



mongoDB

Hybrid Datastore

Chris Harris

Email : charris@10gen.com

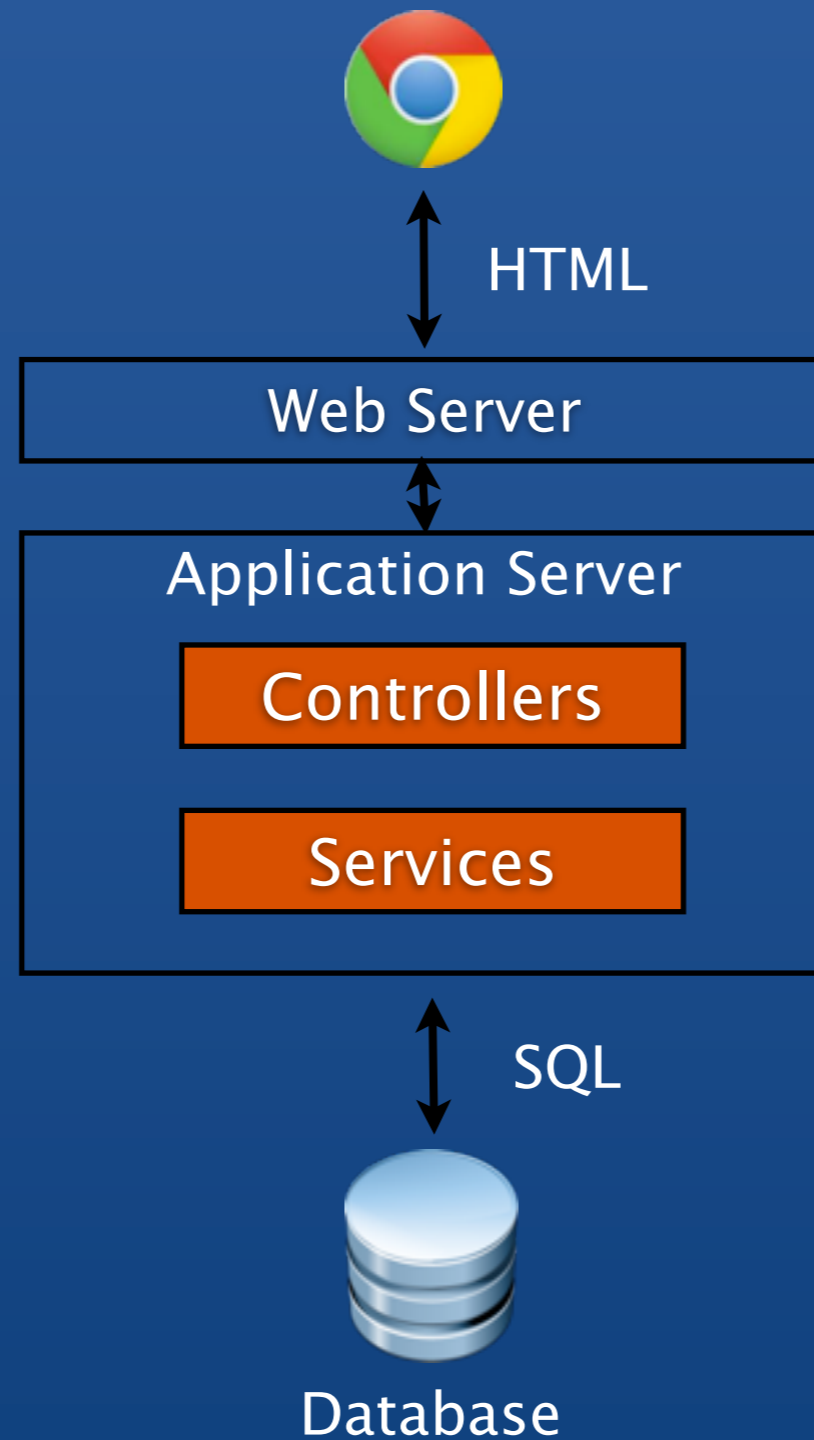
Twitter : [cj_harris5](https://twitter.com/cj_harris5)



