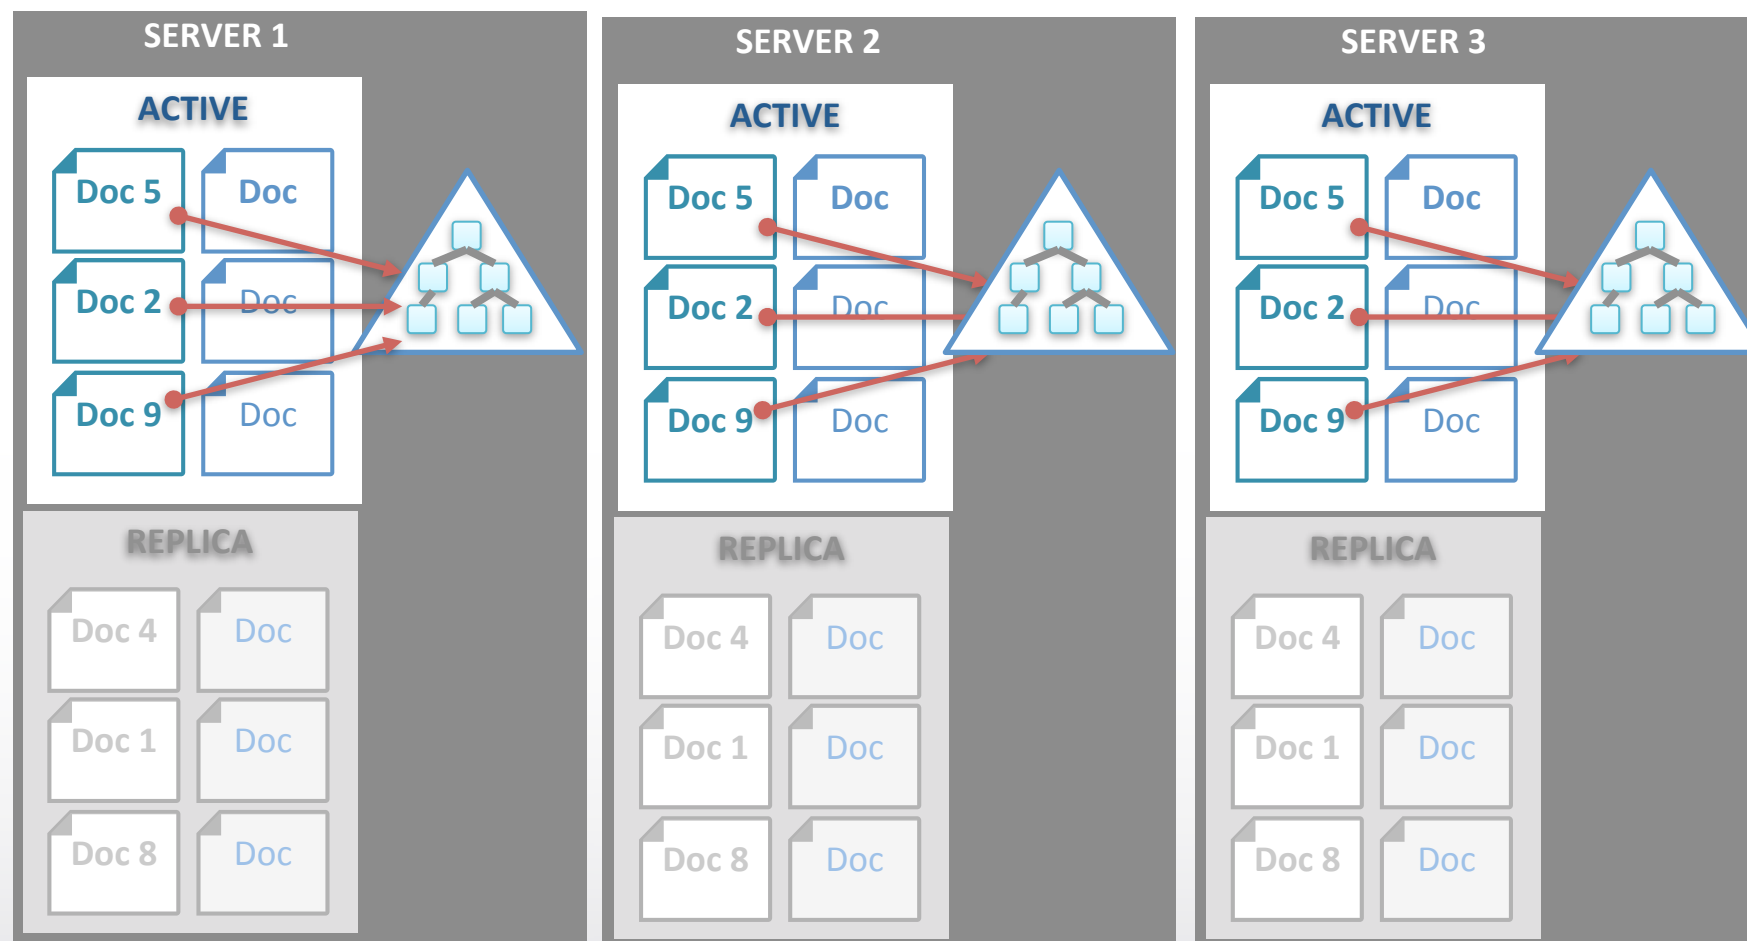

COUCHBASE SERVER CLUSTER

Couchbase

- 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



Indexing and Querying