Traditional Architecture



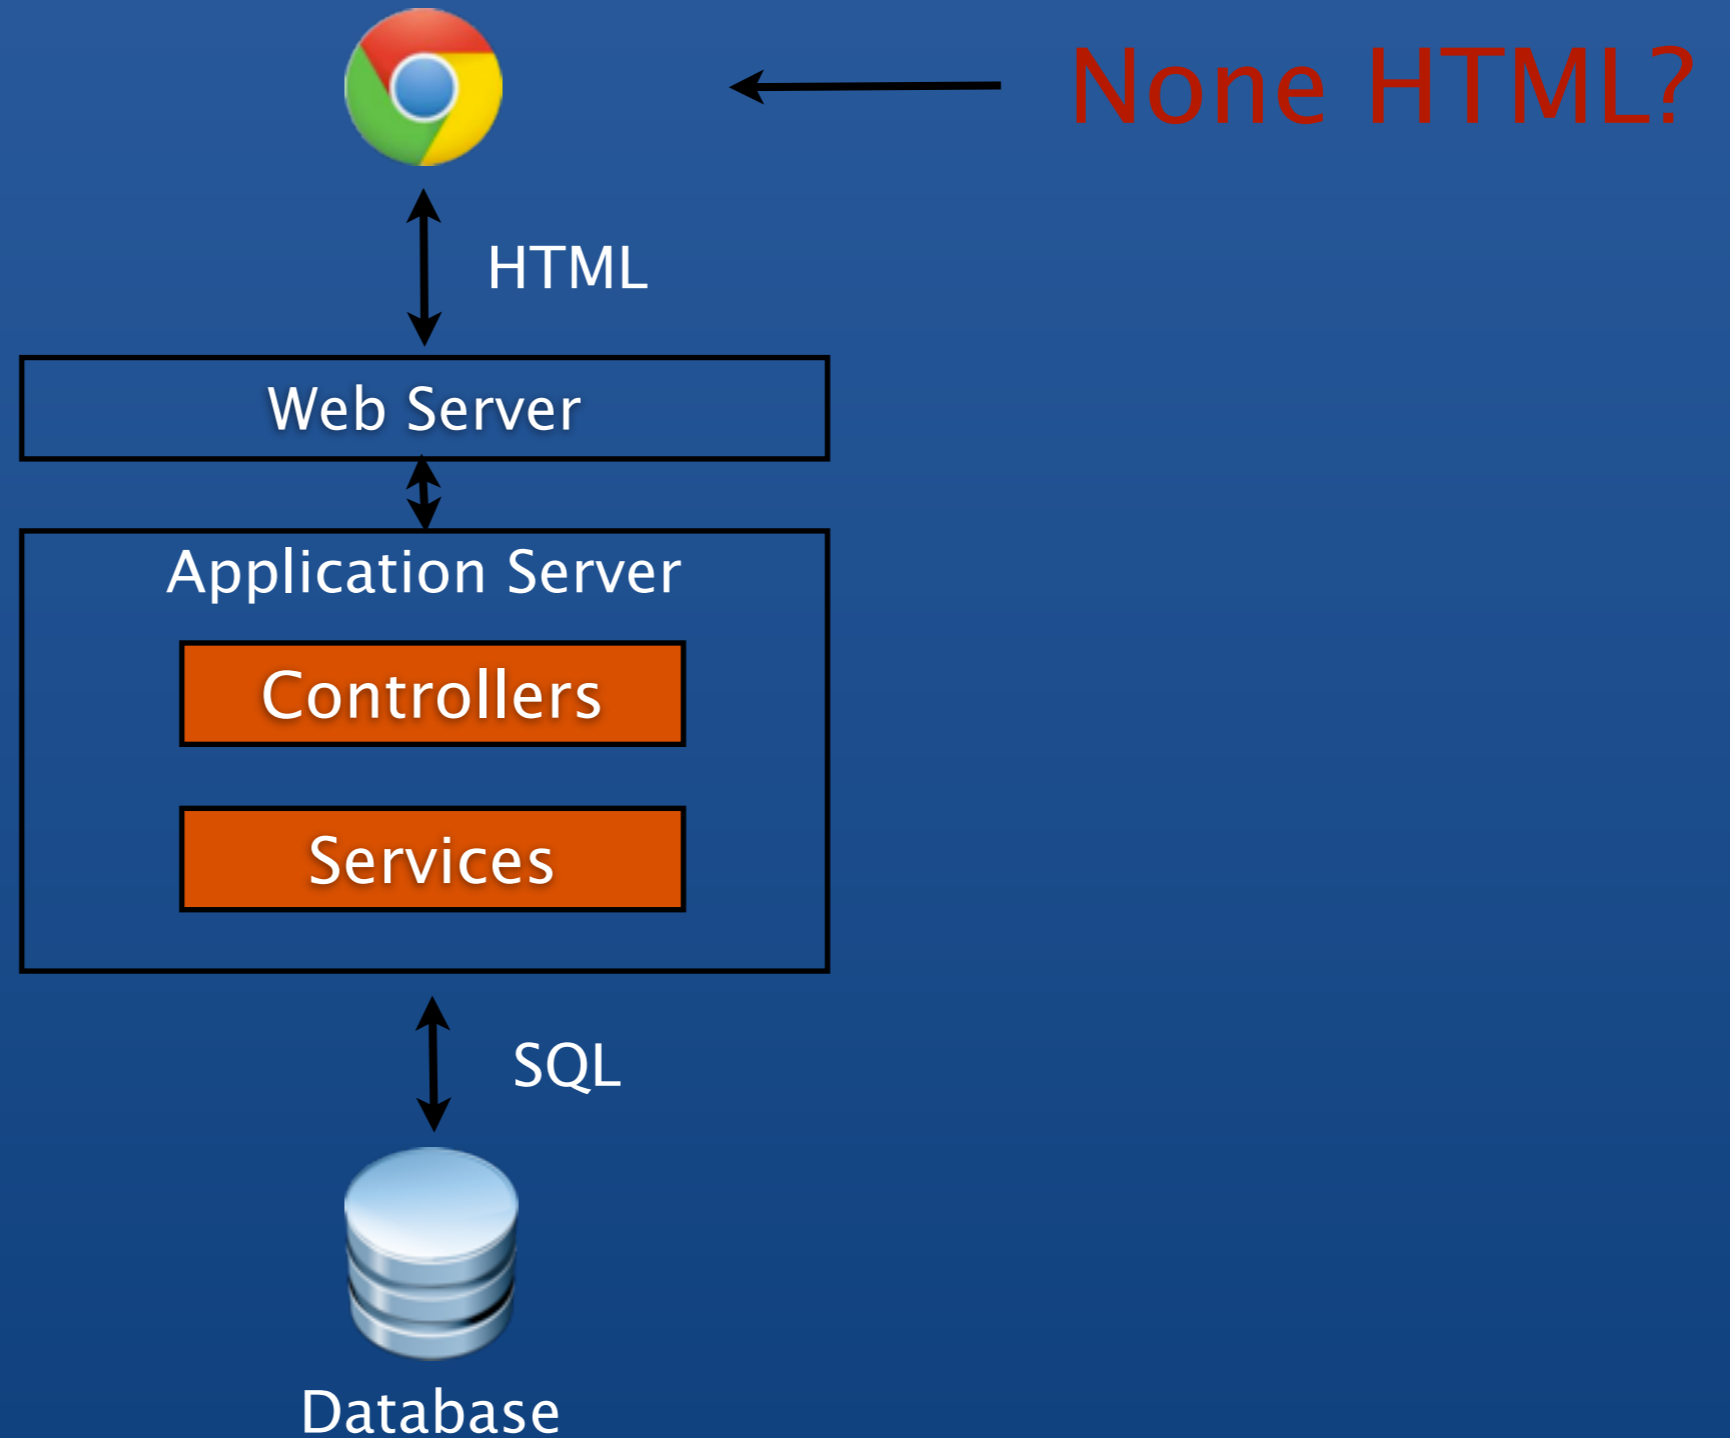
Traditional Architecture



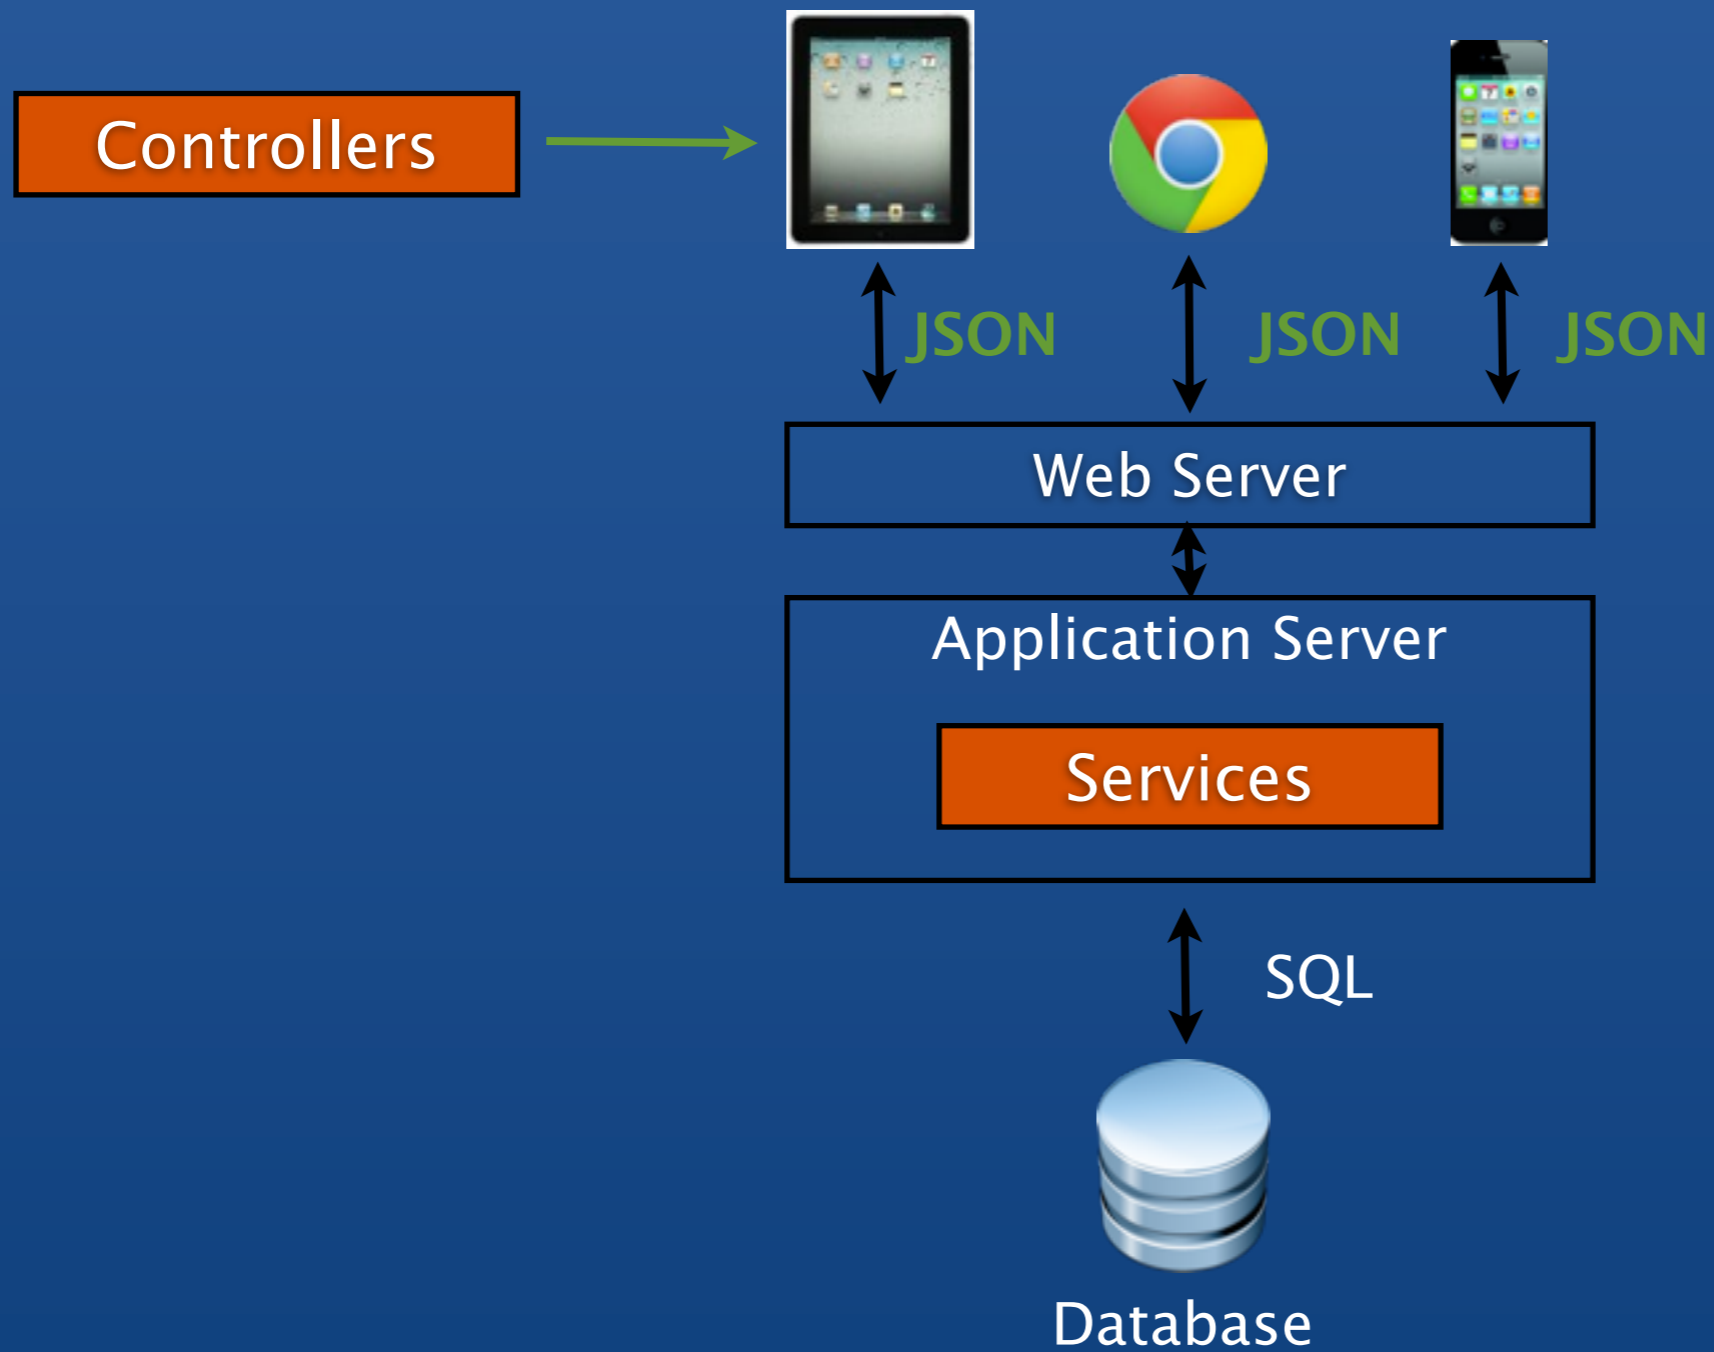
Challenge #1 - HTML 5



Multiple Client Types



HTML 5



The Move to Services

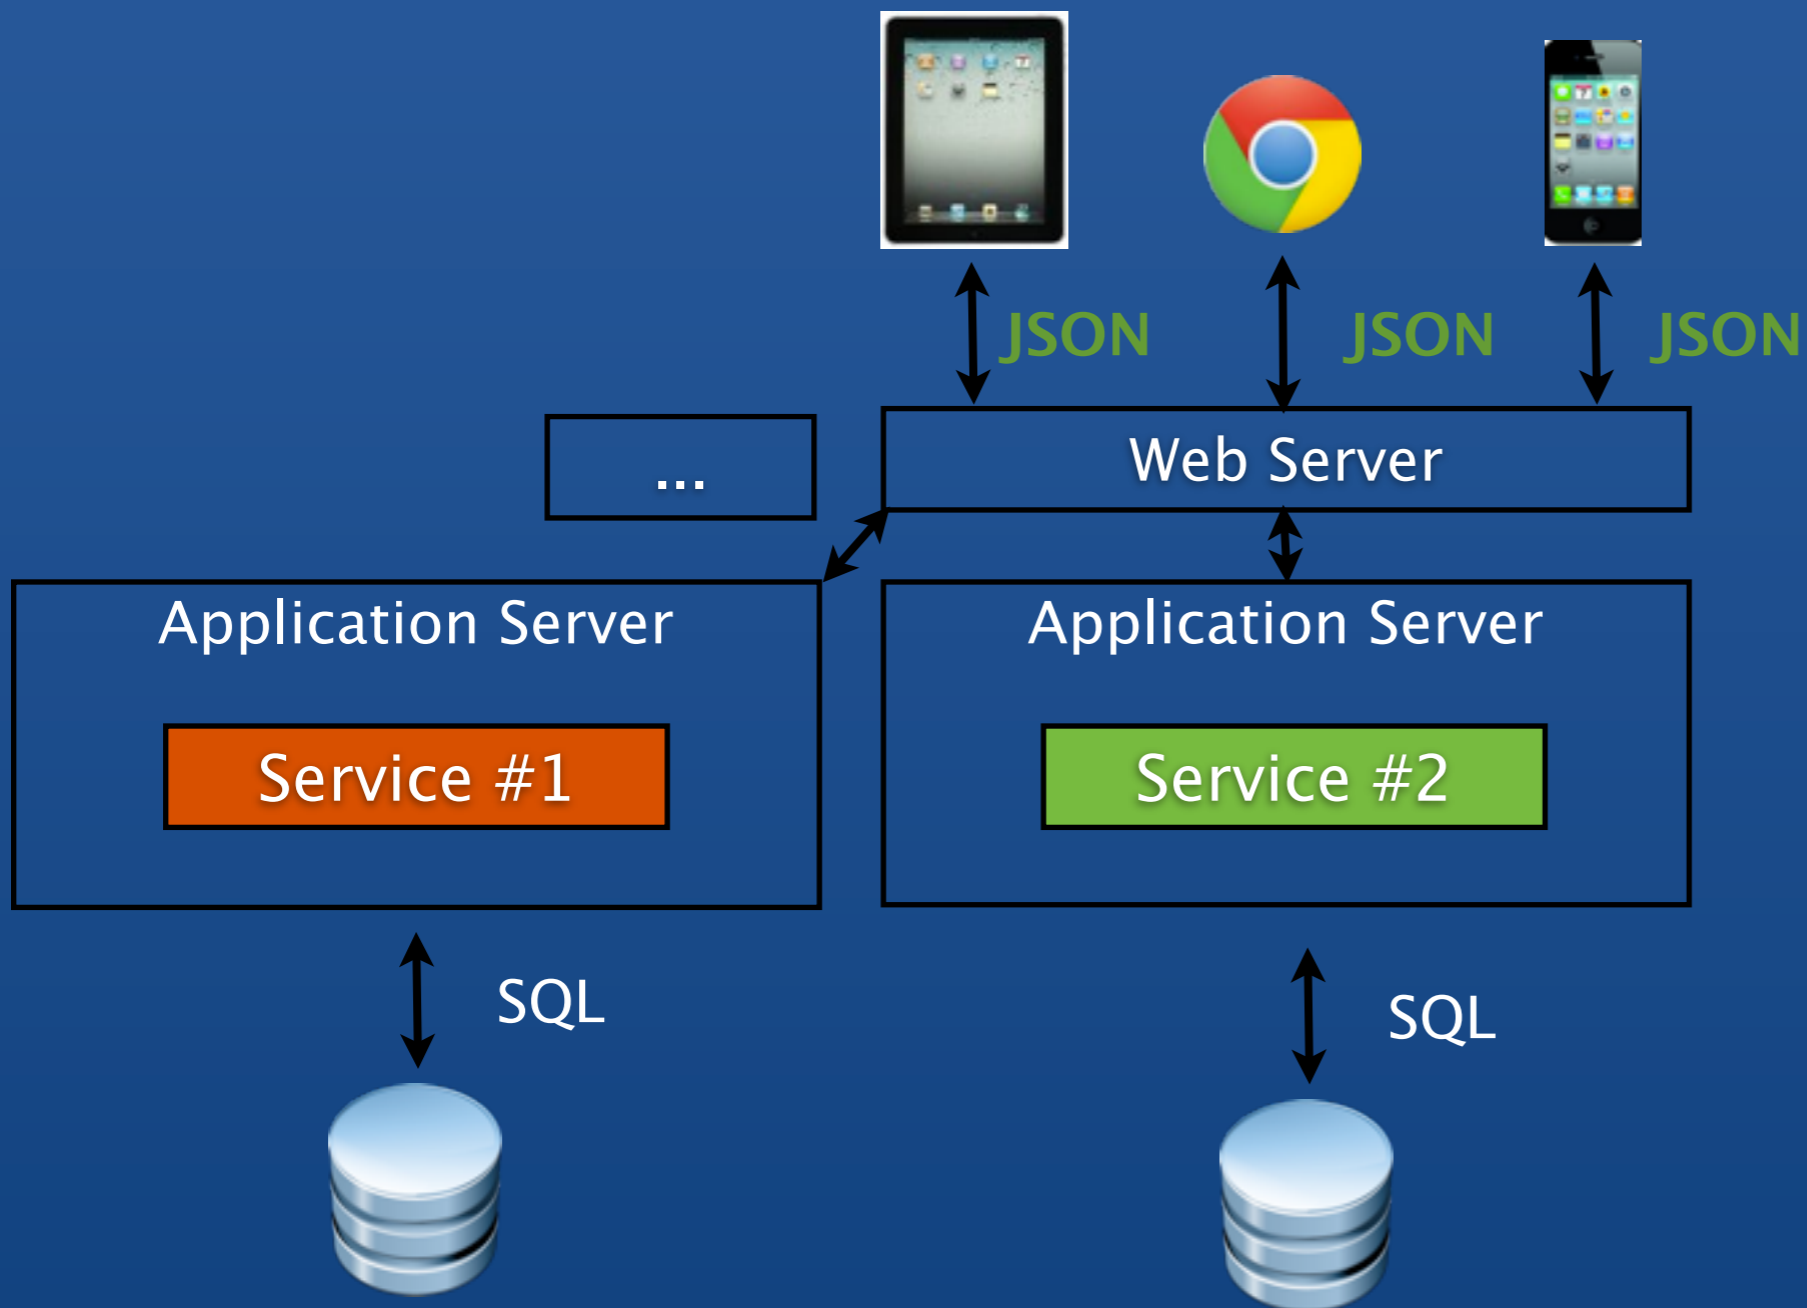
The controllers have moved to the client

Expose small JSON Services

+ Death to the monolithic deployment



Service Architecture



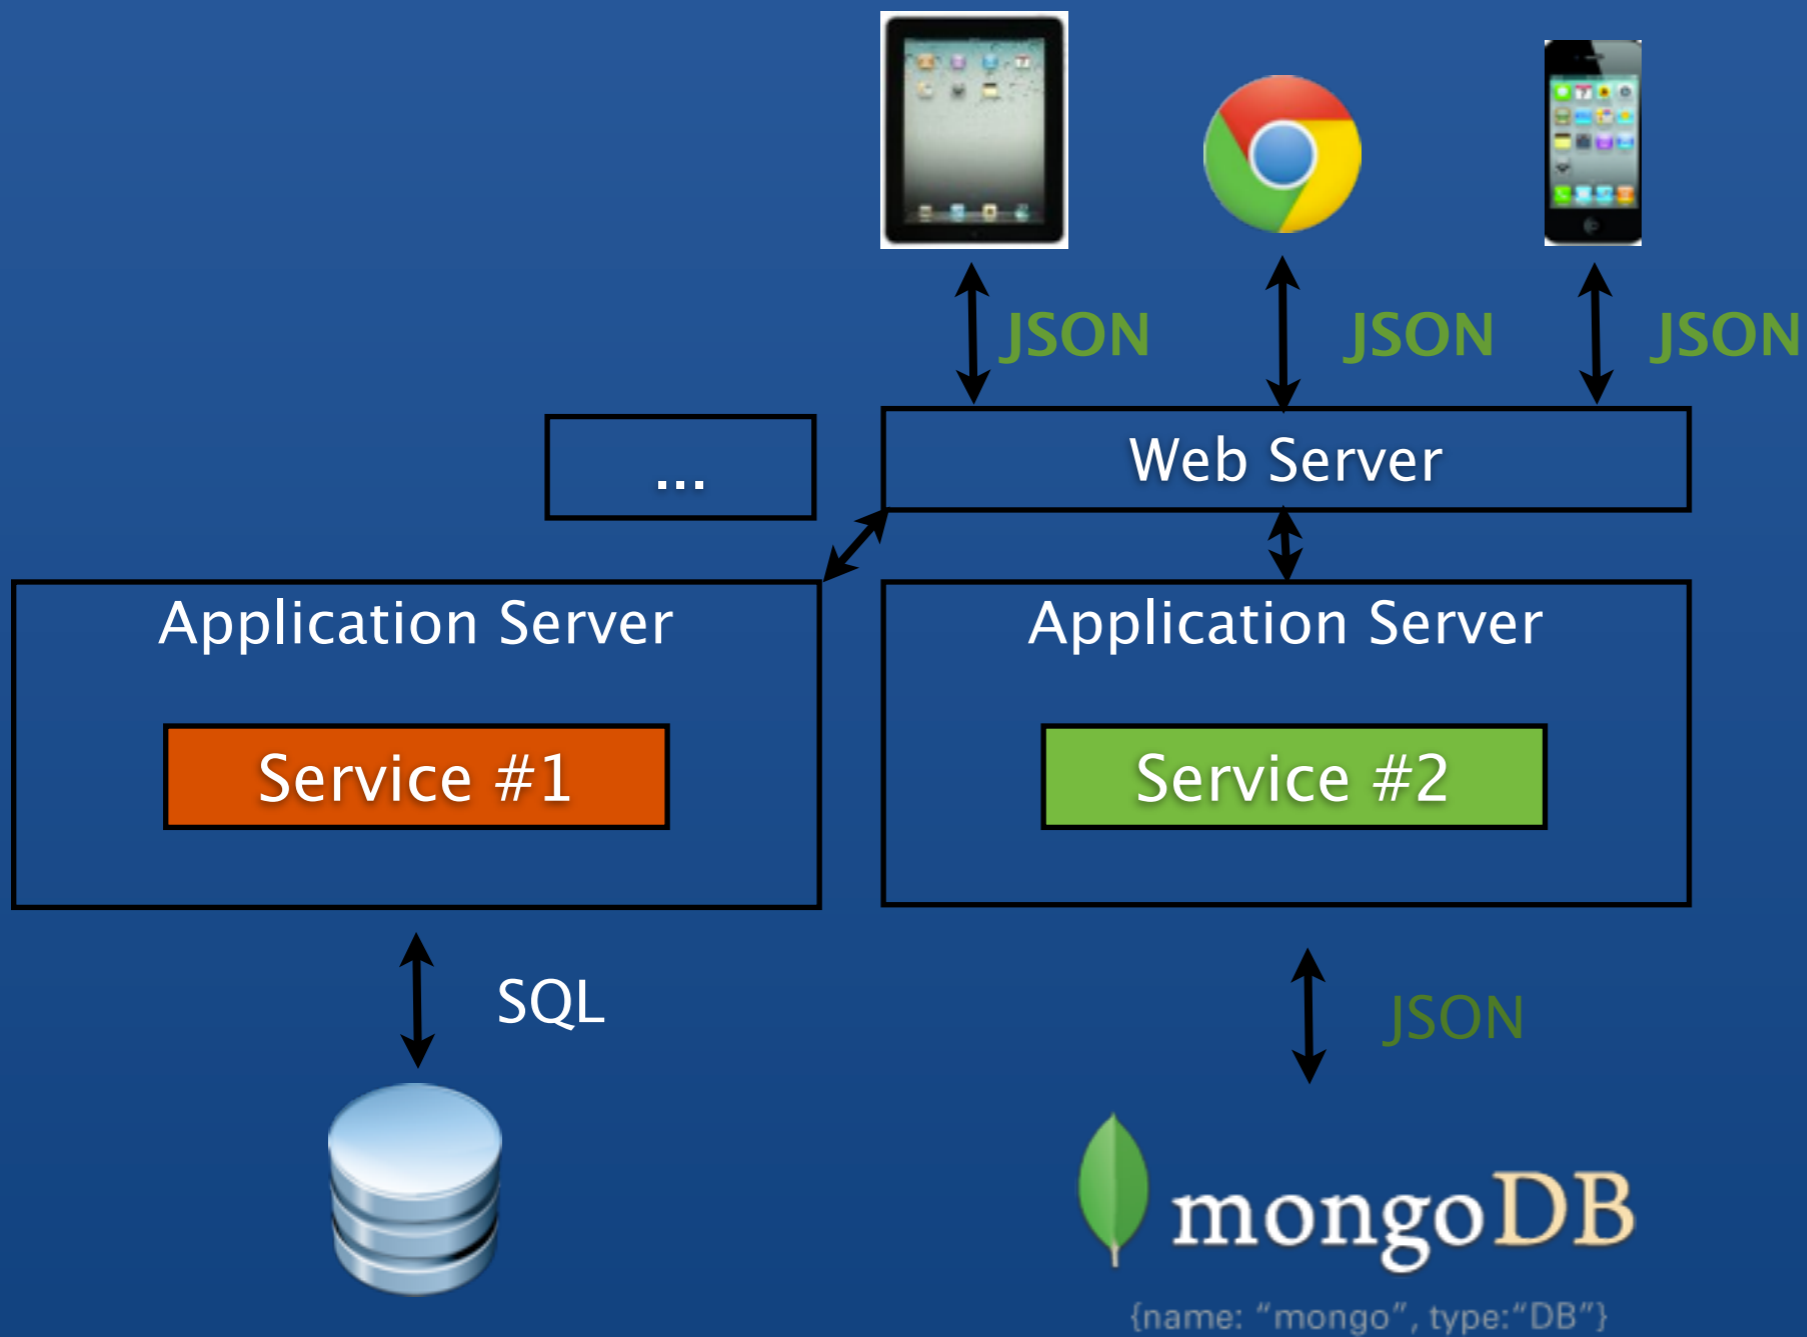
The Service becomes the API

JSON Service becomes a API to client and other applications.

The clients are not bound to the underlying data store

- Complexity due to mismatch between the JSON and SQL

Migrating to MongoDB



Here is a “simple” SQL Model

```
mysql> select * from book;
```

```
+-----+-----+
| id | title |
+-----+-----+
| 1 | The Demon-Haunted World: Science as a Candle in the Dark |
| 2 | Cosmos |
| 3 | Programming in Scala |
+-----+-----+
3 rows in set (0.00 sec)
```

```
mysql> select * from bookauthor;
```

```
+-----+-----+
| book_id | author_id |
+-----+-----+
| 1 | 1 |
| 2 | 1 |
| 3 | 2 |
| 3 | 3 |
| 3 | 4 |
+-----+-----+
5 rows in set (0.00 sec)
```

```
mysql> select * from author;
```

```
+-----+-----+-----+-----+-----+-----+
| id | last_name | first_name | middle_name | nationality | year_of_birth |
+-----+-----+-----+-----+-----+-----+
| 1 | Sagan | Carl | Edward | NULL | 1934 |
| 2 | Odersky | Martin | NULL | DE | 1958 |
| 3 | Spoon | Lex | NULL | NULL | NULL |
| 4 | Venners | Bill | NULL | NULL | NULL |
+-----+-----+-----+-----+-----+-----+
4 rows in set (0.00 sec)
```



The Same Data in MongoDB

```
{
  "_id" : ObjectId("4dfa6baa9c65dae09a4bbda5"),
  "title" : "Programming in Scala",
  "author" : [
    {
      "first_name" : "Martin",
      "last_name" : "Odersky",
      "nationality" : "DE",
      "year_of_birth" : 1958
    },
    {
      "first_name" : "Lex",
      "last_name" : "Spoon"
    },
    {
      "first_name" : "Bill",
      "last_name" : "Venners"
    }
  ]
}
```


Aggregate Example – twitter

```
db.tweets.aggregate(  
  {$match:  
    {"user.friends_count": { $gt: 0 },  
     "user.followers_count": { $gt: 0 }  
  },  
  {$project:  
    { location: "$user.location",  
      friends: "$user.friends_count",  
      followers: "$user.followers_count"  
    }  
  },  
  {$group:  
    { _id: "$location",  
      friends: { $sum: "$friends" },  
      followers: { $sum: "$followers" }  
    }  
  }  
);
```

Predicate

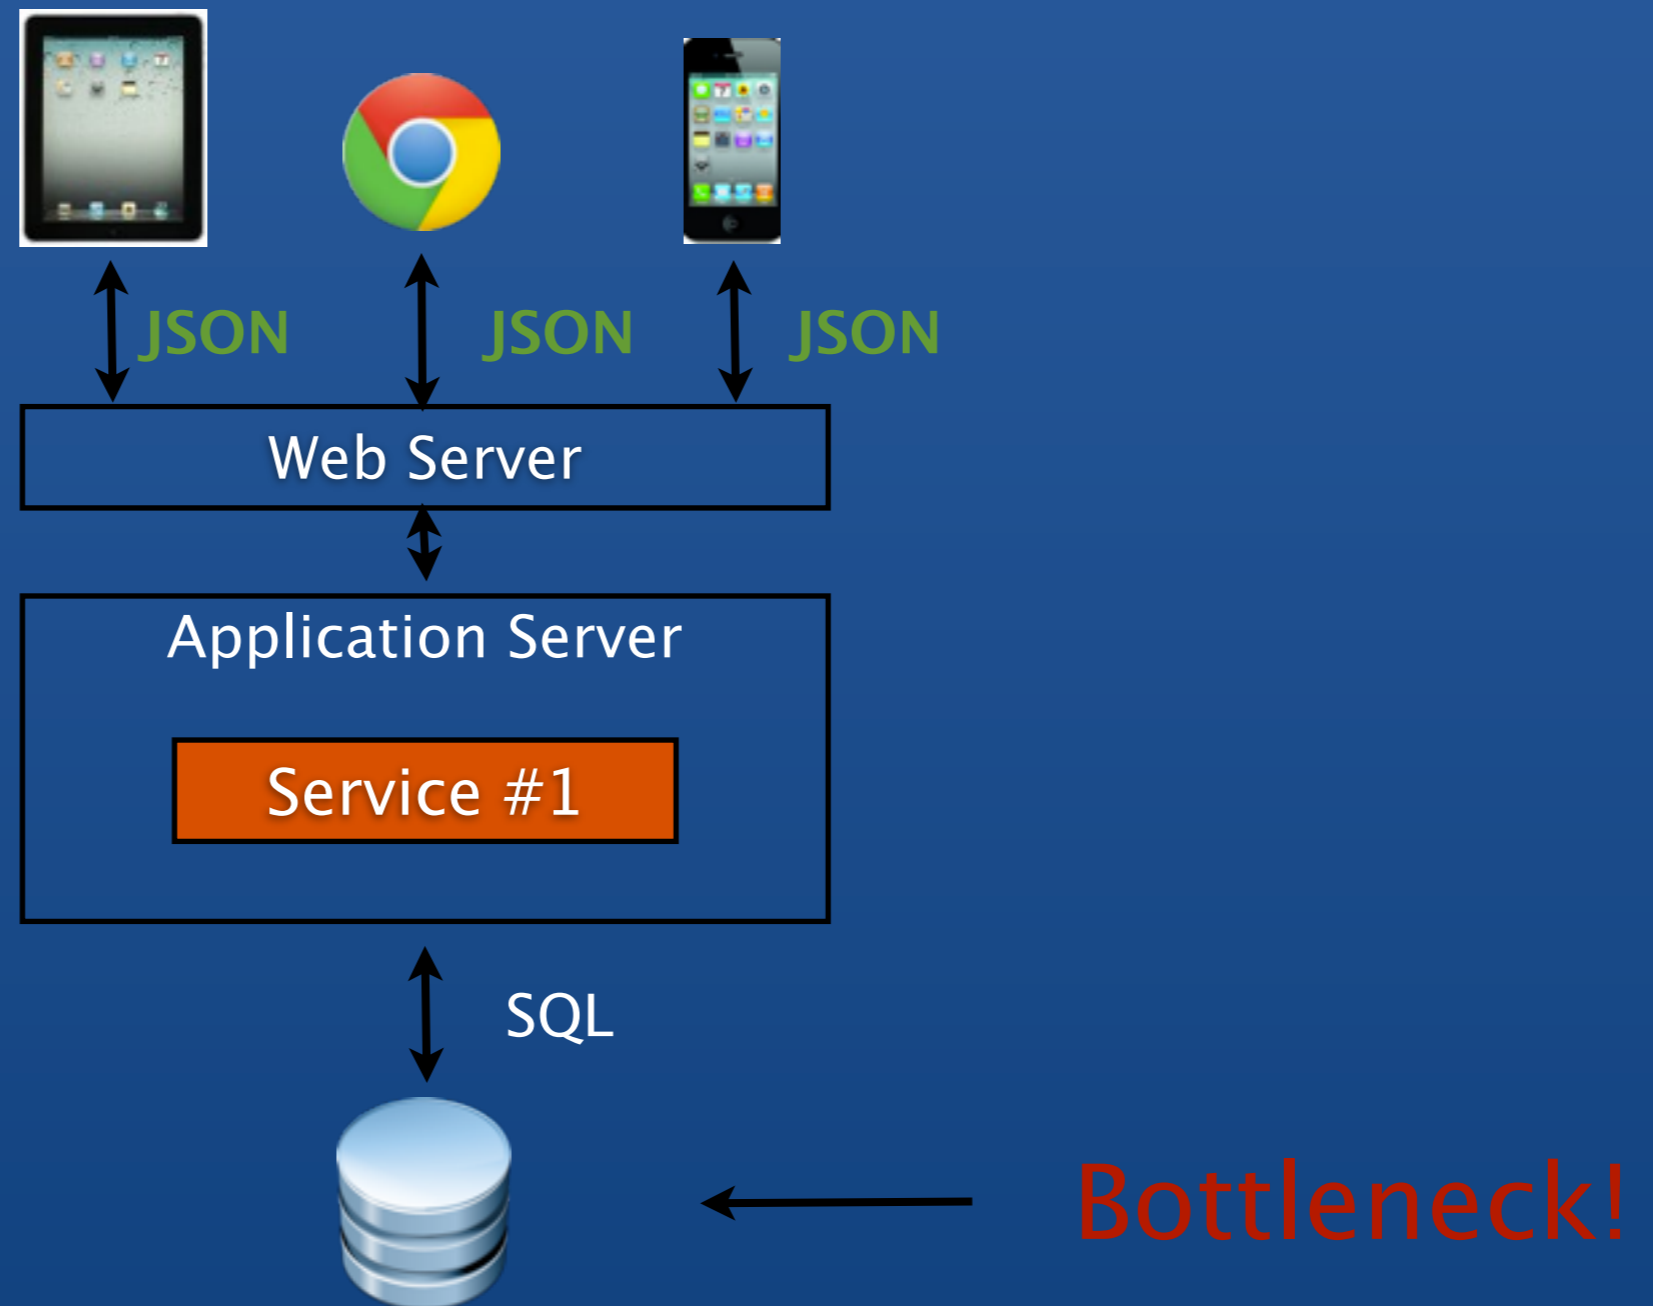
Parts of the document you want to project