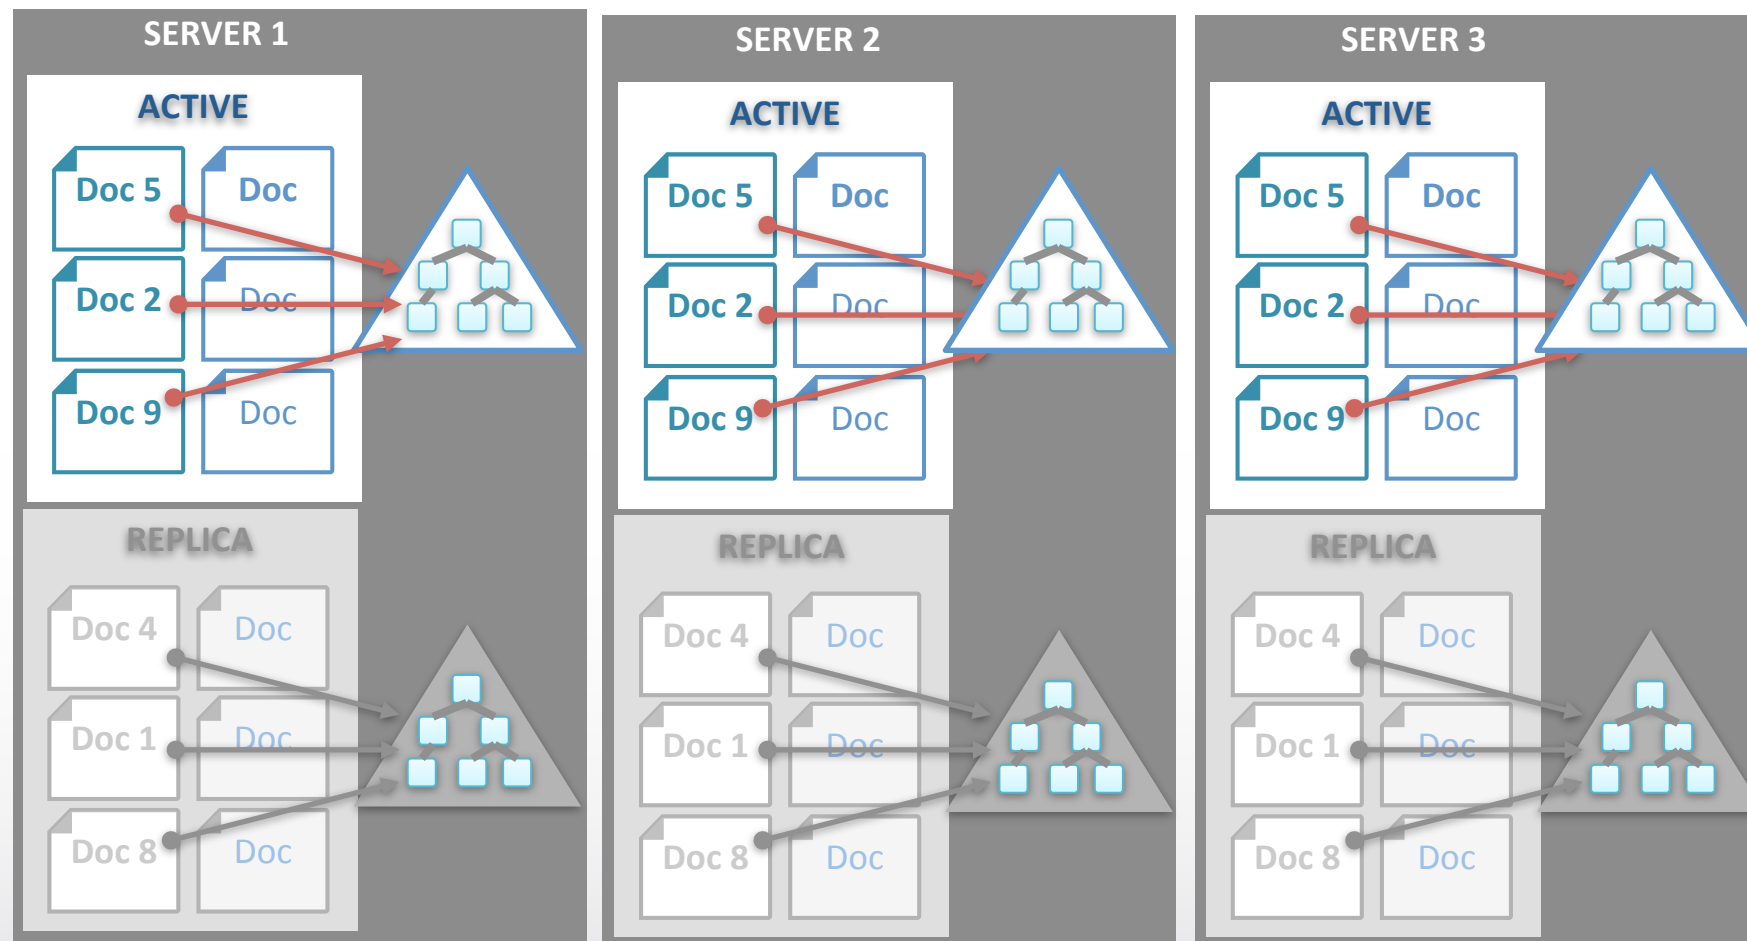

COUCHBASE SERVER CLUSTER

Couchbase

- 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



Indexing and Querying



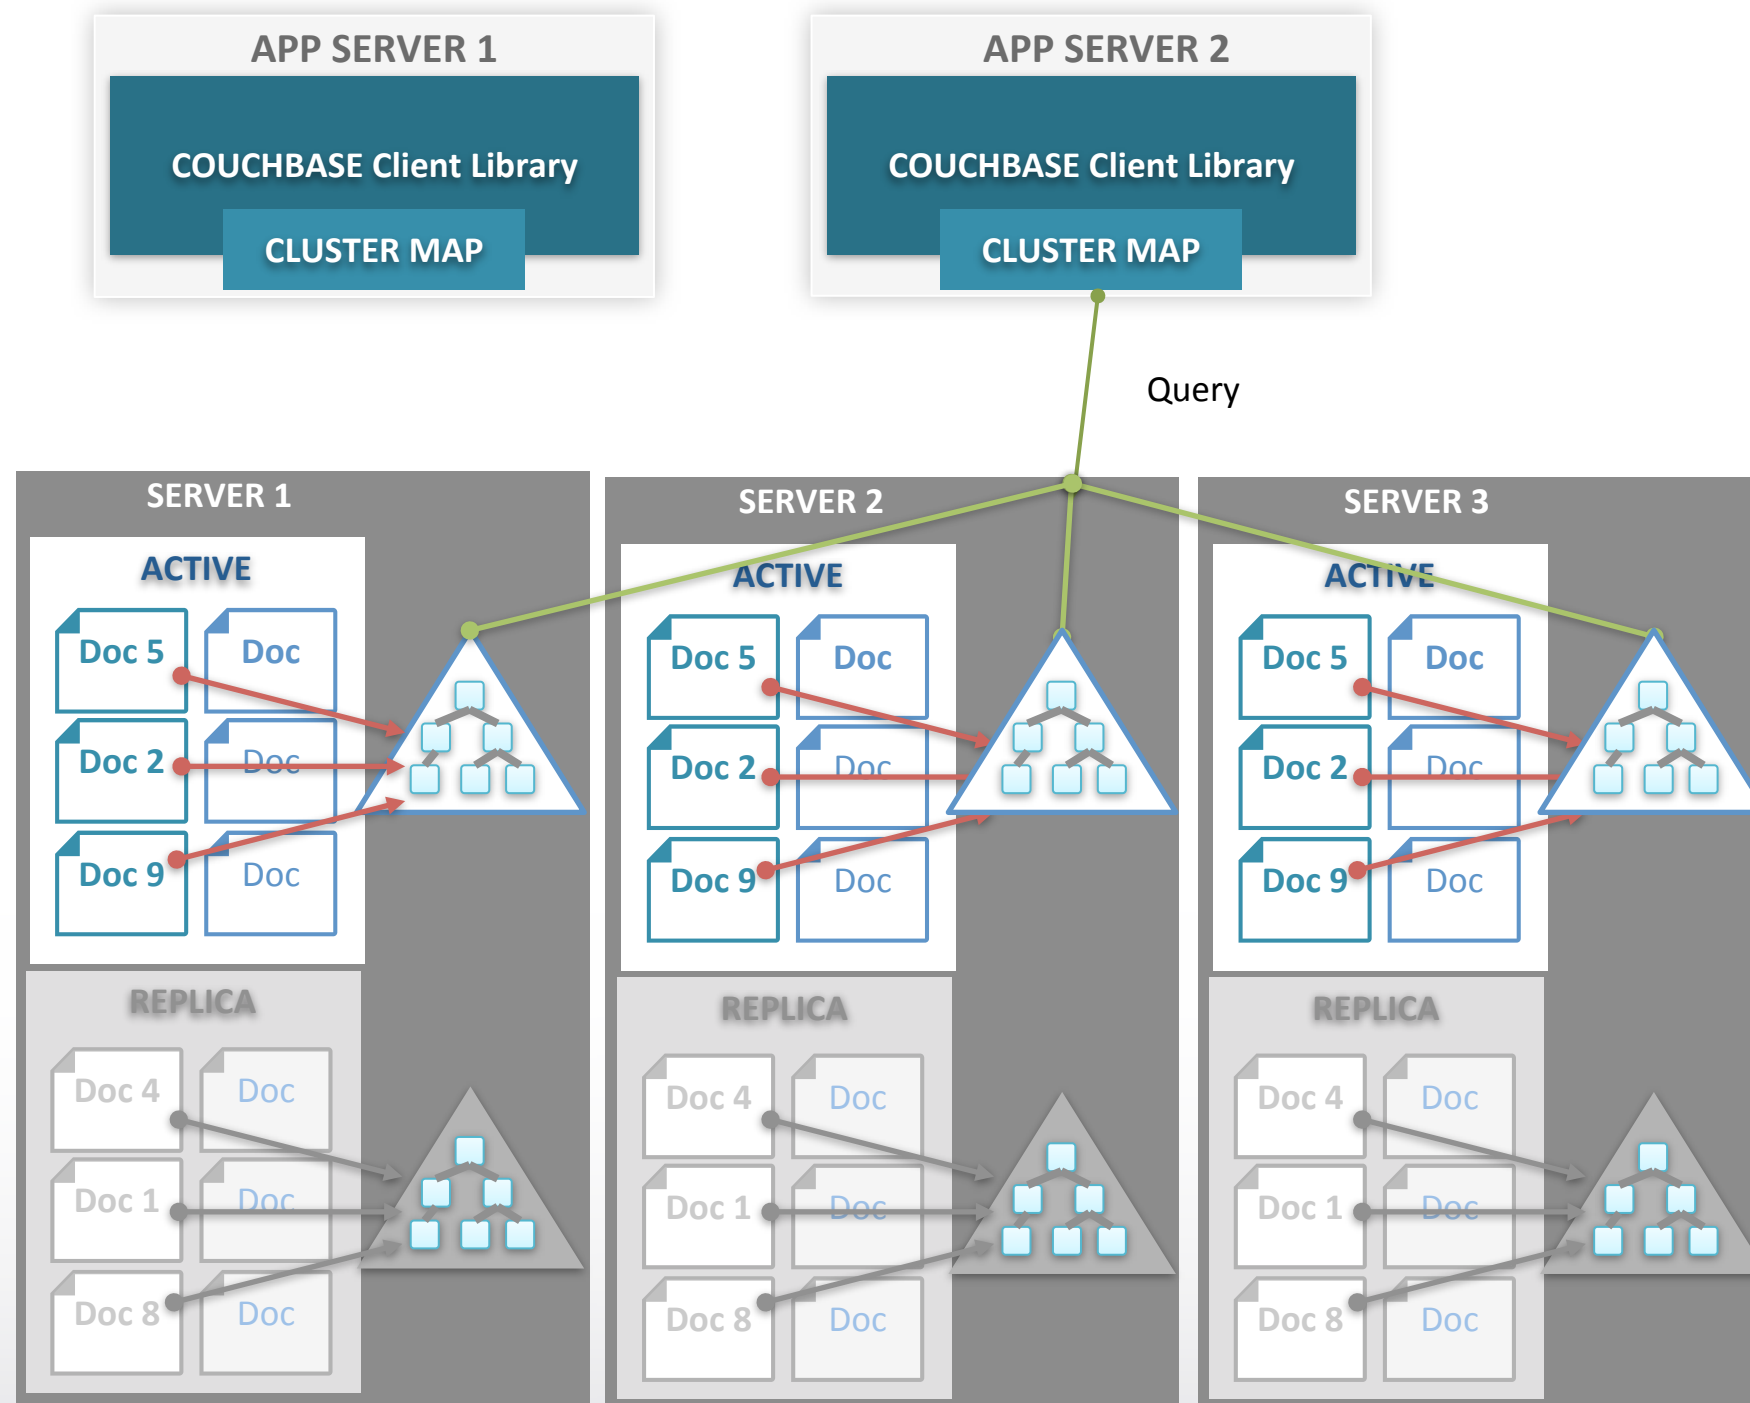
COUCHBASE SERVER CLUSTER

Couchbase

- 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



Indexing and Querying



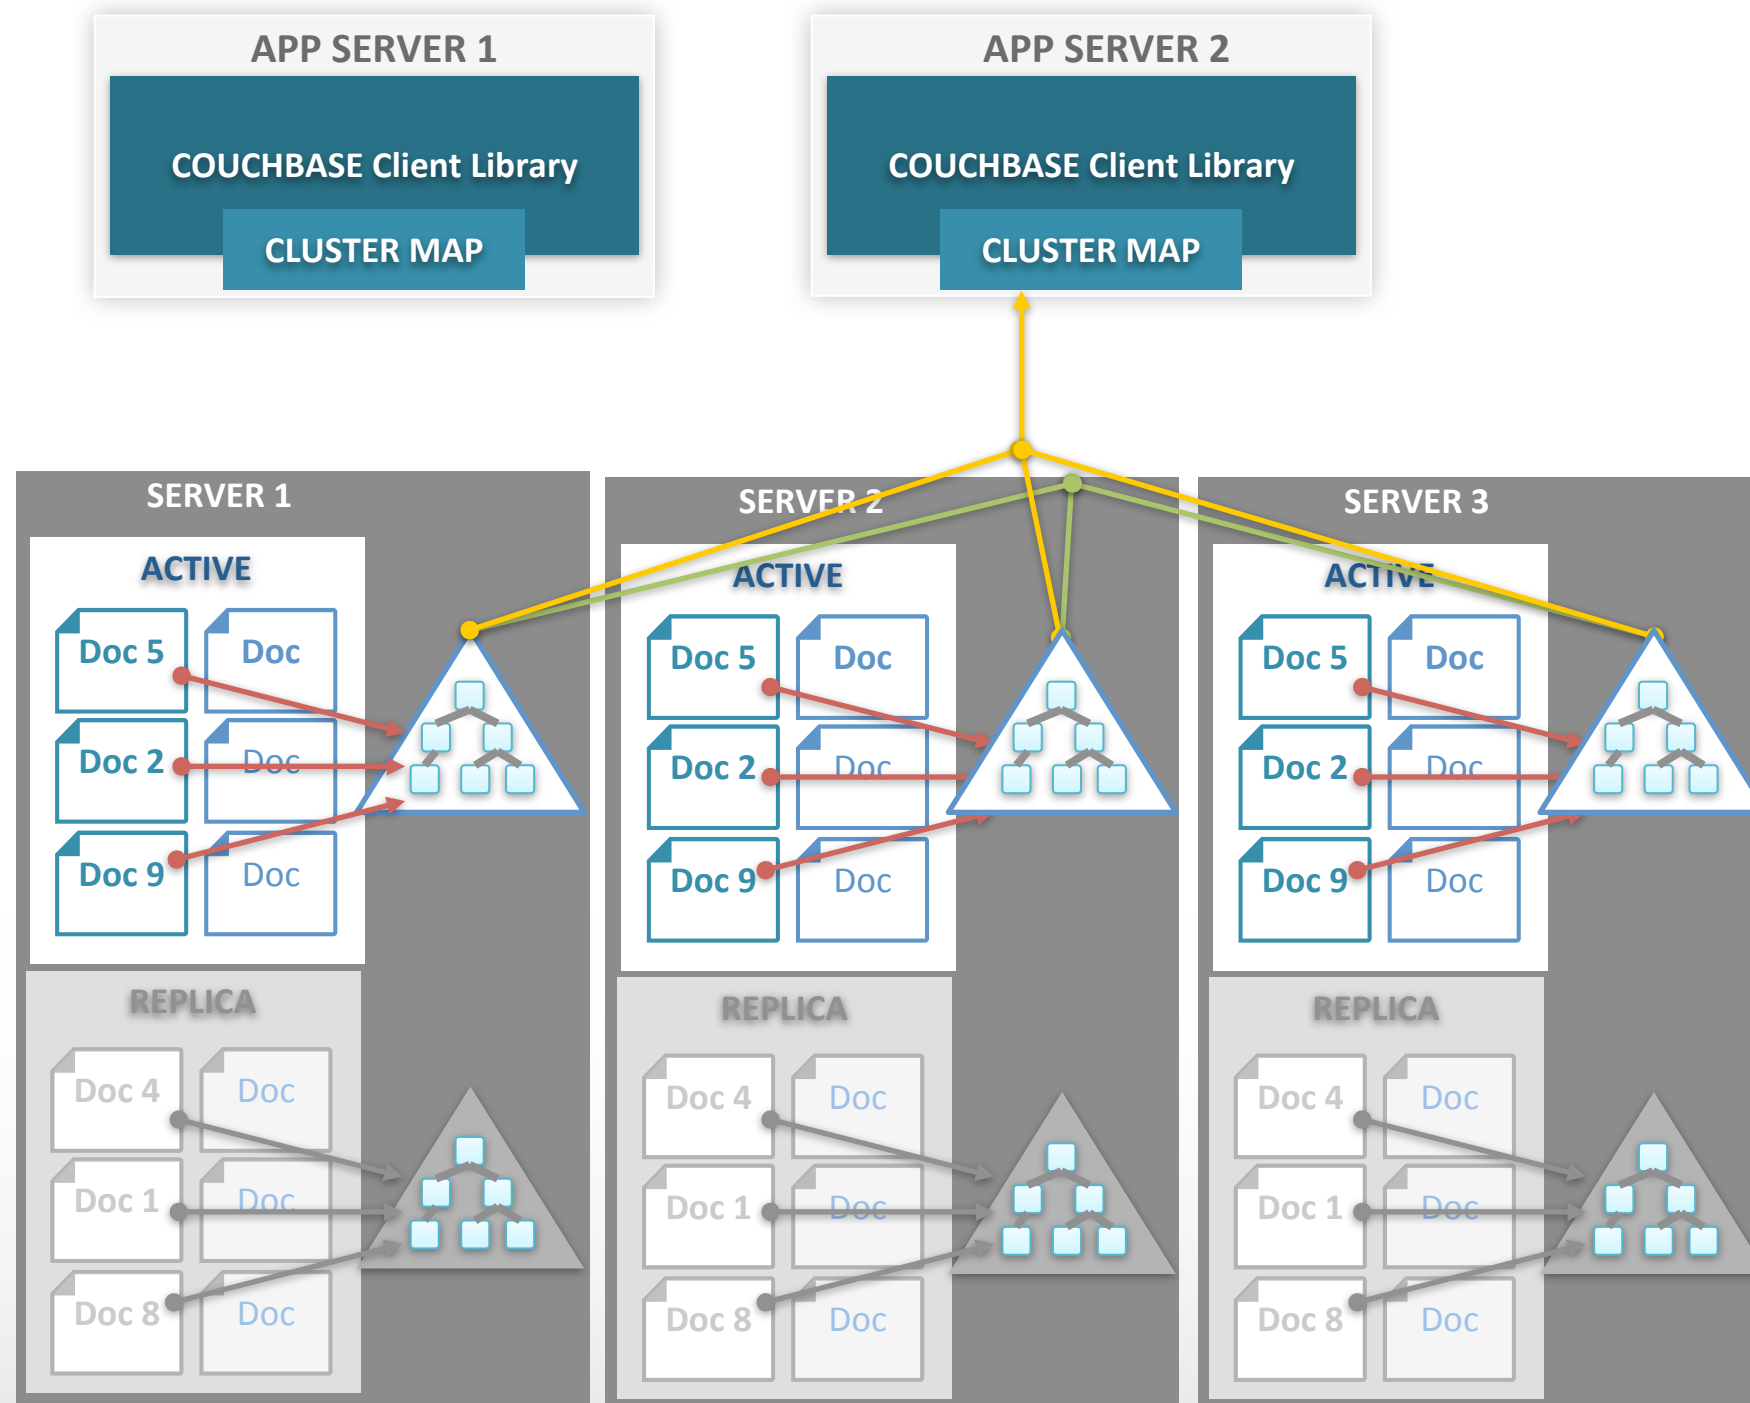