Function to apply to the result set

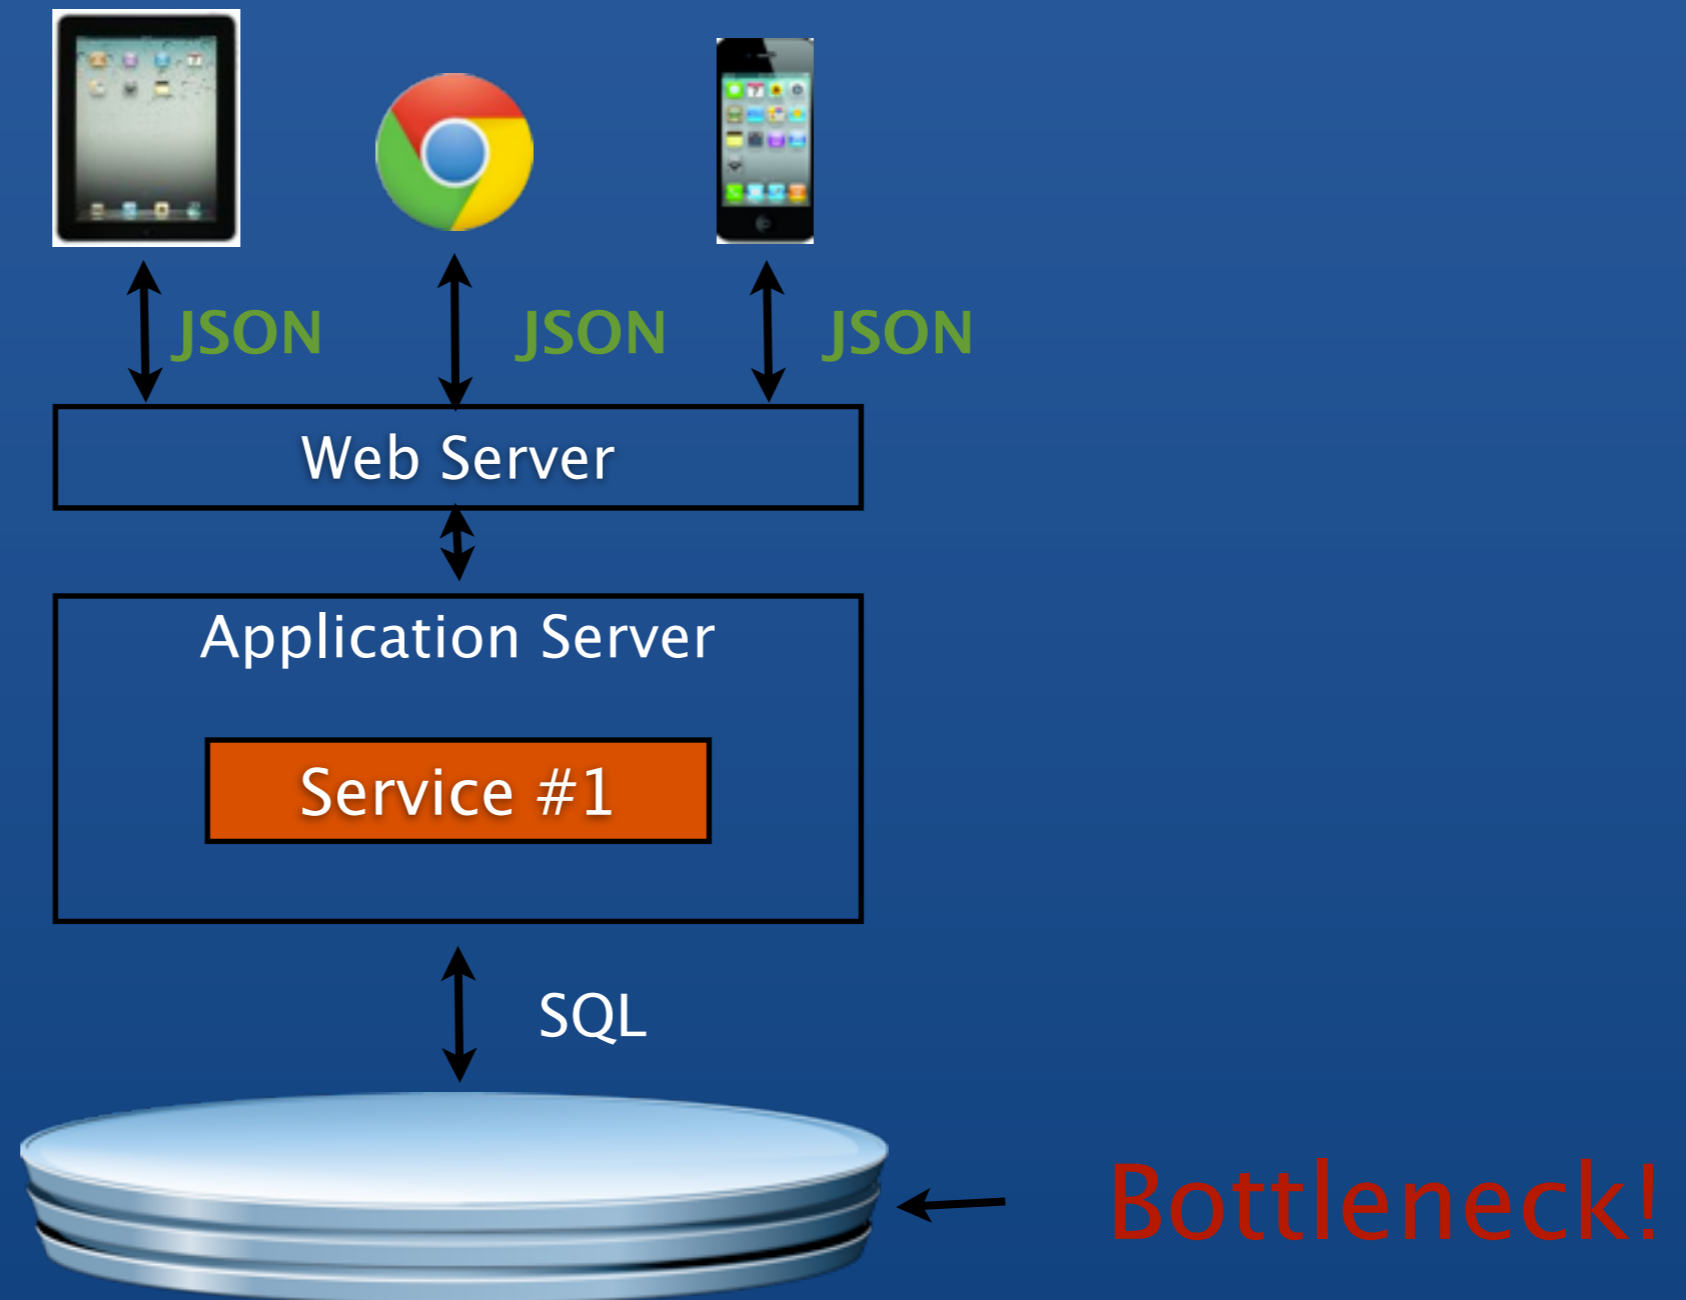
Challenge #2 - Write Volumes



Need to Scale Datasource

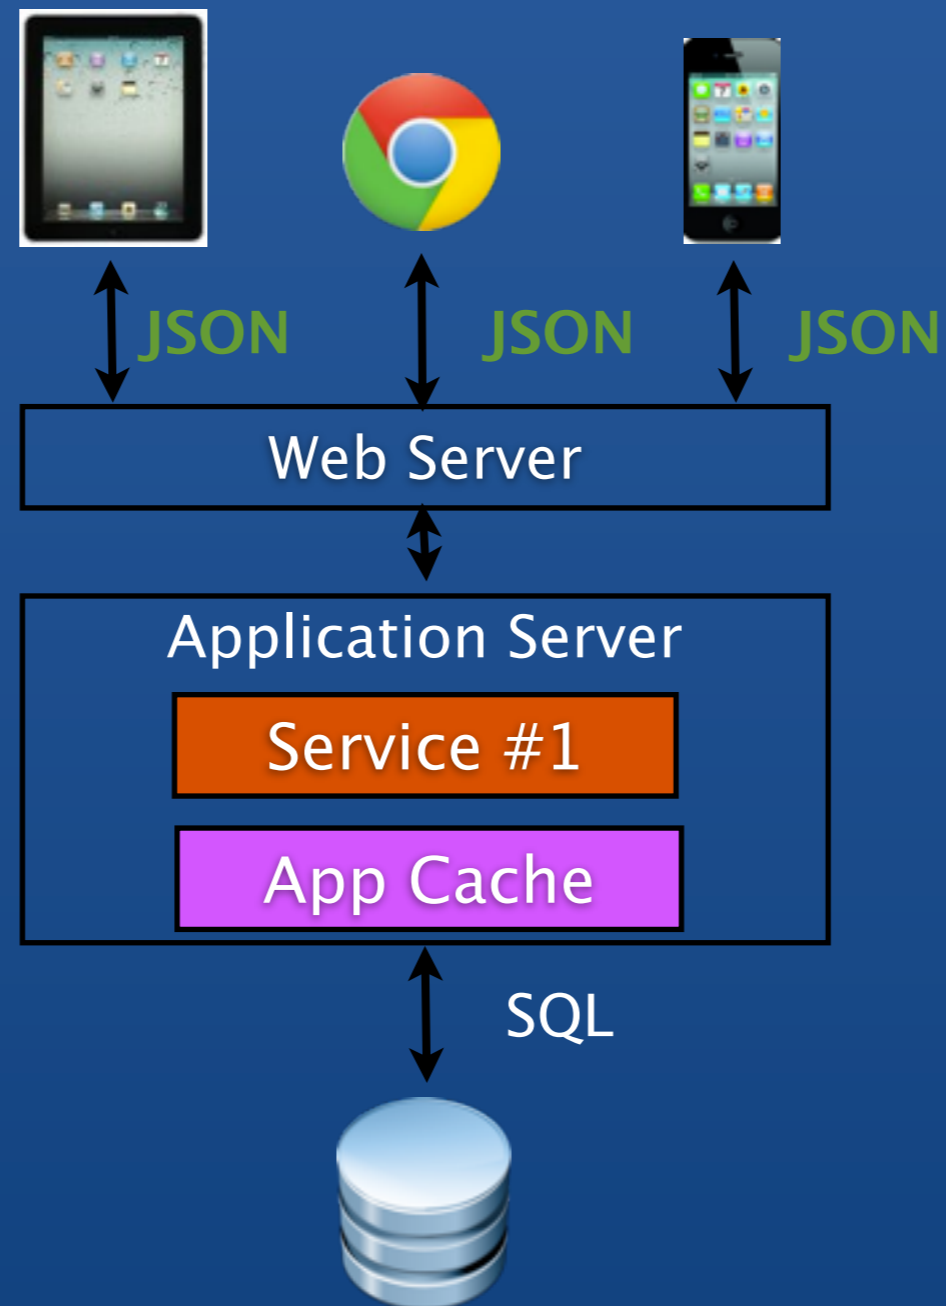


Need to Scale Datasource



10gen  mongoDB

Application Cache?



Issues

- + Read Only data comes from a Cache
- Writes slow down as need to update the Cache and the Database
- Need to keep cache data in sync between Application Servers



http://community.qlikview.com/cfs-filesystemfile.ashx/___key/CommunityServer.Blogs.Components.WeblogFiles/theqlikviewblog/Cutting-Grass-with-Scissors-_2D00_-2.jpg



http://www.bitquill.net/blog/wp-content/uploads/2008/07/pack_of_harvesters.jpg

Big Data at a Glance

Large Dataset
Primary Key as "username"

Big Data at a Glance

Large Dataset
Primary Key as “username”

- Systems like Google File System (which inspired Hadoop’s HDFS) and MongoDB’s Sharding handle the scale problem by chunking

Big Data at a Glance

Large Dataset
Primary Key as “username”

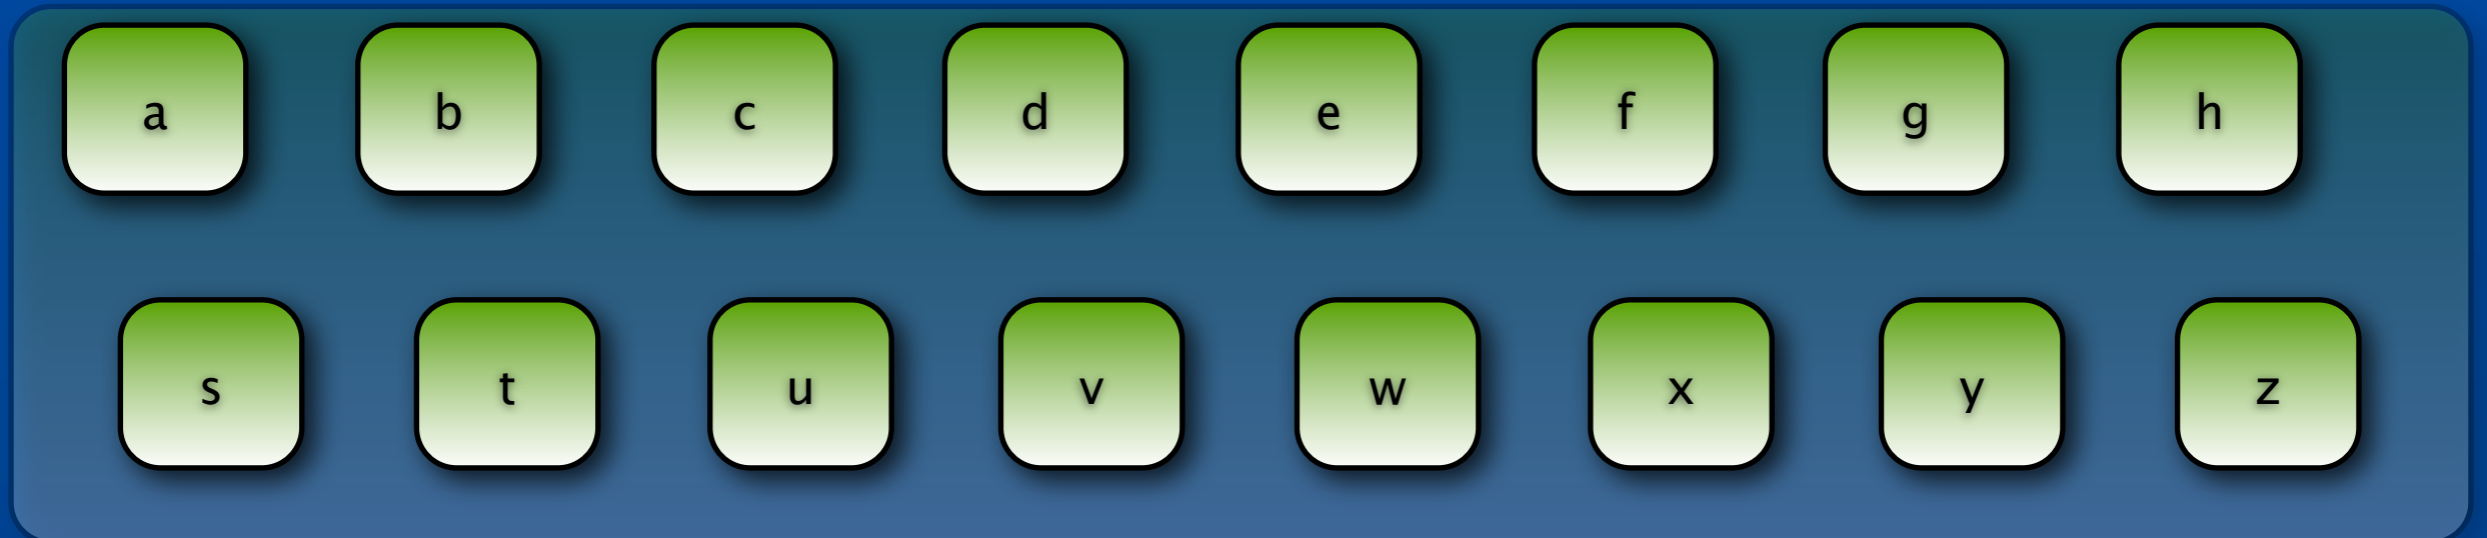
- Systems like Google File System (which inspired Hadoop’s HDFS) and MongoDB’s Sharding handle the scale problem by chunking
- Break up pieces of data into smaller chunks, spread across many data nodes

Big Data at a Glance

Large Dataset
Primary Key as “username”

- Systems like Google File System (which inspired Hadoop’s HDFS) and MongoDB’s Sharding handle the scale problem by chunking
- Break up pieces of data into smaller chunks, spread across many data nodes
 - Each data node contains many chunks

Big Data at a Glance




- Systems like Google File System (which inspired Hadoop's HDFS) and MongoDB's Sharding handle the scale problem by chunking
- Break up pieces of data into smaller chunks, spread across many data nodes
 - Each data node contains many chunks
 - If a chunk gets too large or a node overloaded, data can be rebalanced

Big Data at a Glance



Big Data at a Glance


Large Dataset
Primary Key as "username"



MongoDB Sharding (as well as HDFS) breaks data into chunks (~64 mb)



Scaling

Large Dataset
Primary Key as "username"