COUCHBASE SERVER CLUSTER

Couchbase

- 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



Indexing and Querying



COUCHBASE SERVER CLUSTER

Couchbase

- 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



Couchbase



Couchbase

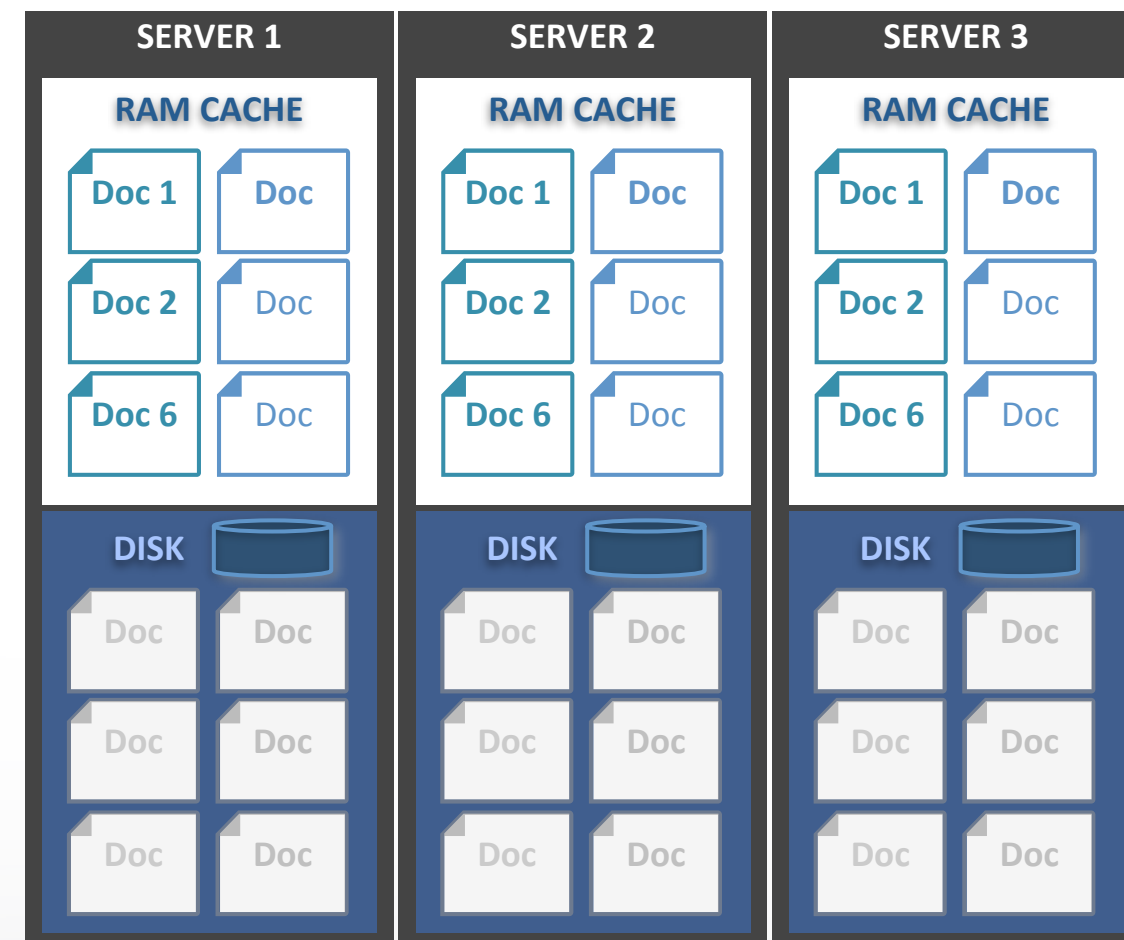
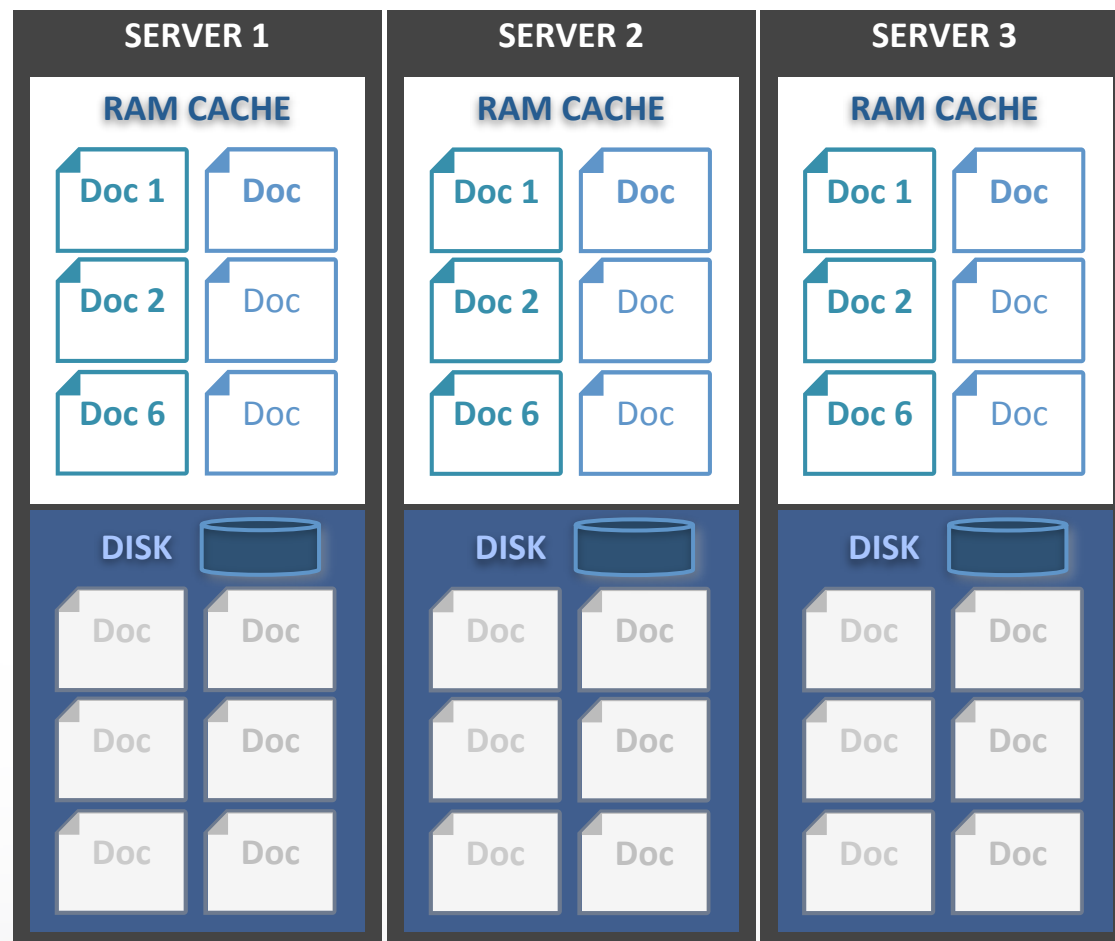
DEMONSTRATION