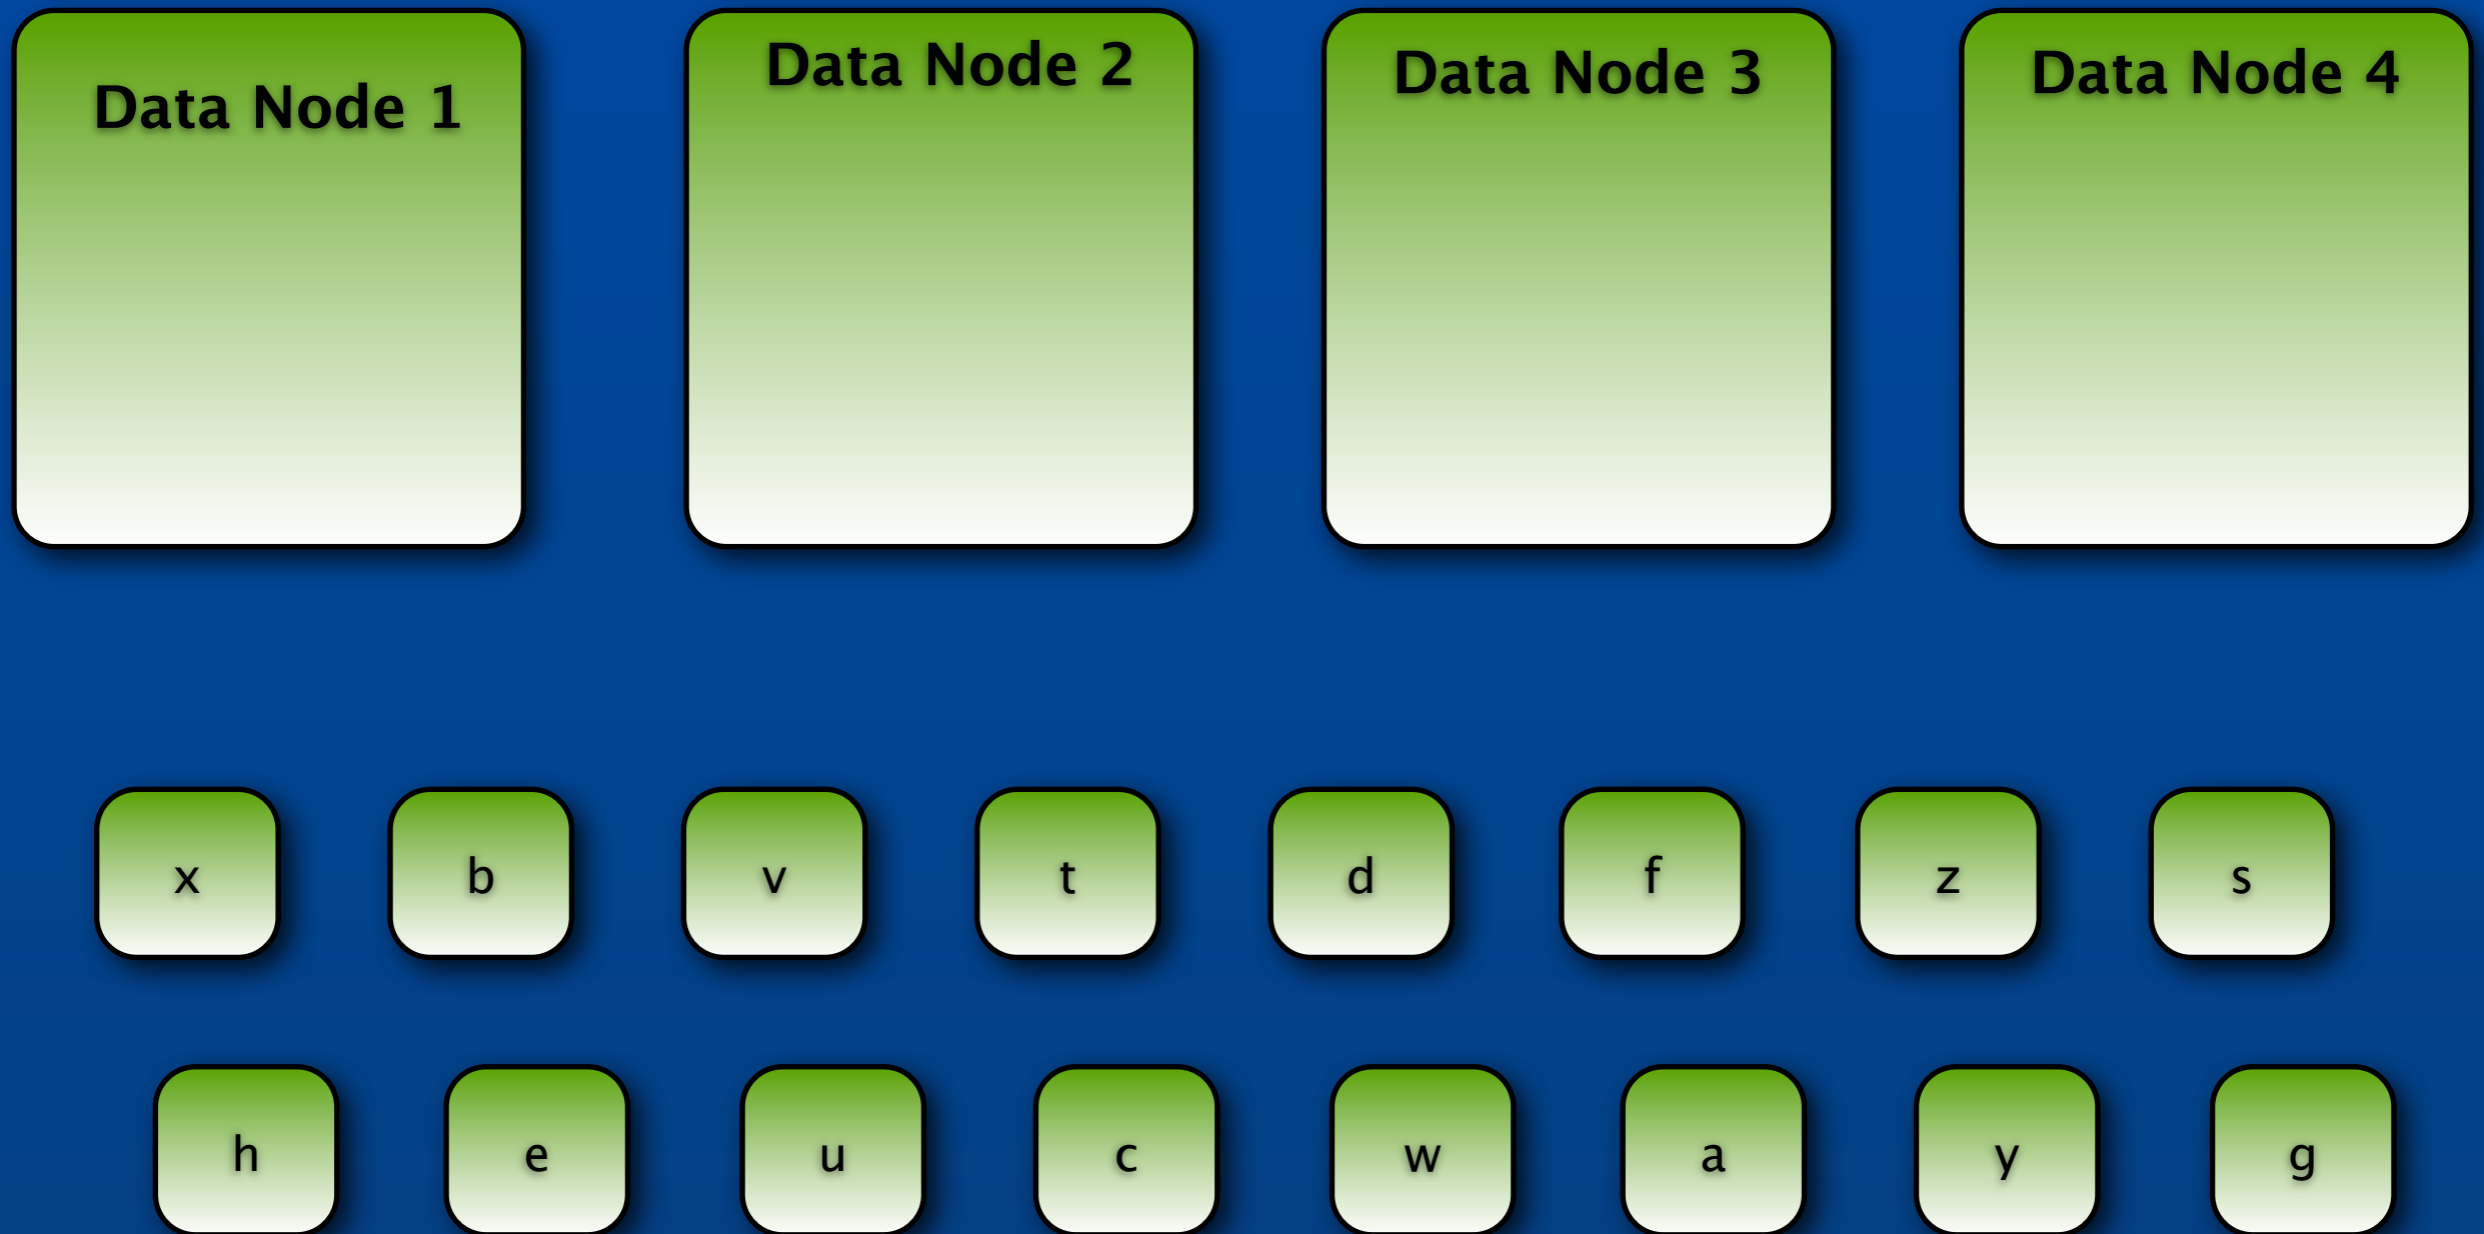
Representing data as chunks allows many levels of scale across n data nodes

Scaling




Representing data as chunks allows many levels of scale across n data nodes

Scaling



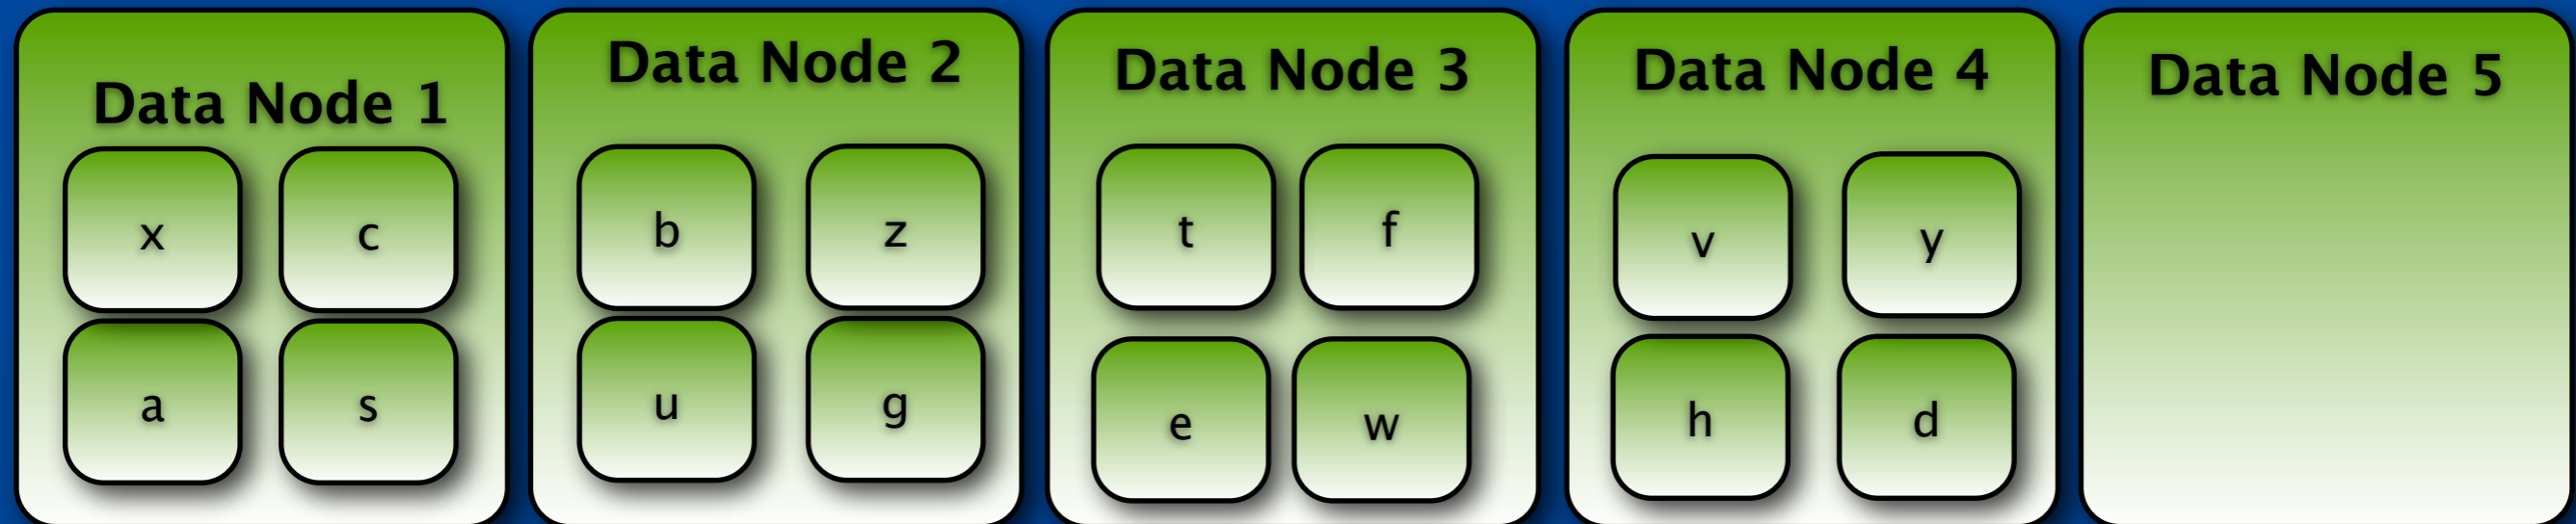
The set of chunks can be evenly distributed across n data nodes

Scaling




The set of chunks can be evenly distributed across n data nodes

Add Nodes: Chunk Rebalancing



The goal is equilibrium – an equal distribution.
As nodes are added (or even removed)
chunks can be redistributed for balance.

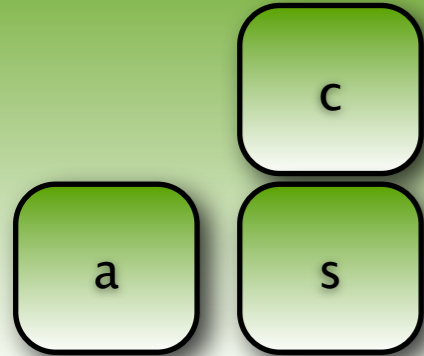
Add Nodes: Chunk Rebalancing




The goal is equilibrium – an equal distribution.
As nodes are added (or even removed)
chunks can be redistributed for balance.

Writes Routed to Appropriate Chunk

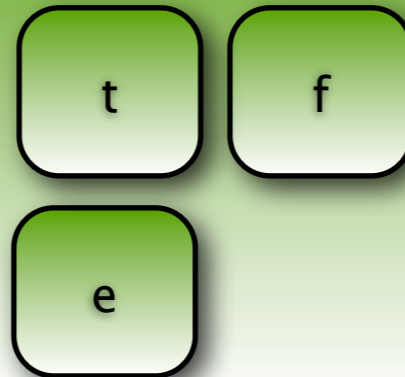
Data Node 1




Data Node 2




Data Node 3



Data Node 4




Data Node 5




Writes Routed to Appropriate Chunk

Data Node 1



Data Node 2



Data Node 3




Data Node 4



Data Node 5



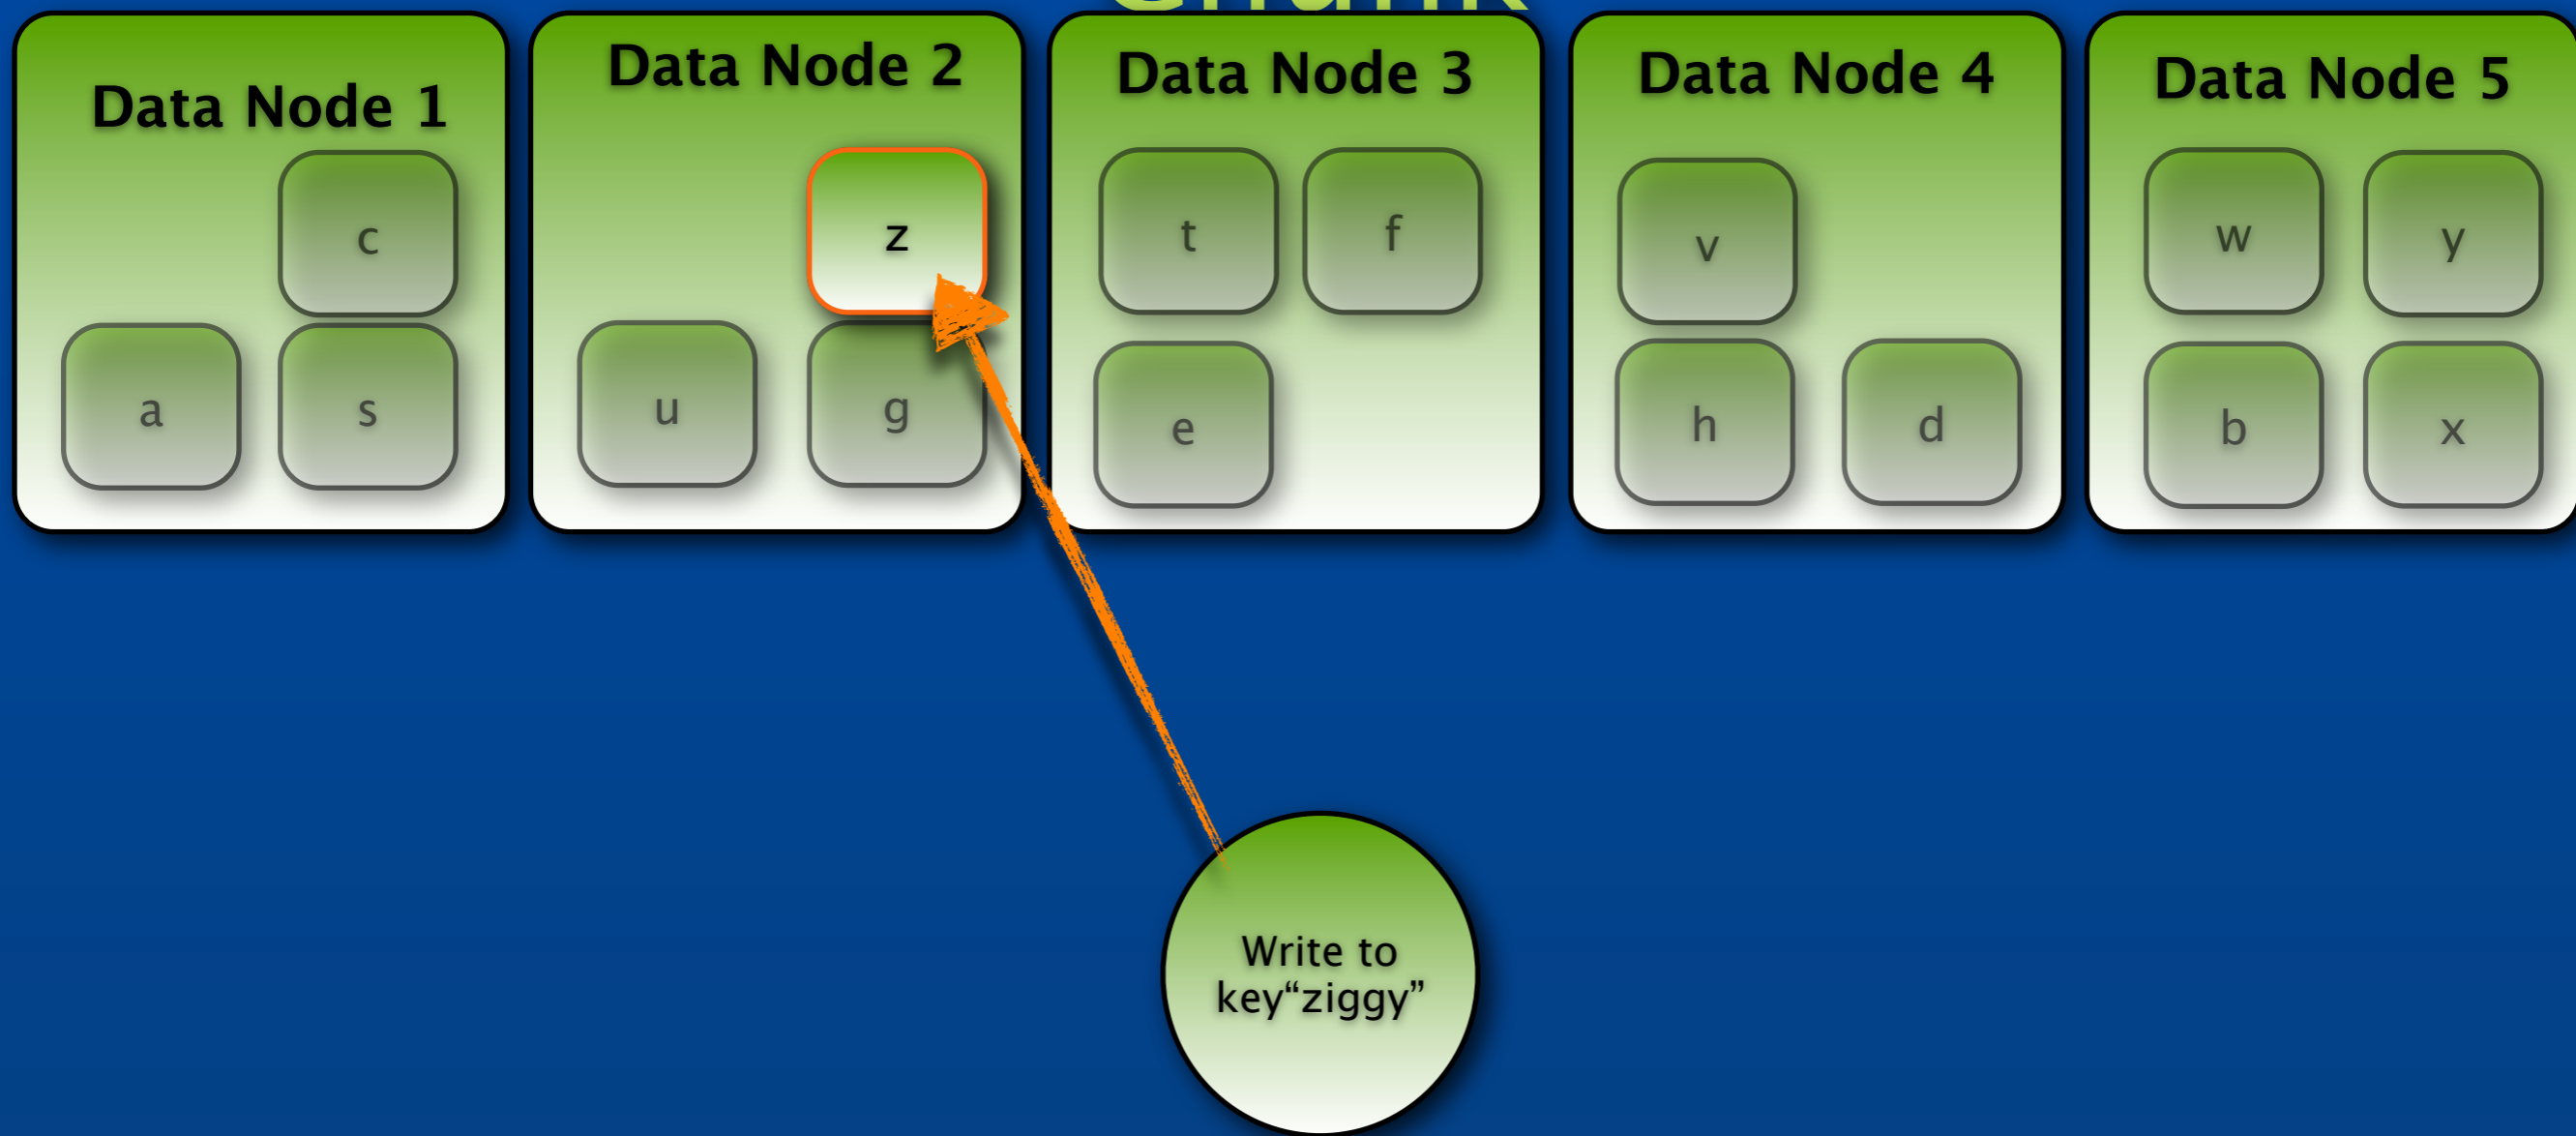
Writes Routed to Appropriate Chunk



Write to
key "ziggy"