Couchbase

Cross Data Center Replication (XDCR)

Couchbase Cluster
West Coast Data Center

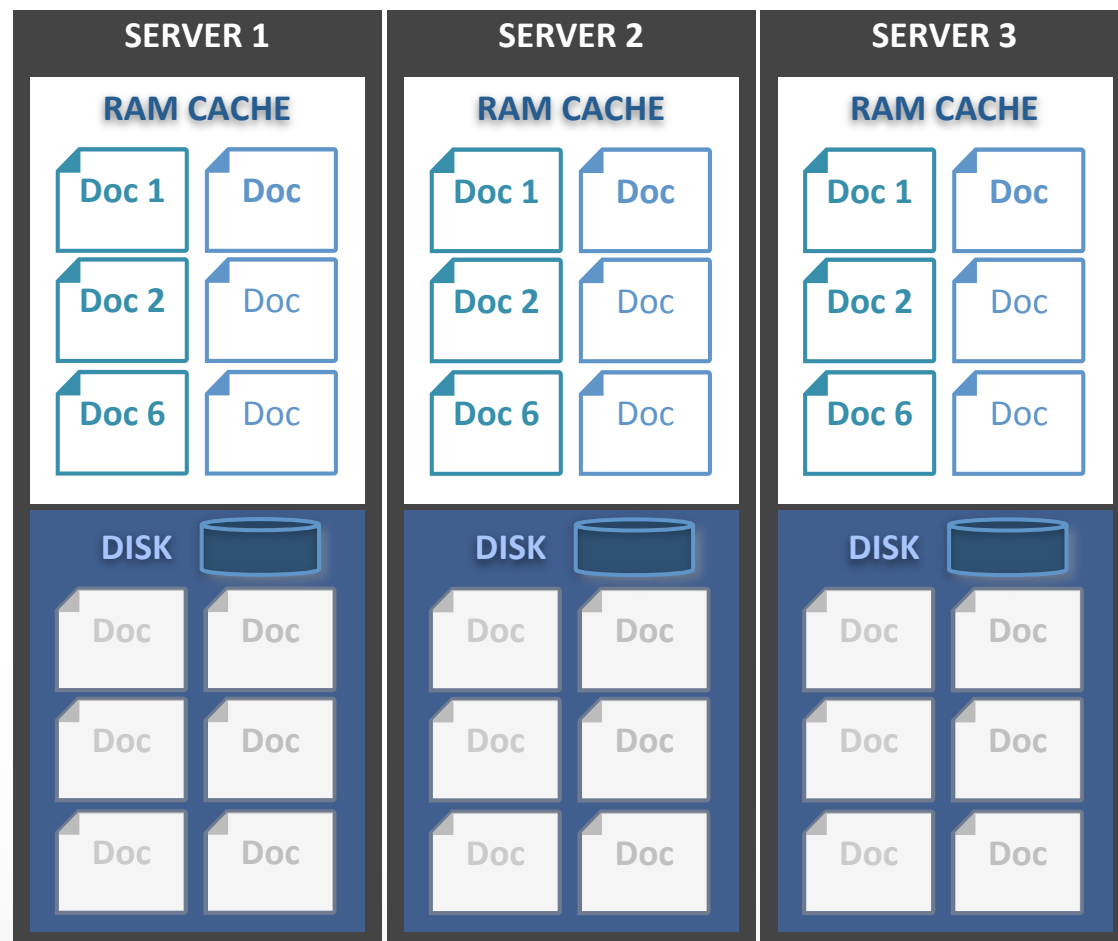
Couchbase Cluster
East Coast Data Center




Couchbase

Cross Data Center Replication (XDCR)