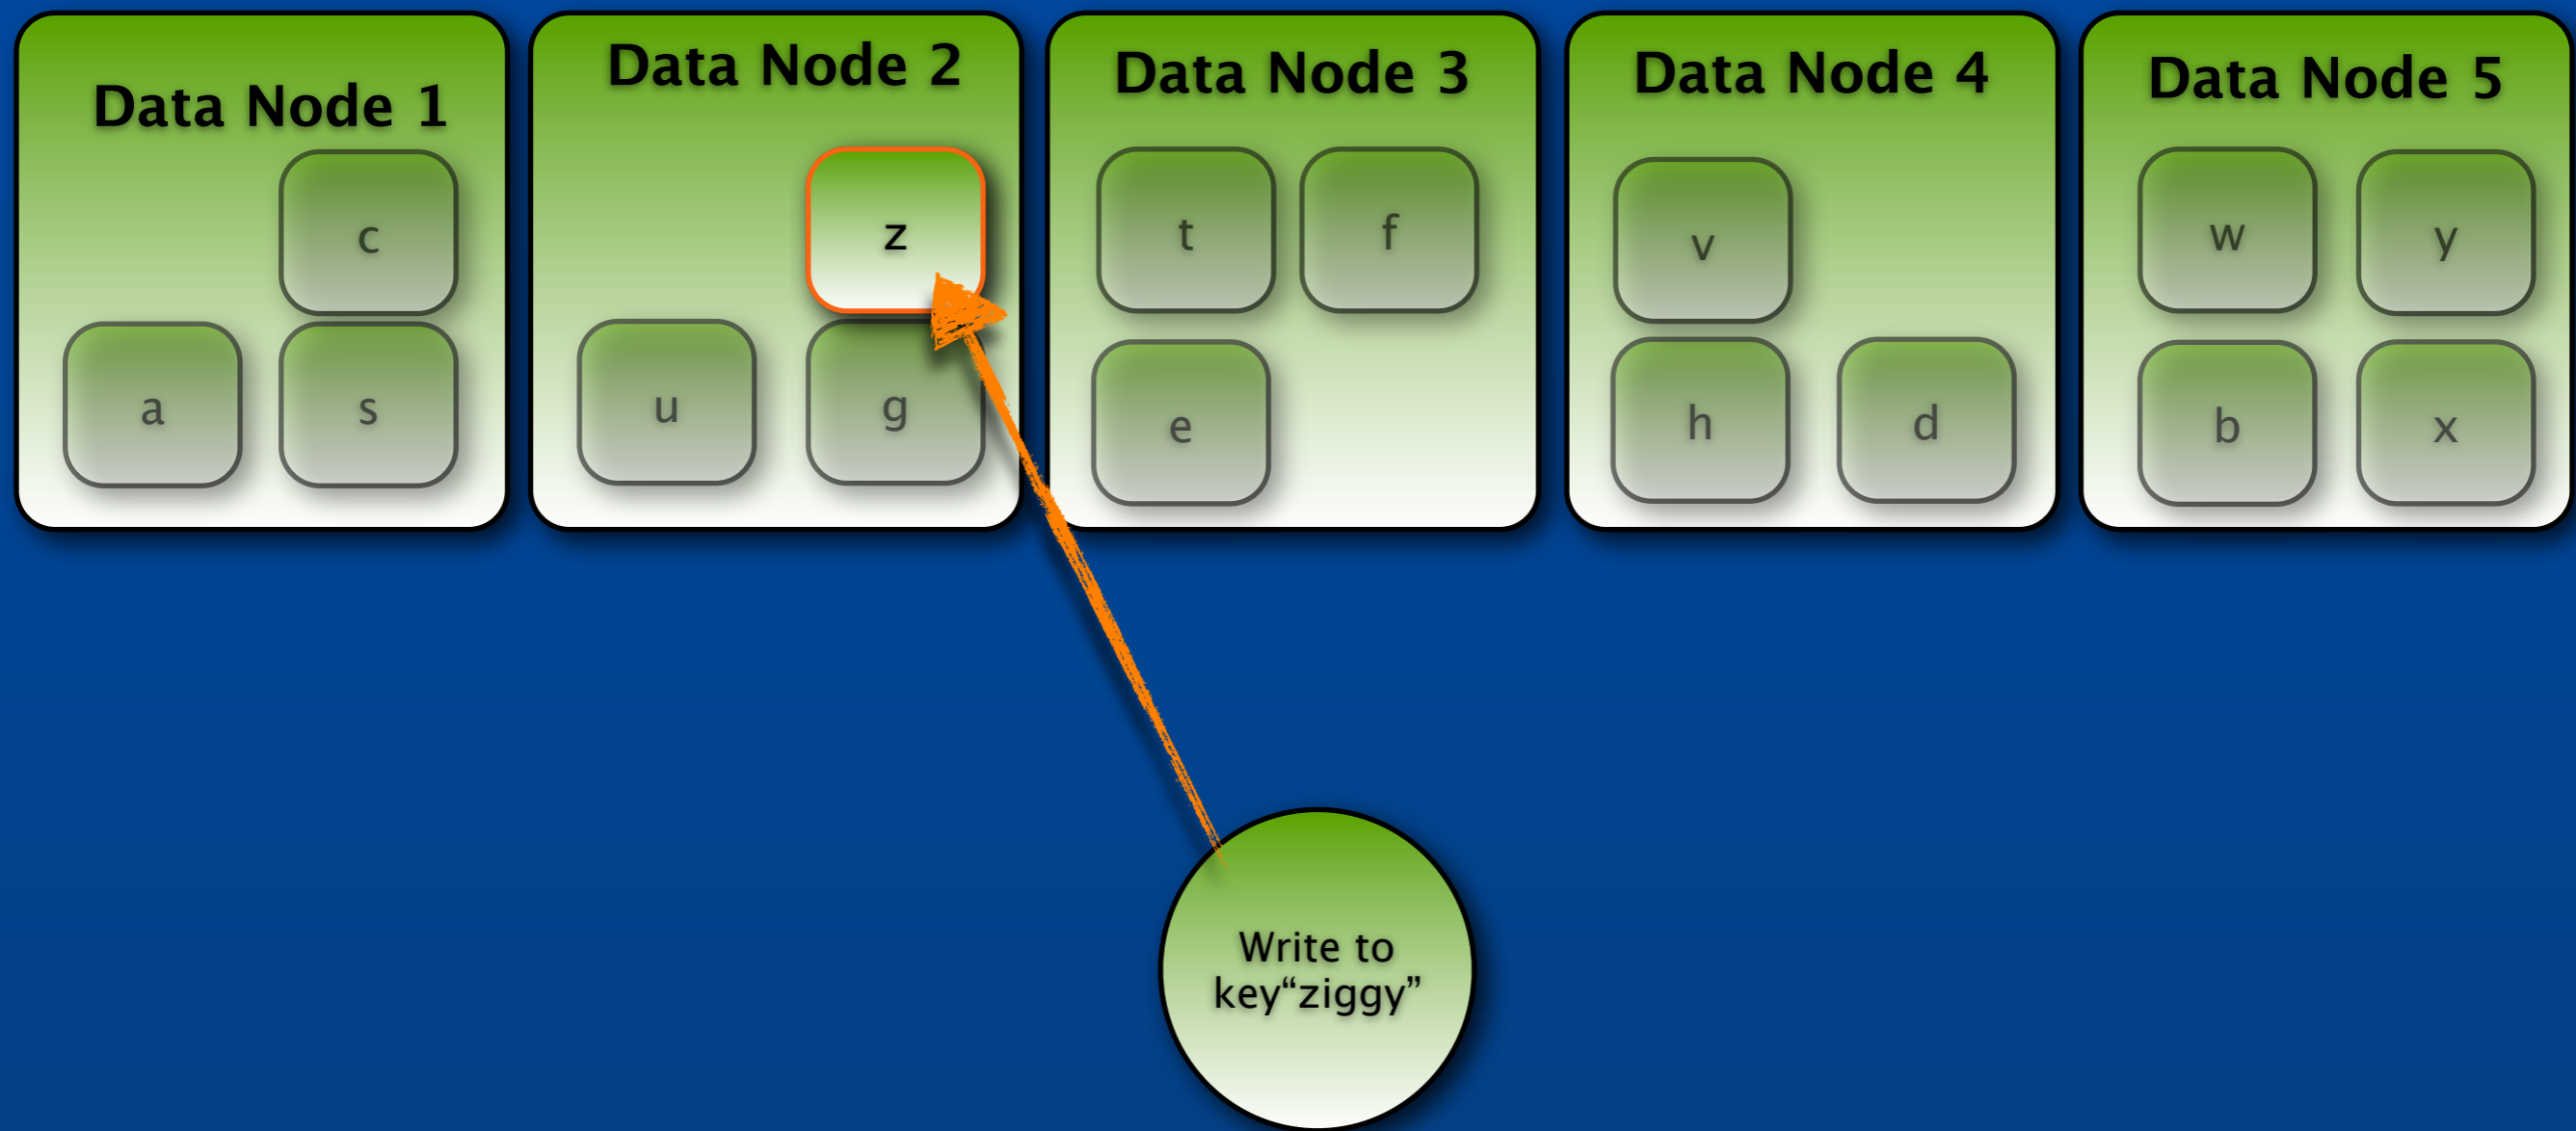
Writes are efficiently routed to the appropriate node & chunk

Writes Routed to Appropriate Chunk




Writes are efficiently routed to the appropriate node & chunk

Chunk Splitting & Balancing




If a chunk gets too large (default in MongoDB – 64mb per chunk),
It is split into two new chunks

Chunk Splitting & Balancing



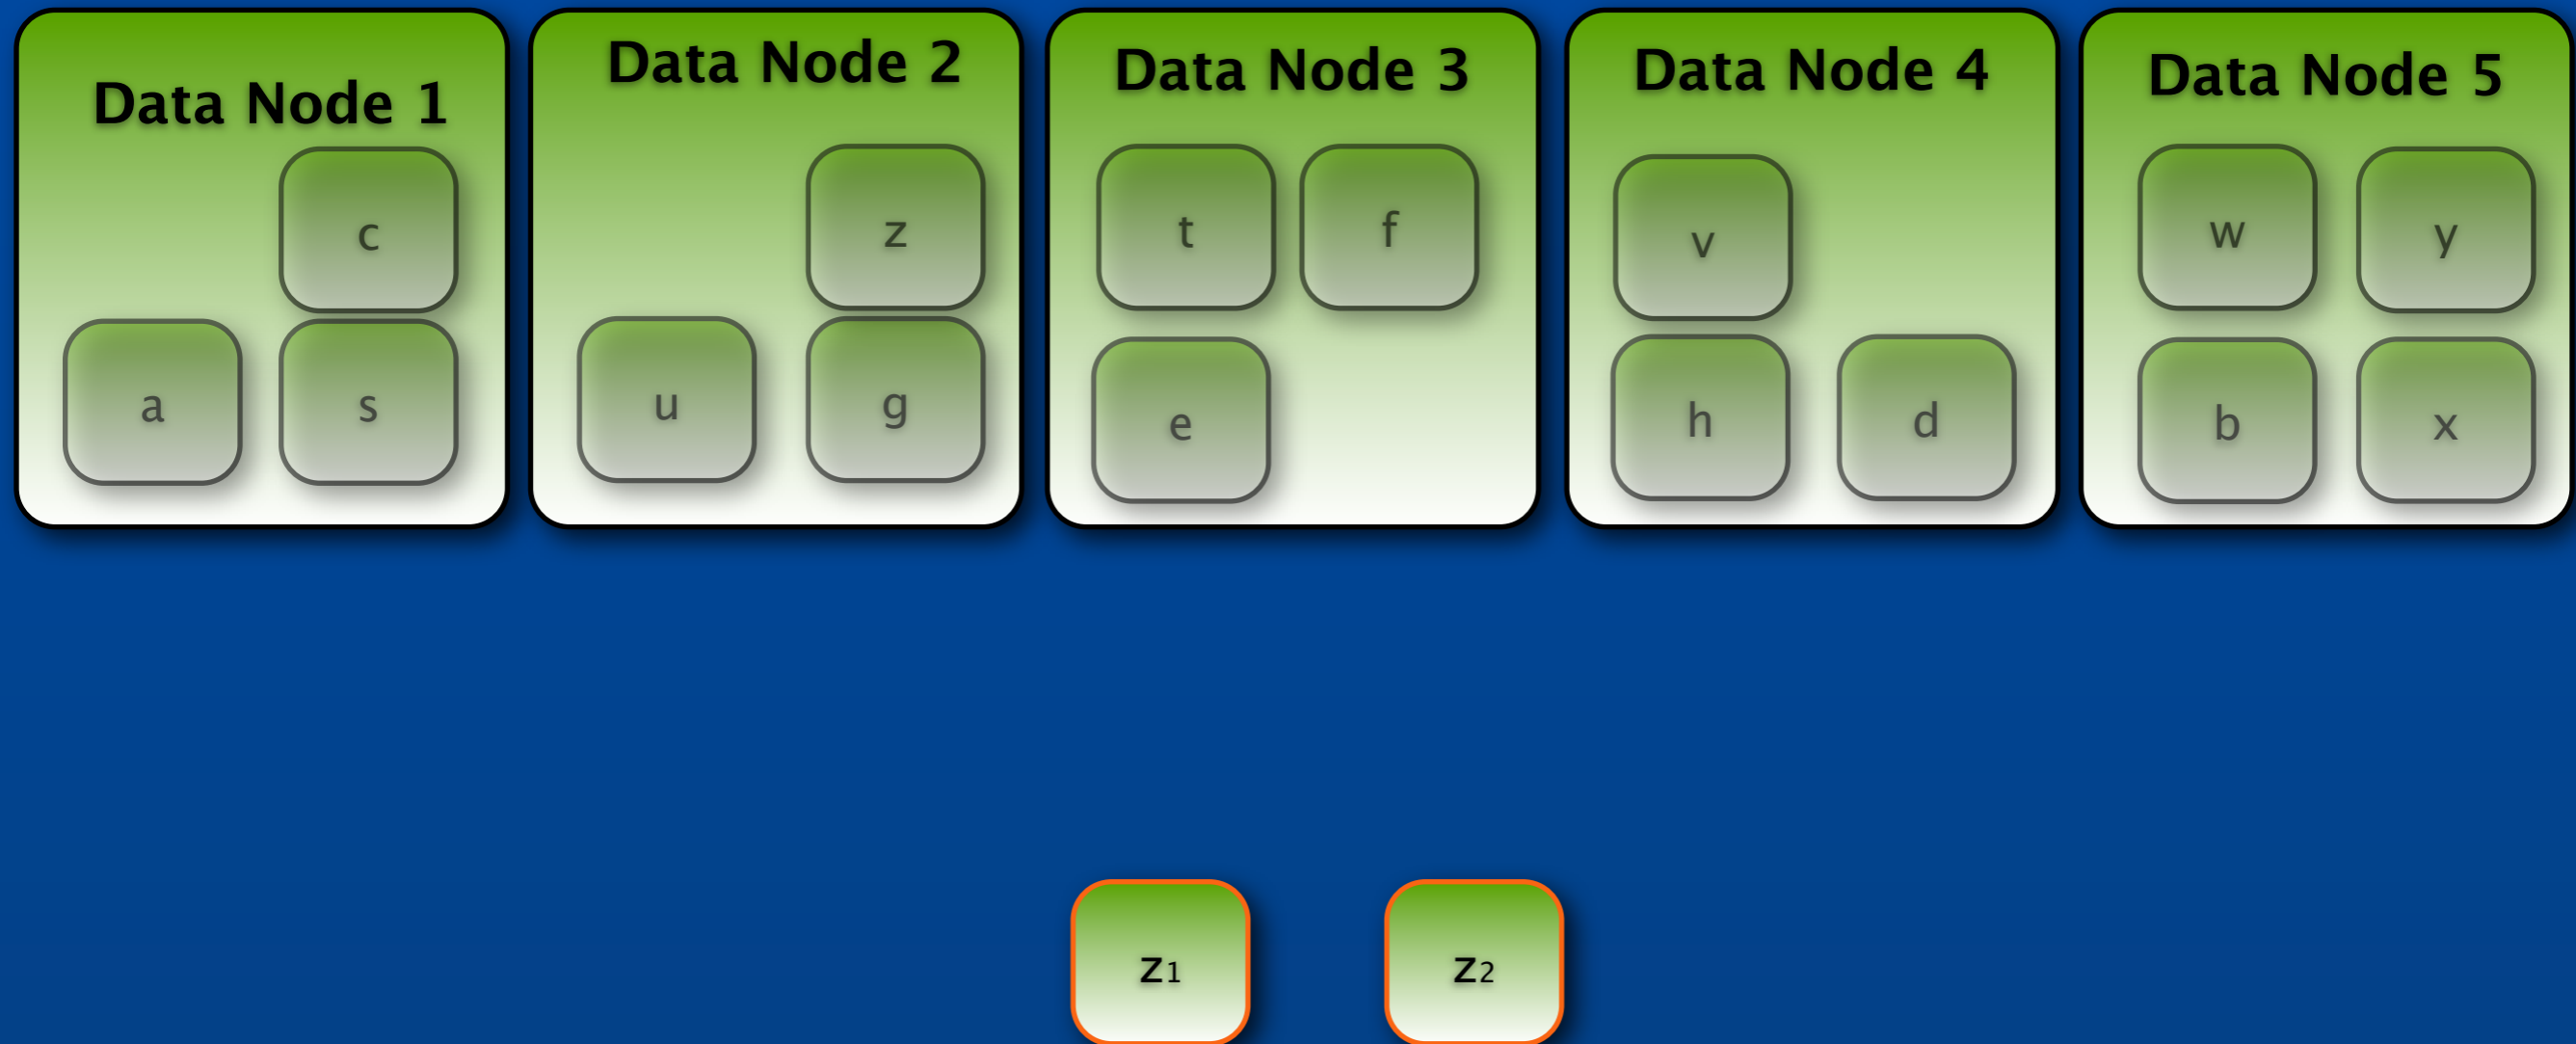
If a chunk gets too large (default in MongoDB – 64mb per chunk),
It is split into two new chunks

Chunk Splitting & Balancing



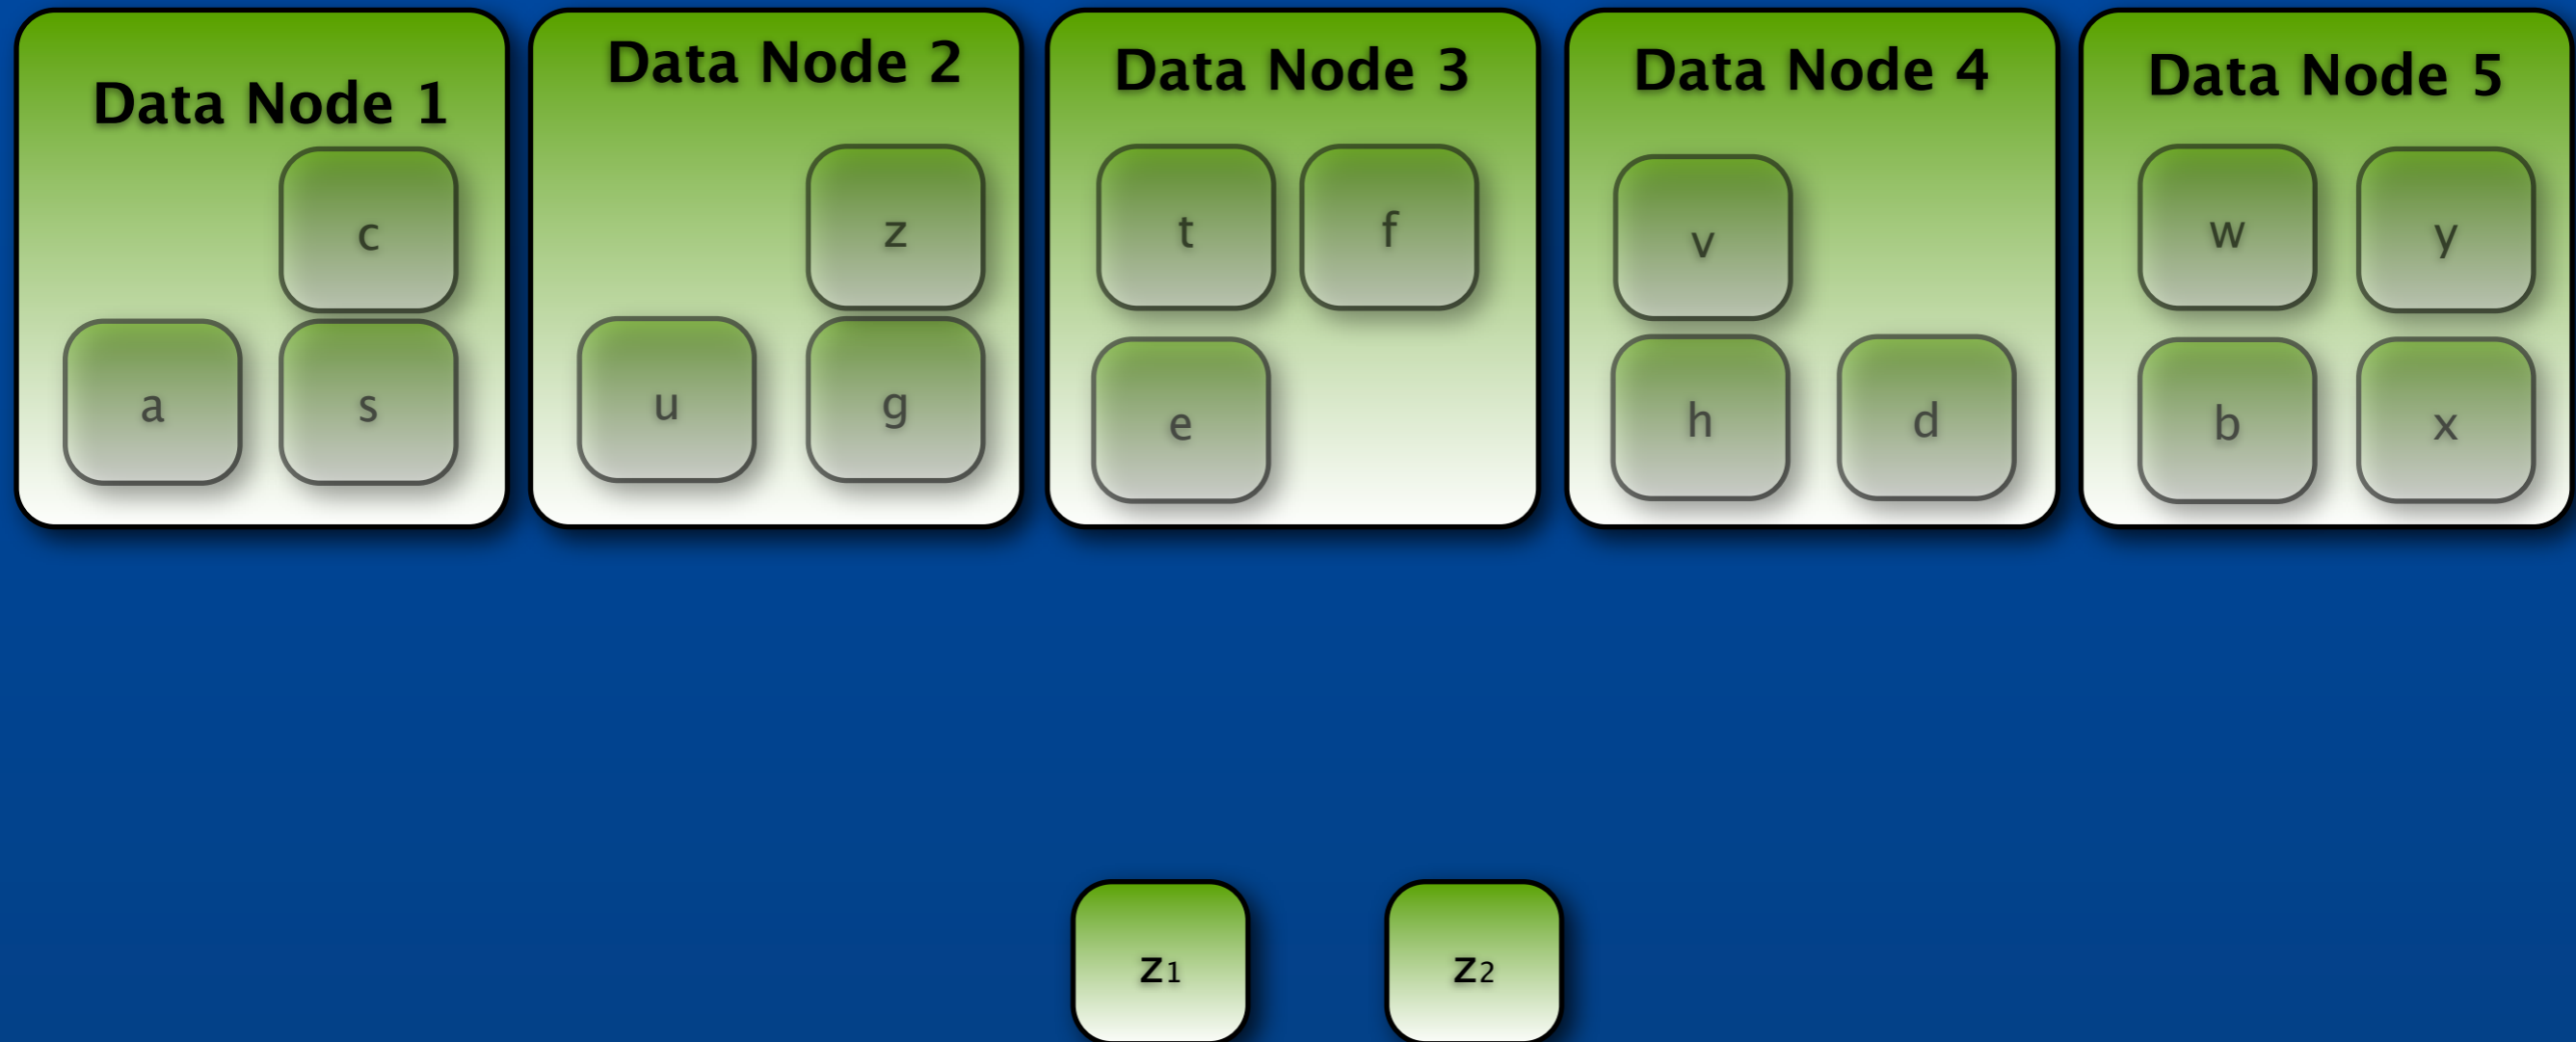
If a chunk gets too large (default in MongoDB – 64mb per chunk),
It is split into two new chunks

Chunk Splitting & Balancing




If a chunk gets too large (default in MongoDB – 64mb per chunk),
It is split into two new chunks

Chunk Splitting & Balancing




If a chunk gets too large (default in MongoDB – 64mb per chunk),
It is split into two new chunks

Chunk Splitting & Balancing



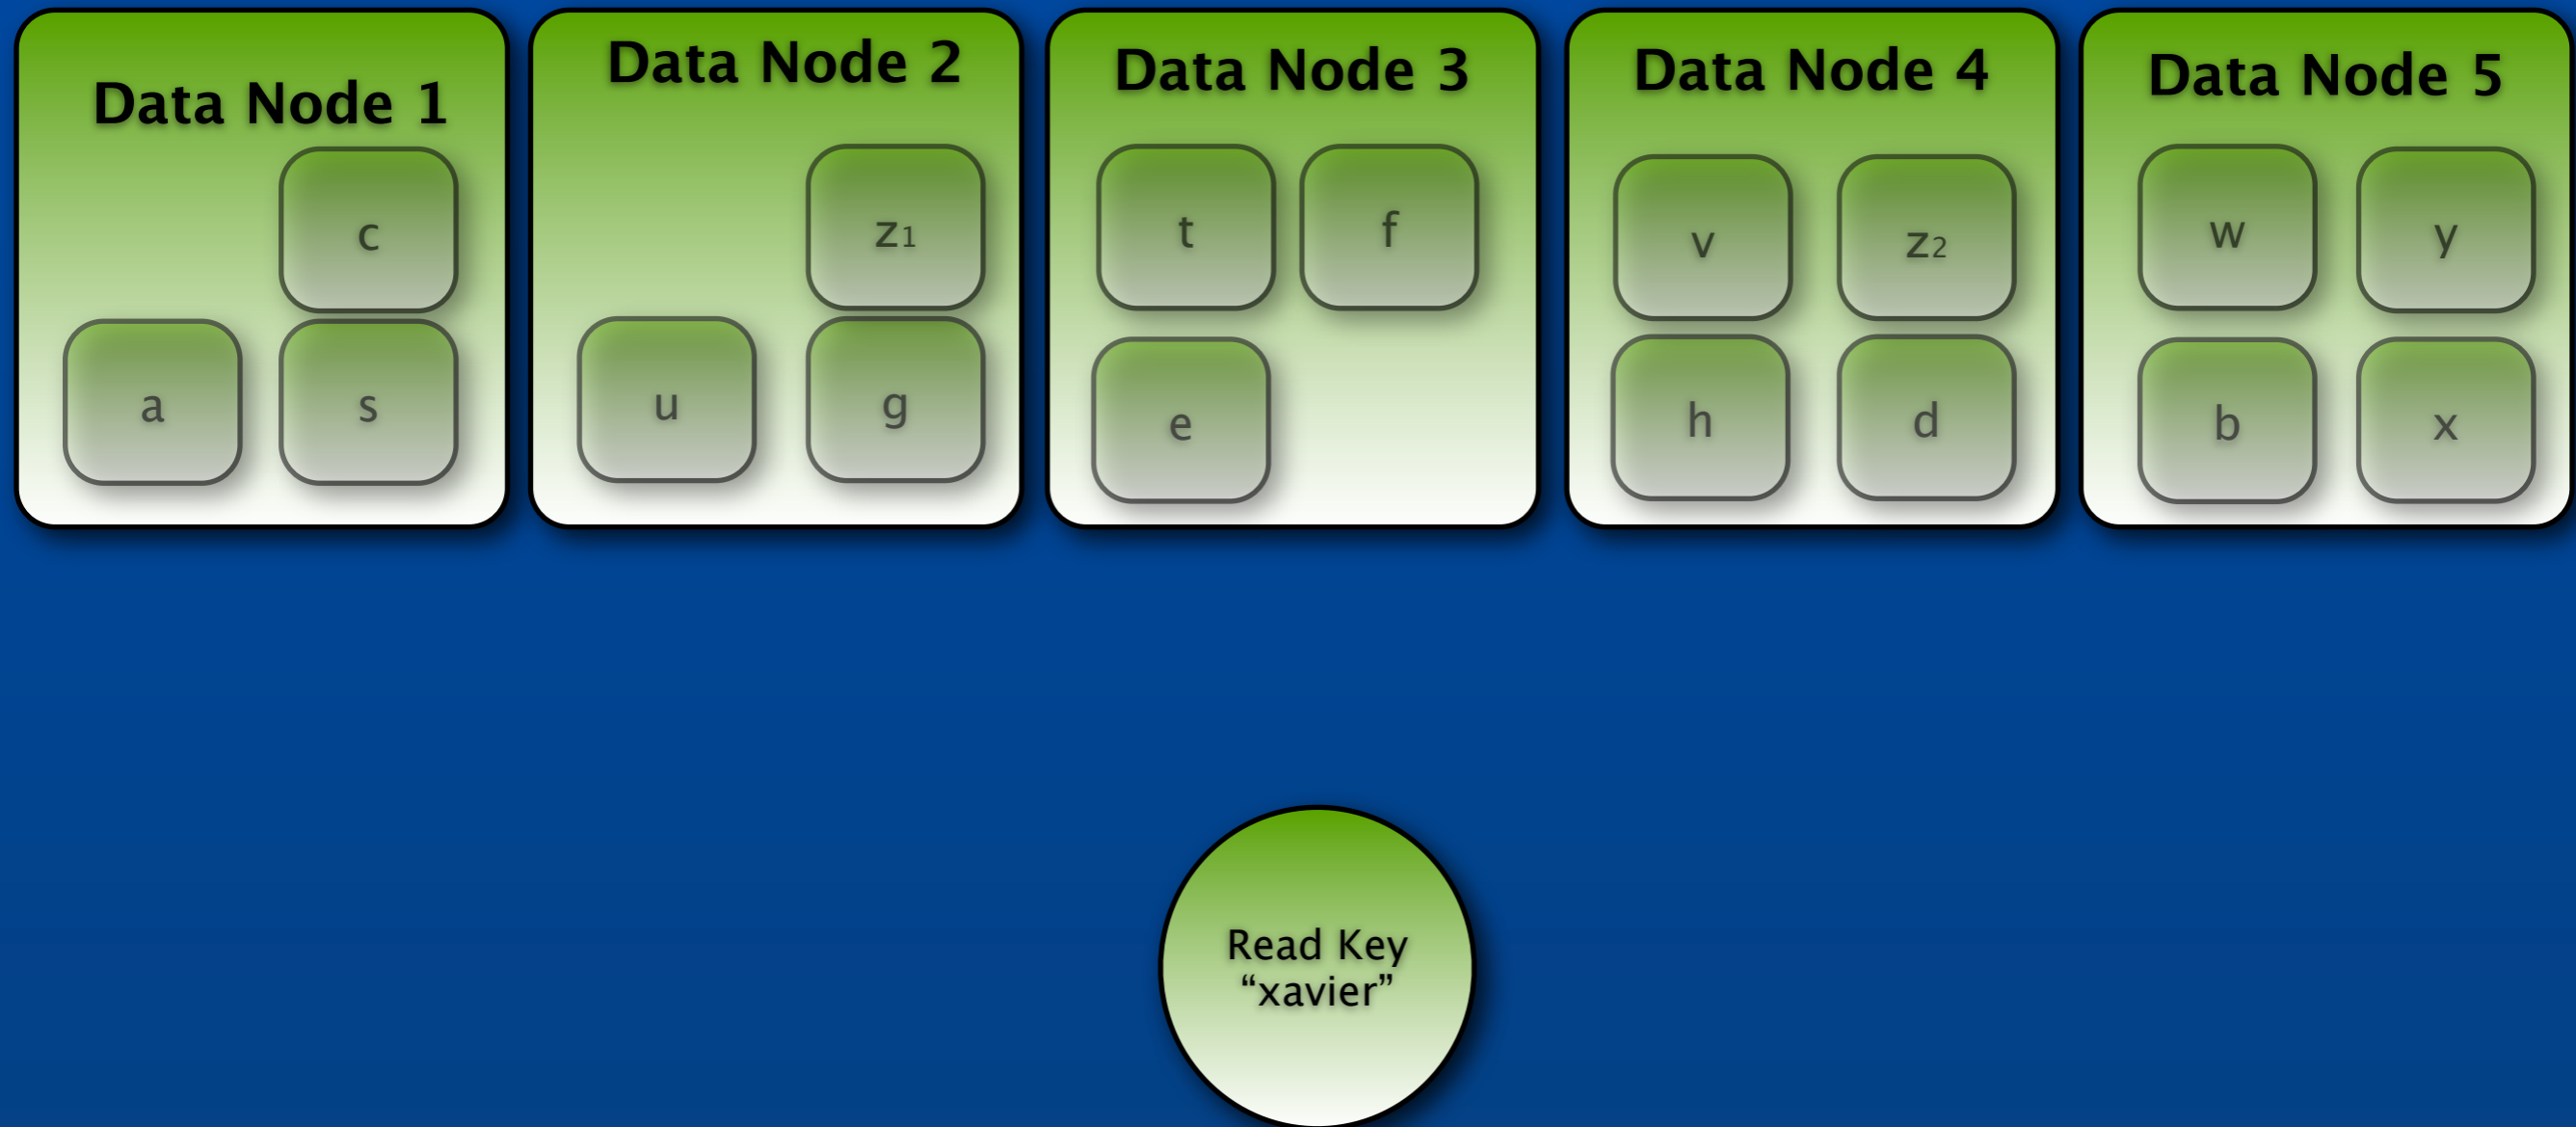
Each new part of the **Z** chunk (left & right) now contains half of the keys

Chunk Splitting & Balancing



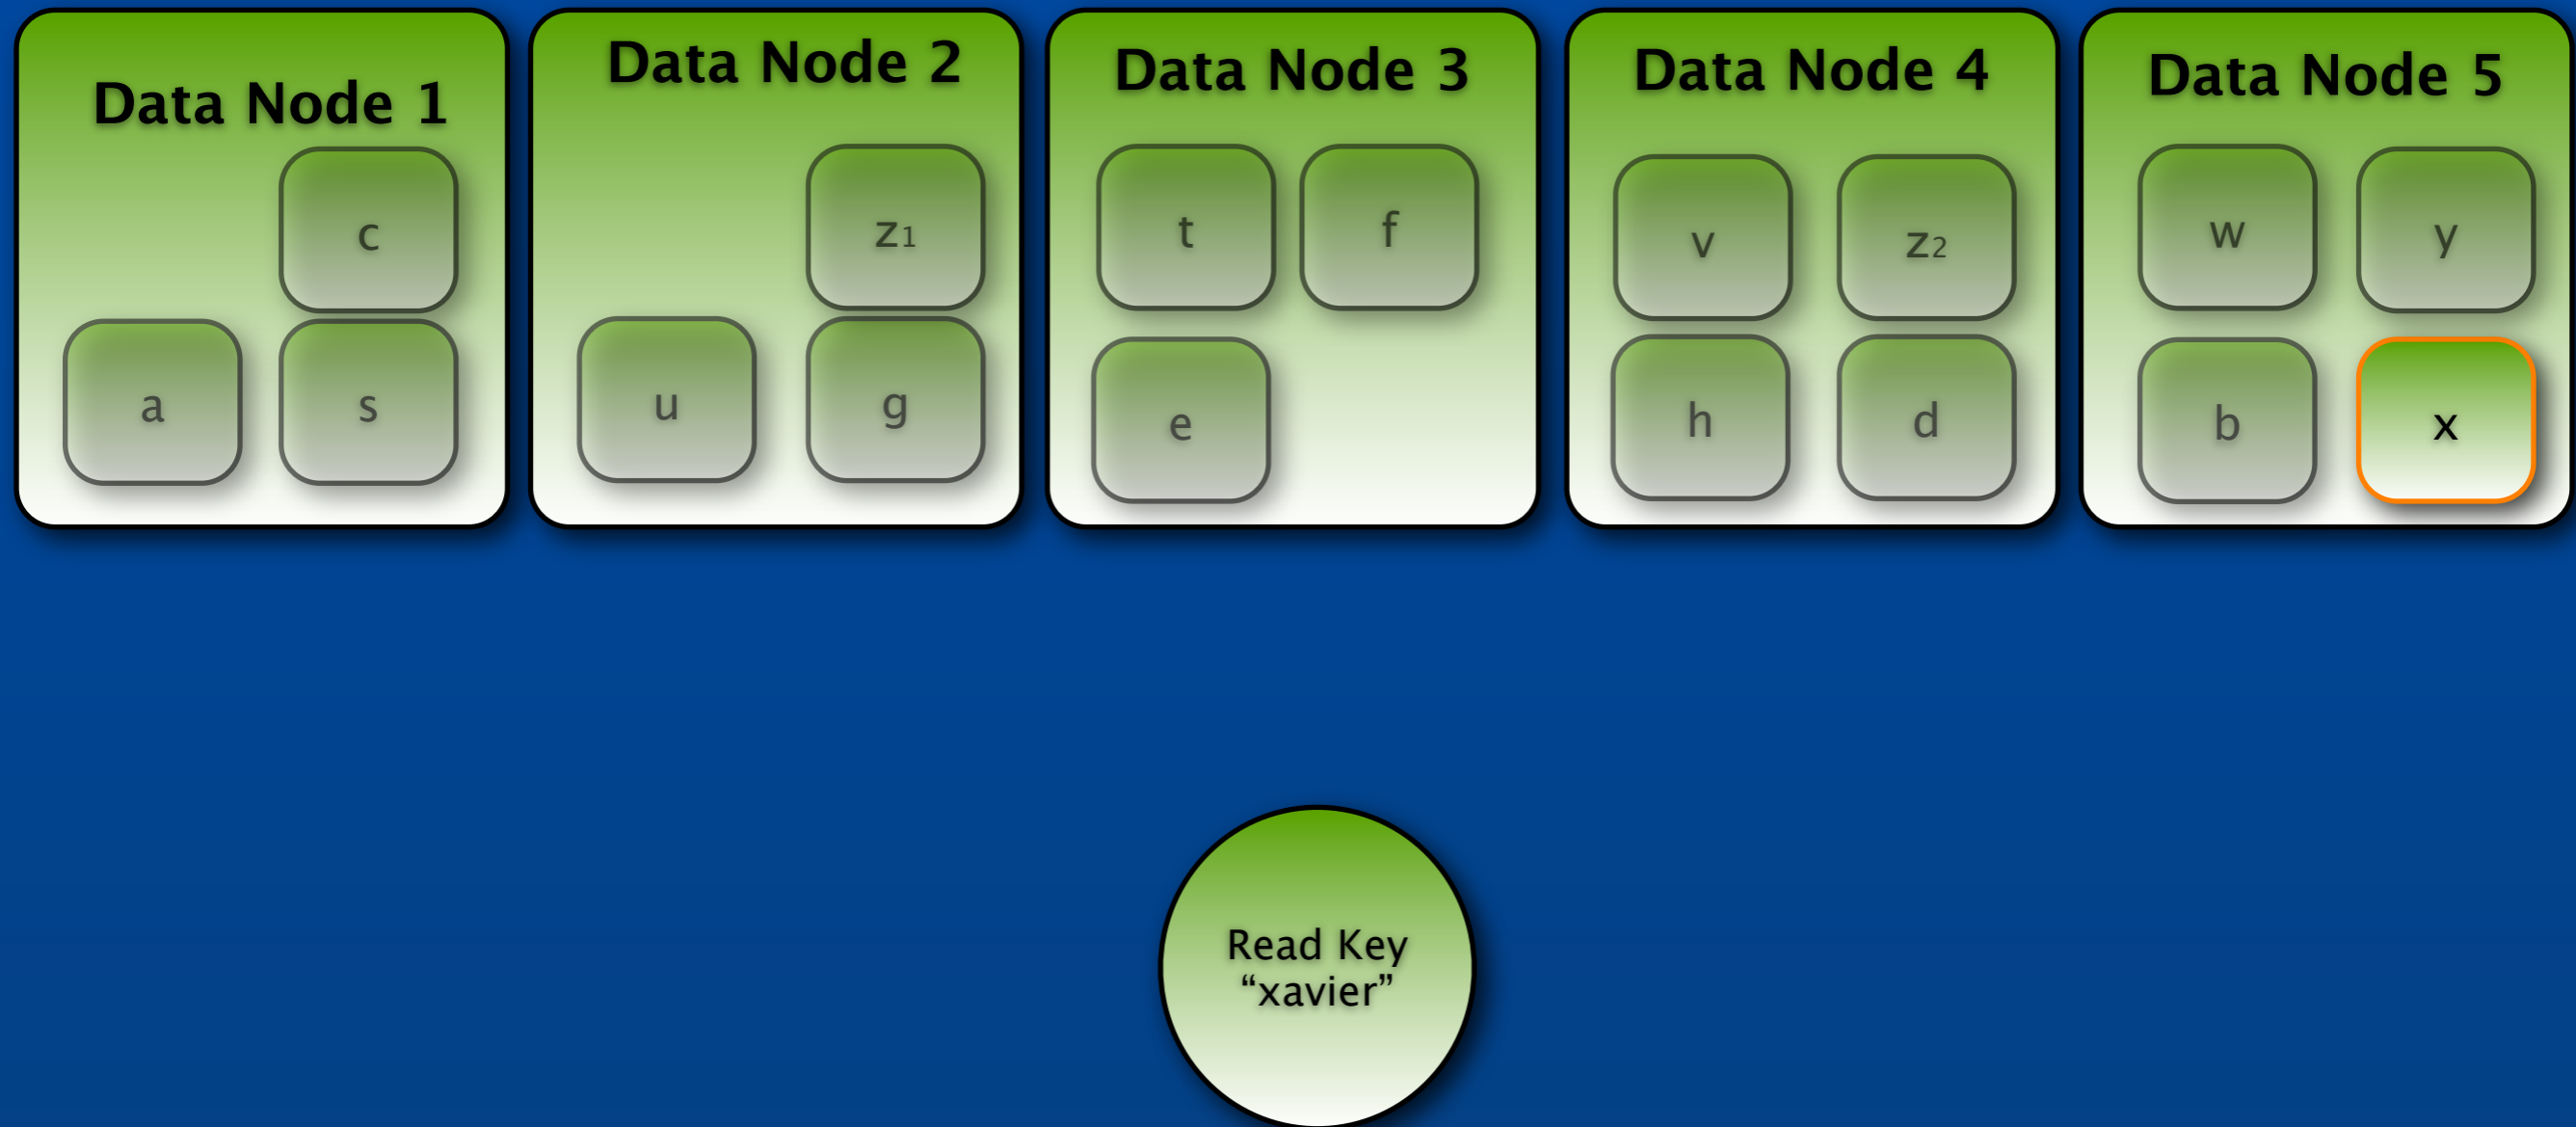
As chunks continue to grow and split, they can be rebalanced to keep an equal share of data on each server.