Couchbase Cluster
West Coast Data Center



Couchbase Cluster
East Coast Data Center



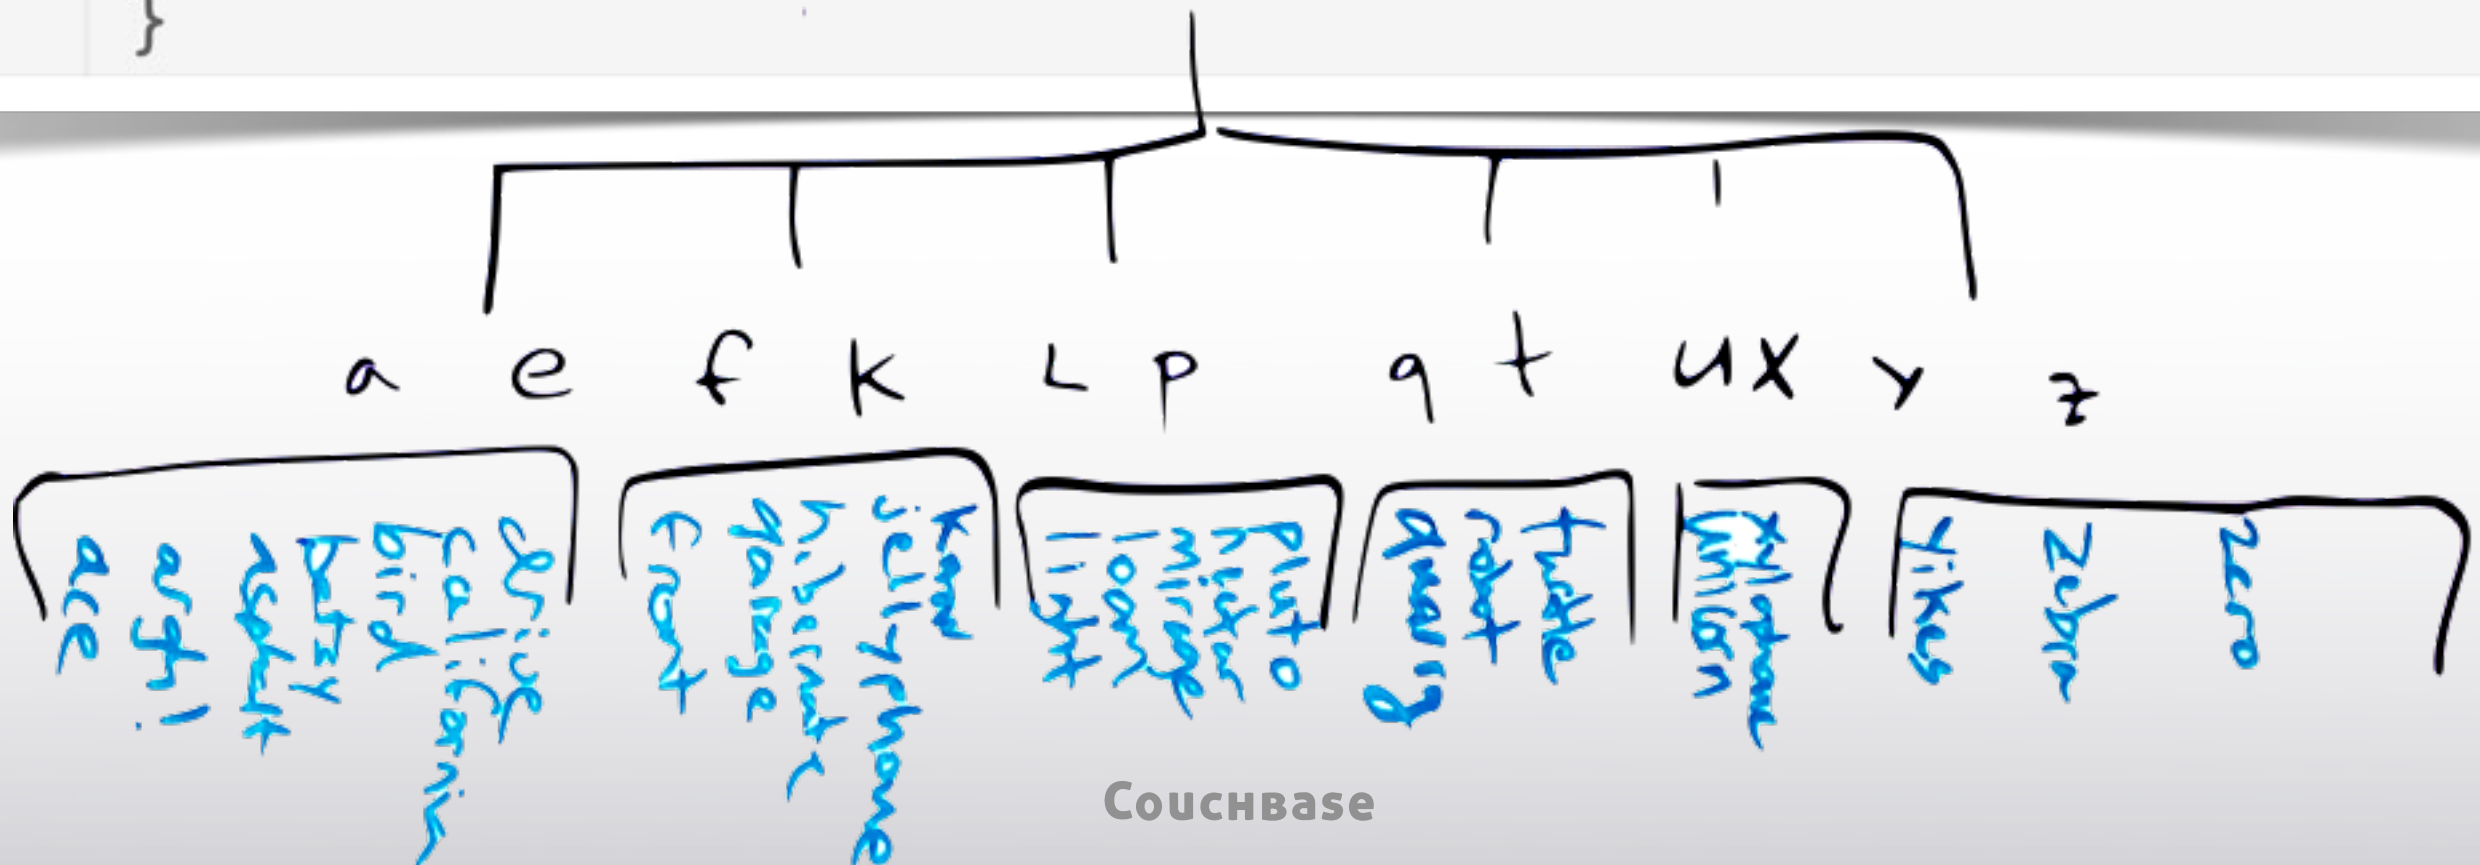
Couchbase

Map Function

```

1  function (doc, meta) {
2      if (doc.country, doc.state, doc.city) {
3          emit([doc.country, doc.state, doc.city], 1);
4      } else if (doc.country, doc.state) {
5          emit([doc.country, doc.state], 1);
6      } else if (doc.country) {
7          emit([doc.country], 1);
8      }
9  }

```



Couchbase

Couchbase



Couchbase


DEMONSTRATION

Couchbase

GeoCouch R-Tree Index

- Optimized for bulk loading of large data sets
- Simplified query model (bounding box, nearest neighbor)