Reads with Key Routed Efficiently



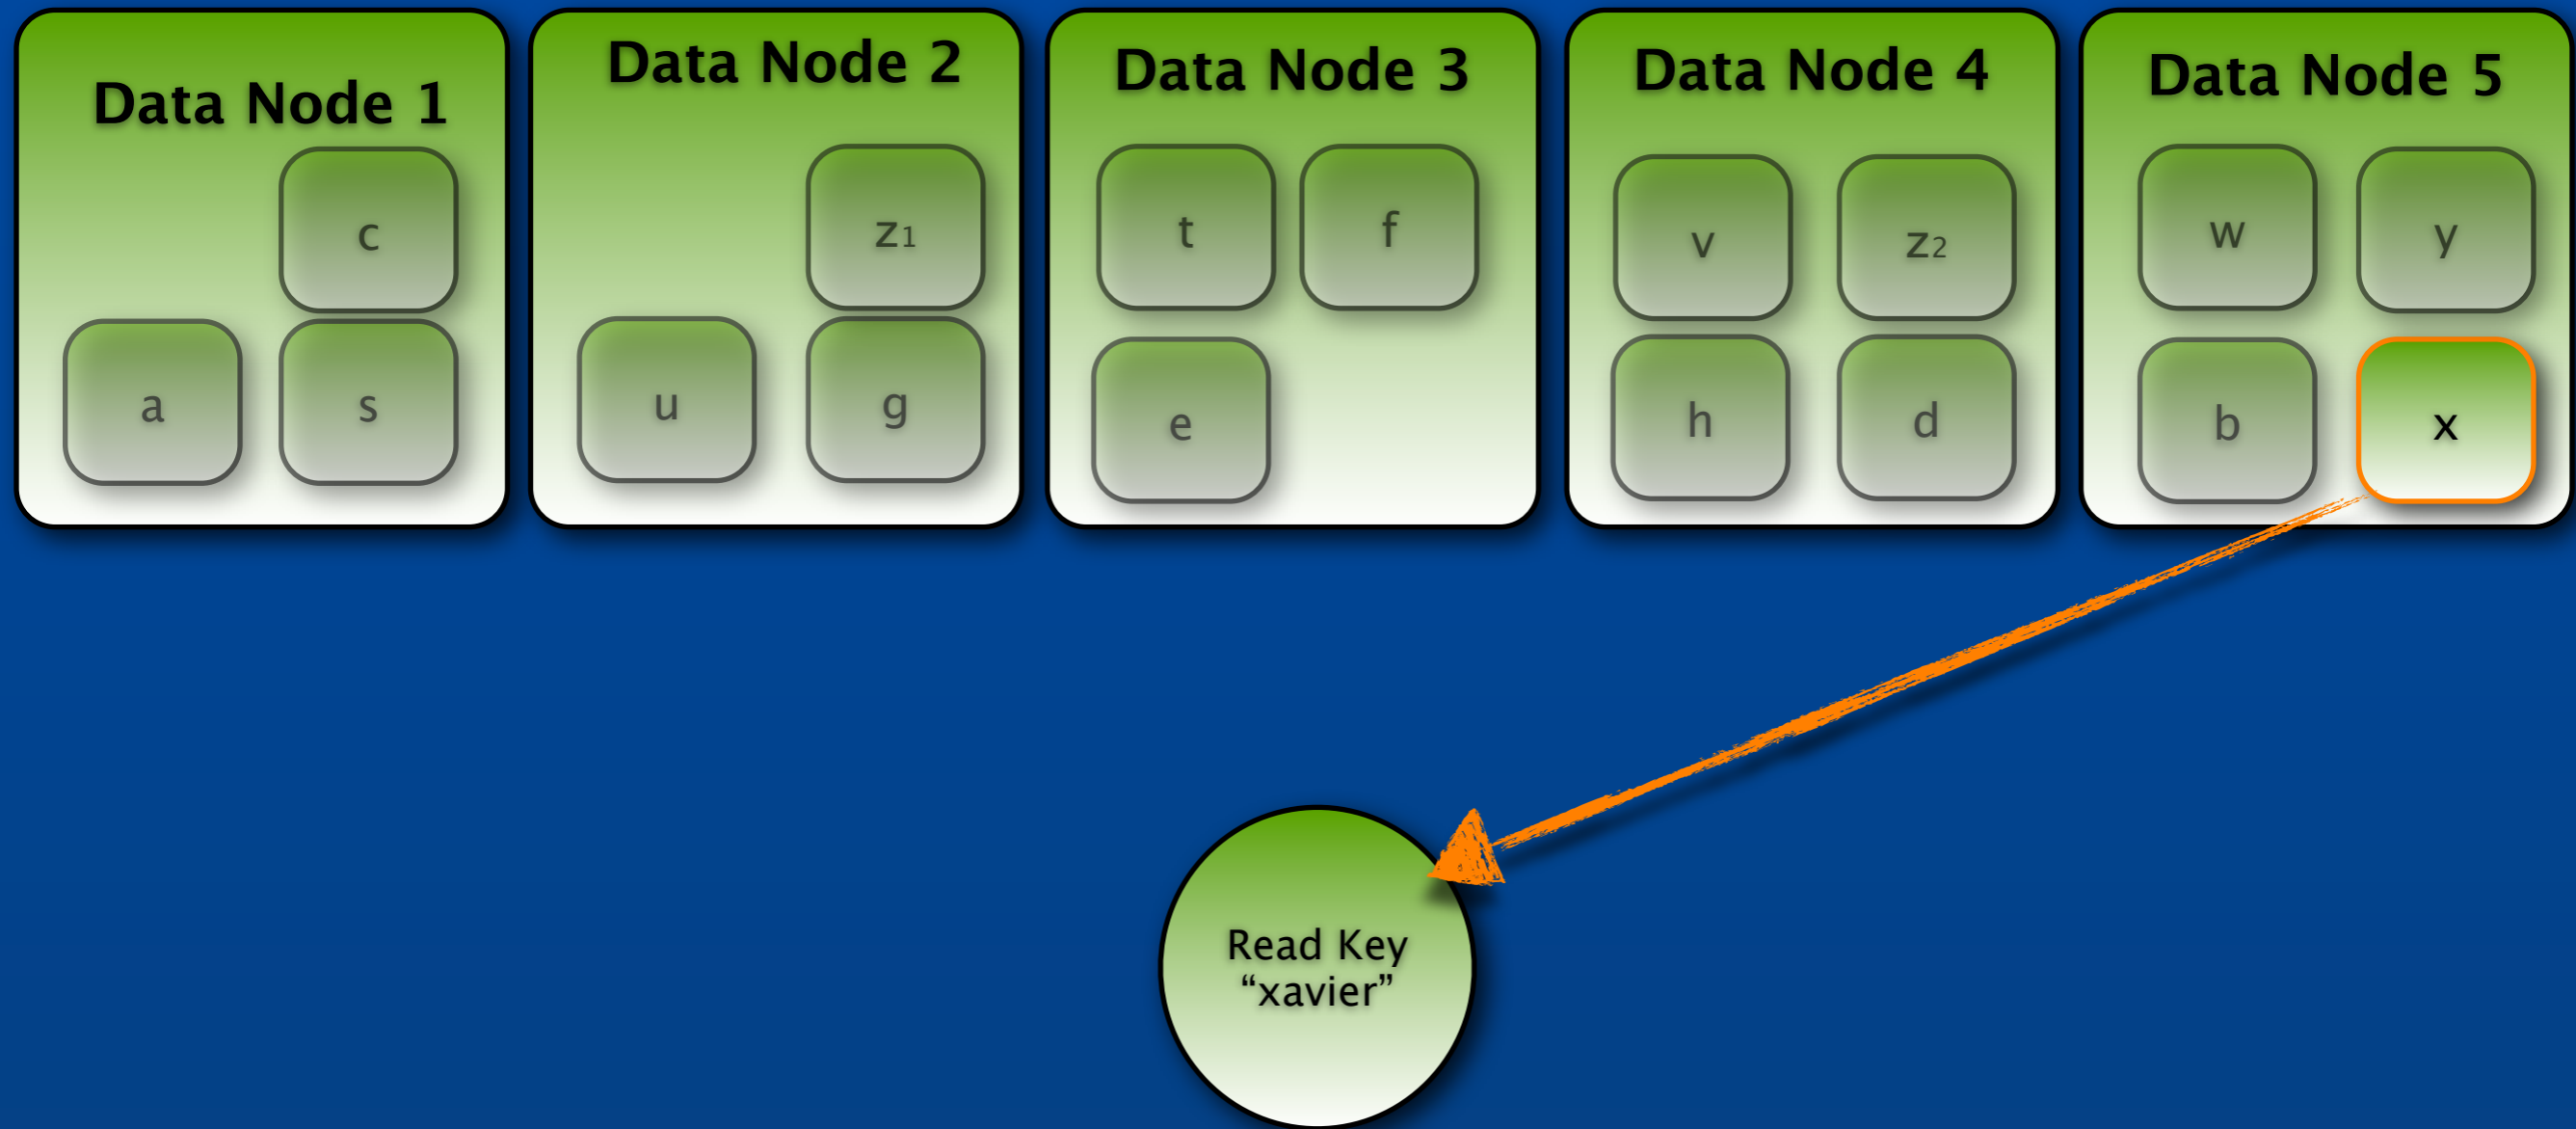
Reading a single value by Primary Key
Read routed efficiently to specific chunk containing key

Reads with Key Routed Efficiently



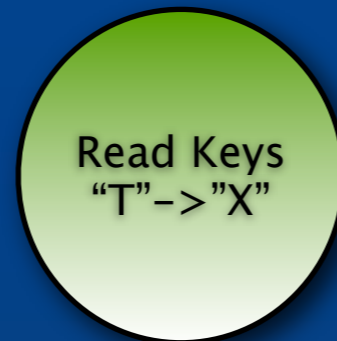

Reading a single value by Primary Key
Read routed efficiently to specific chunk containing key

Reads with Key Routed Efficiently




Reading a single value by Primary Key
Read routed efficiently to specific chunk containing key

Reads with Key Routed Efficiently



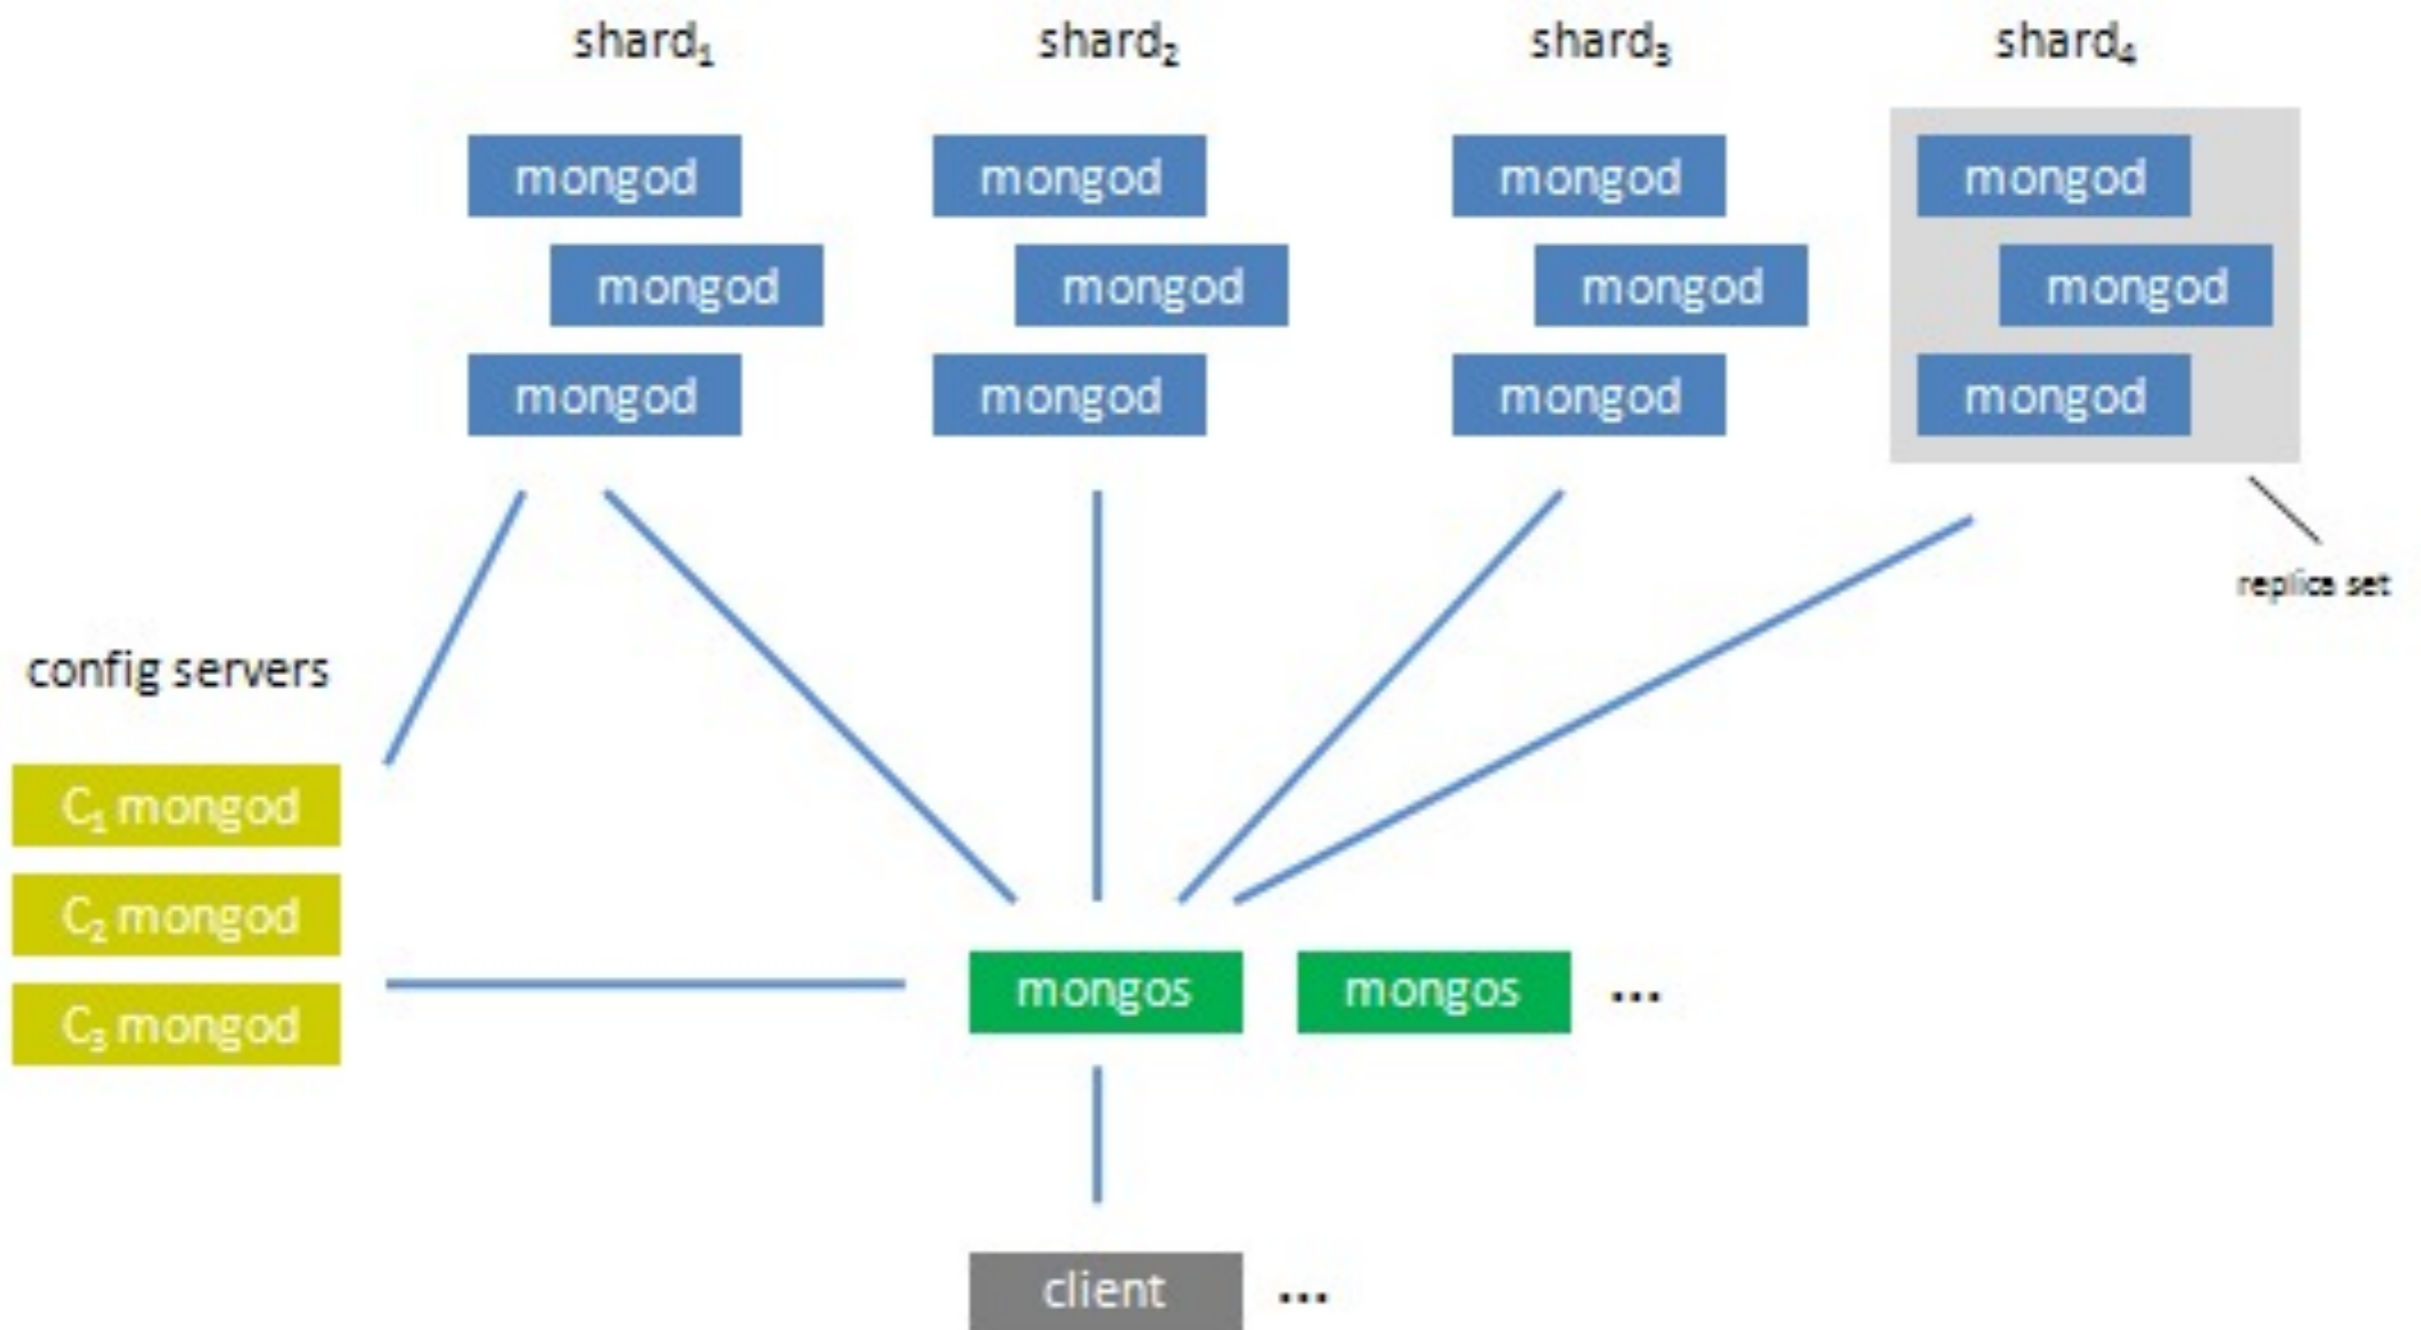
Reading multiple values by Primary Key
Reads routed efficiently to specific chunks in range