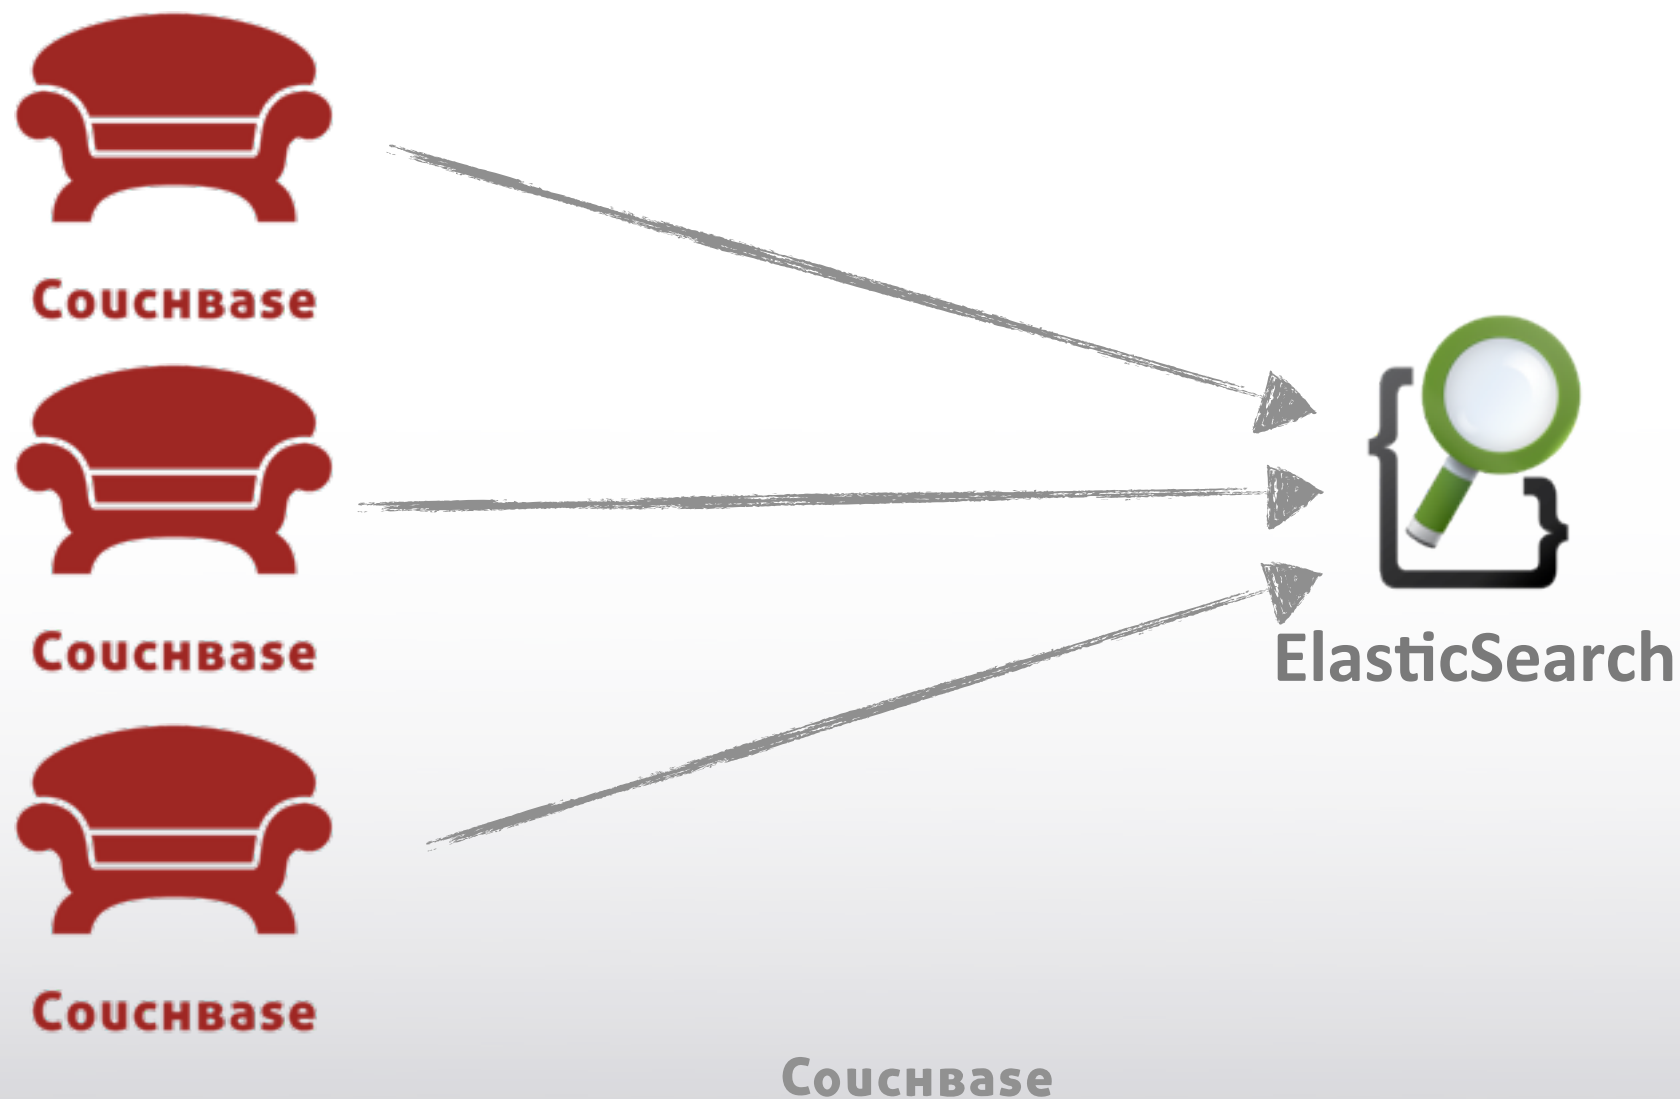
```
1 function (doc) {  
2   if (doc.geo) {  
3     emit( {type: "Point",  
4           coordinates: [doc.geo.lng, doc.geo.lat]},  
5           doc.name );  
6   }  
7 }
```



Couchbase

Elastic Search Adaptor

- Elastic Search is good for ad-hoc queries and faceted browsing
- Our adapter is aware of changing Couchbase topology
- Indexed by Elastic Search after stored to disk in Couchbase



I'm Excited to See What You Build, Q & A

Couchbase Docs

www.couchbase.com/docs/index-full.html

Couchbase Forums

www.couchbase.com/forums

IRC

#couchbase

#libcouchbase

Contact me on Twitter

@tgrall

Contact me by Email

tug@couchbase.com

Learn More About Design Patterns

CouchbaseModels.com

Setting up for Ruby on Rails

CouchbaseOnRails.com

Couchbase

Couchbase



Couchbase

Q&A

Couchbase

Couchbase



Couchbase



Couchbase