Reads with Key Routed Efficiently

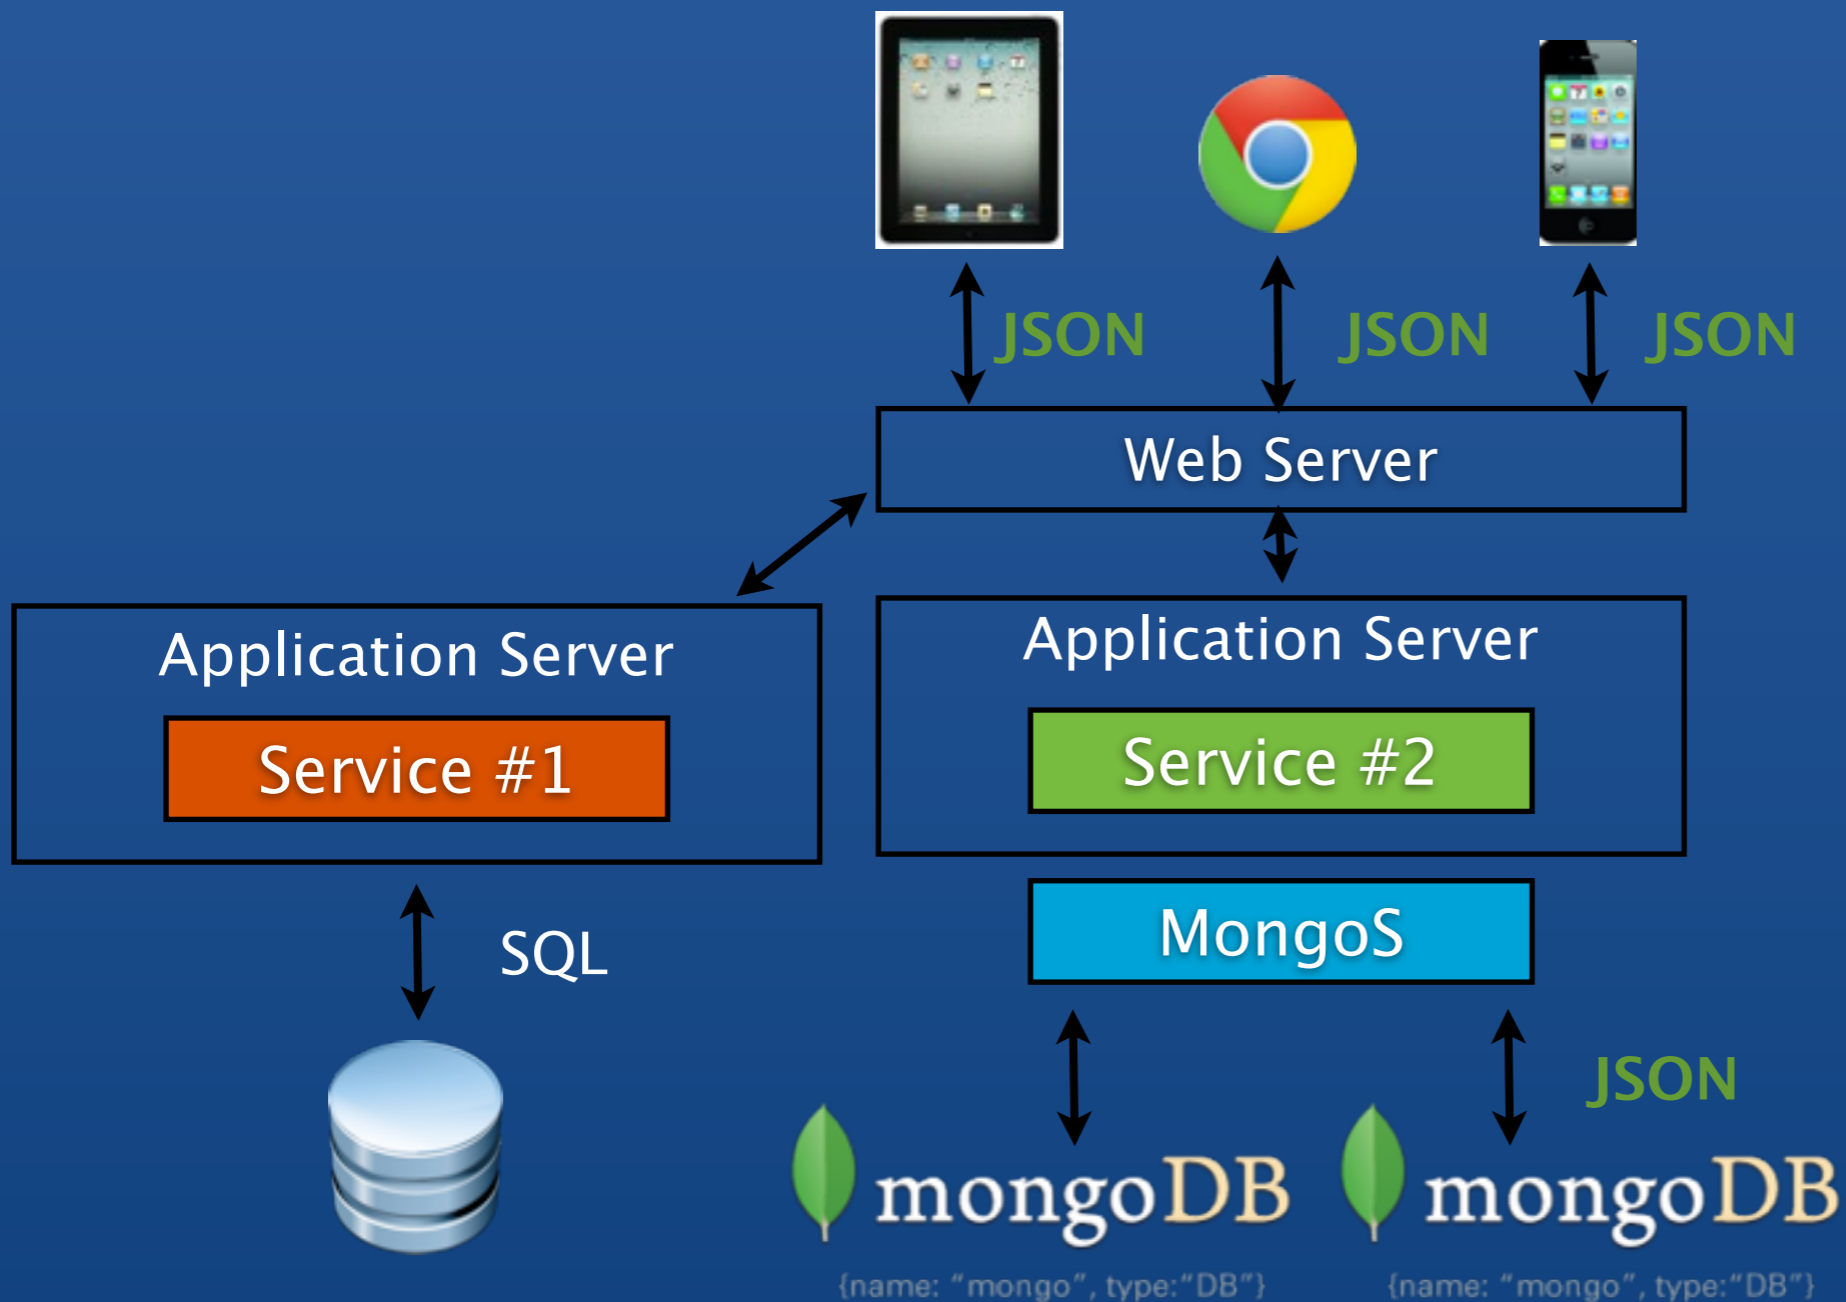


Reading multiple values by Primary Key
Reads routed efficiently to specific chunks in range

Architecture



Adding MongoS

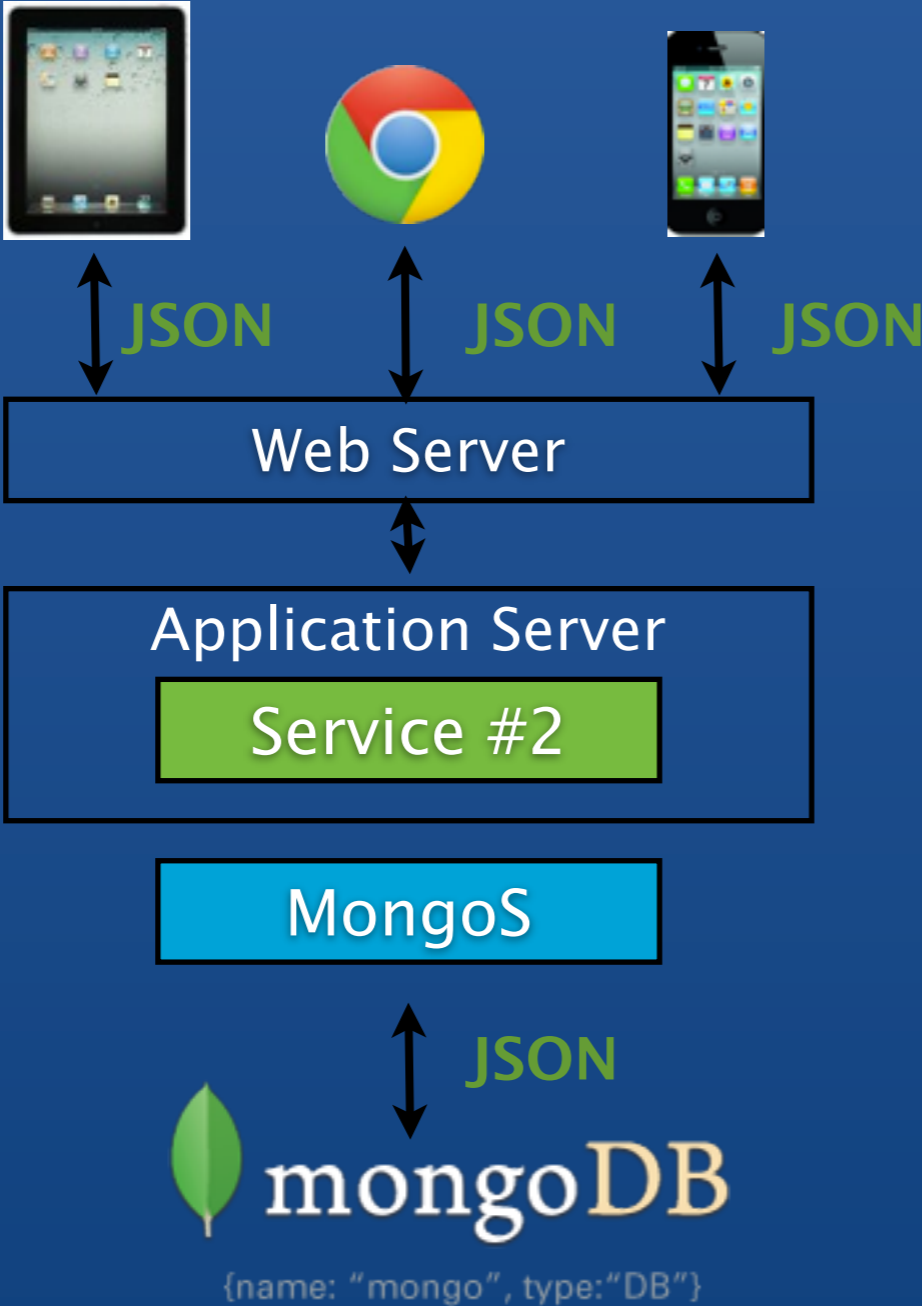


Challenge #3 -Offline Processing




Online / Offline

Online



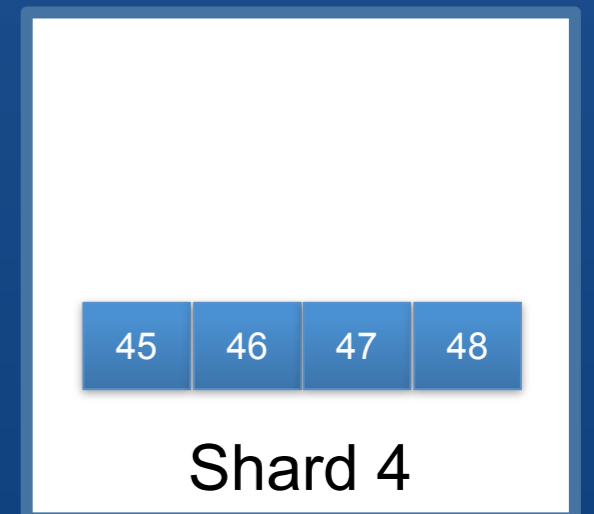

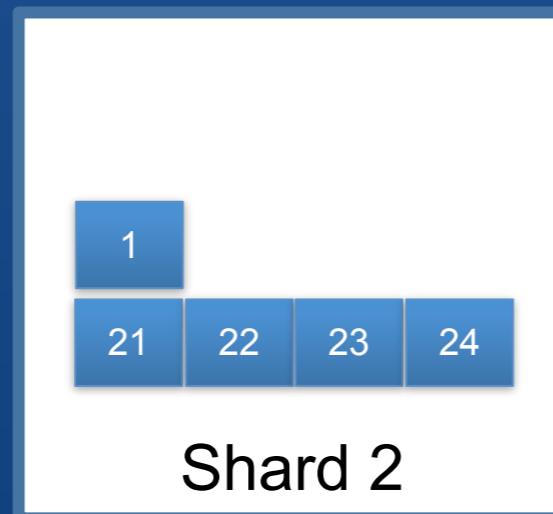

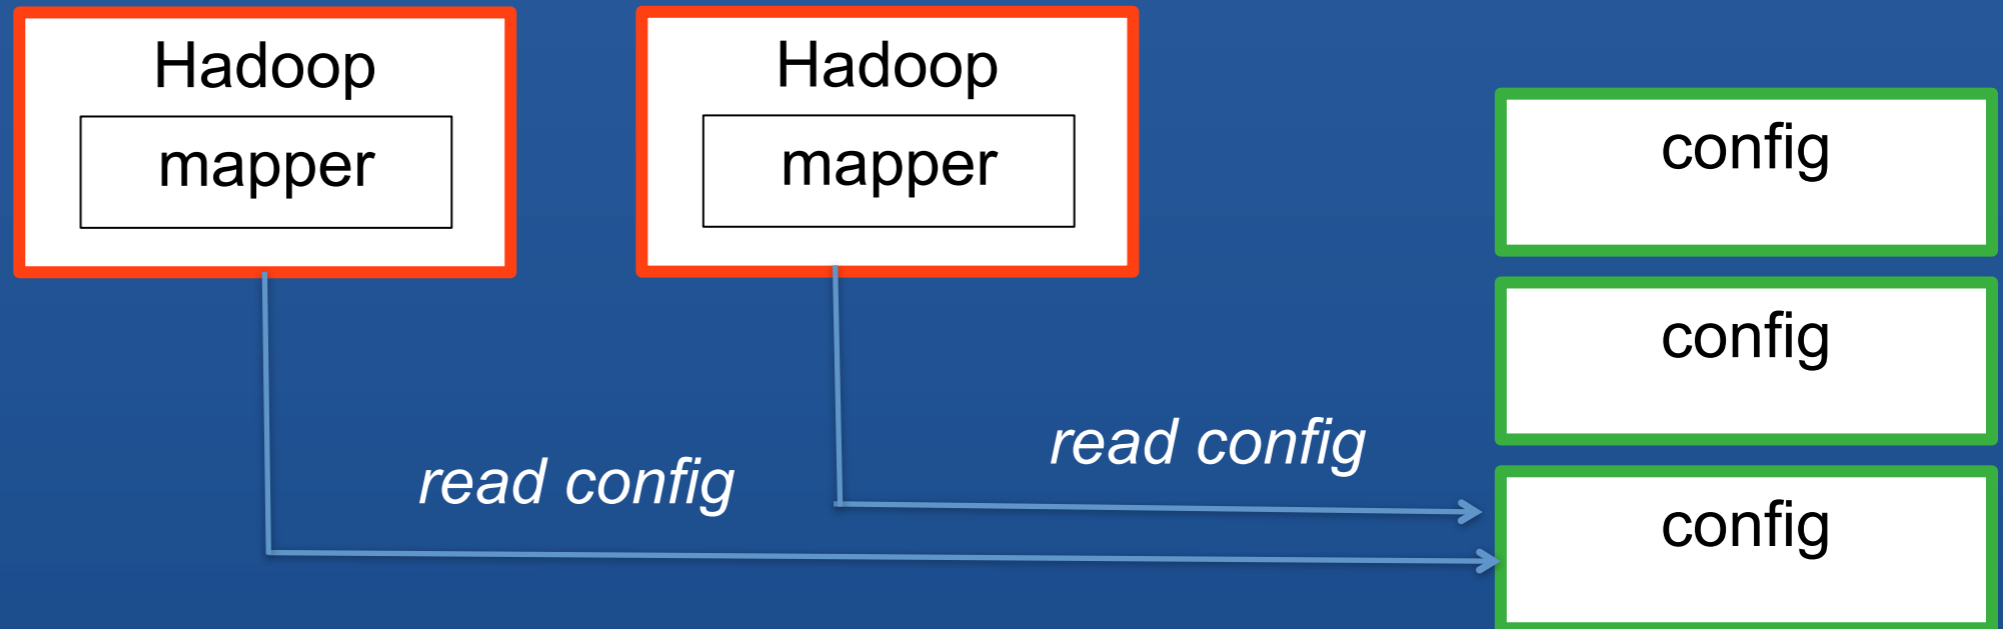
Offline




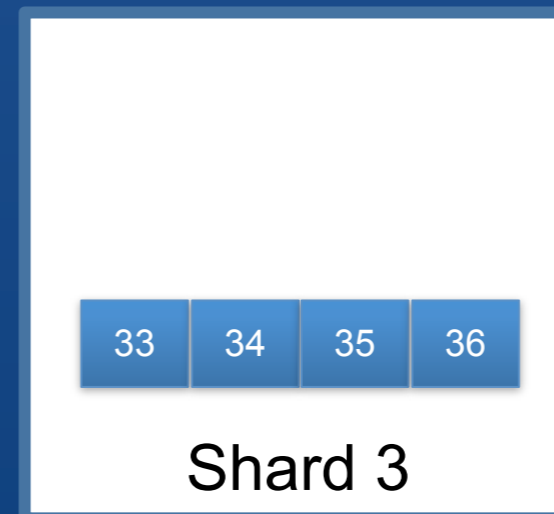

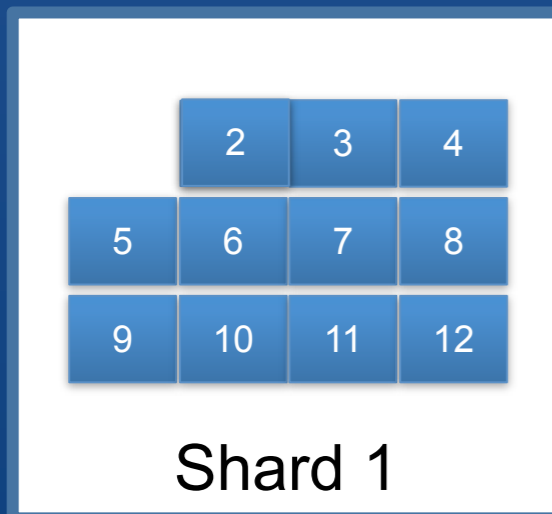



MongoDB and Hadoop









MongoDB & Hadoop



MongoDB & Hadoop



MongoDB & Hadoop



Word Count - Map

Classic Hadoop

```
public void map(LongWritable key, Text value, Context context) throws ..{
    String line = value.toString();
    StringTokenizer tokenizer = new StringTokenizer(line);
    while (tokenizer.hasMoreTokens()) {
        word.set(tokenizer.nextToken());
        context.write(word, one);
    }
}
```

MongoDB Hadoop Adapter

```
public void map(Object key, BSONObject value, Context context ) throws ....{
    StringTokenizer itr = new StringTokenizer( value.get( "line" ).toString() );
    while (tokenizer.hasMoreTokens()) {
        word.set(tokenizer.nextToken());
        context.write(word, one);
    }
}
```

Word Count - Reduce

Classic Hadoop

```
public void reduce( Text key, Iterable<IntWritable> values, Context context )
    throws IOException, InterruptedException{
    int sum = 0;
    for ( final IntWritable val : values ){
        sum += val.get();
    }
    context.write( key, new IntWritable(sum));
}
```

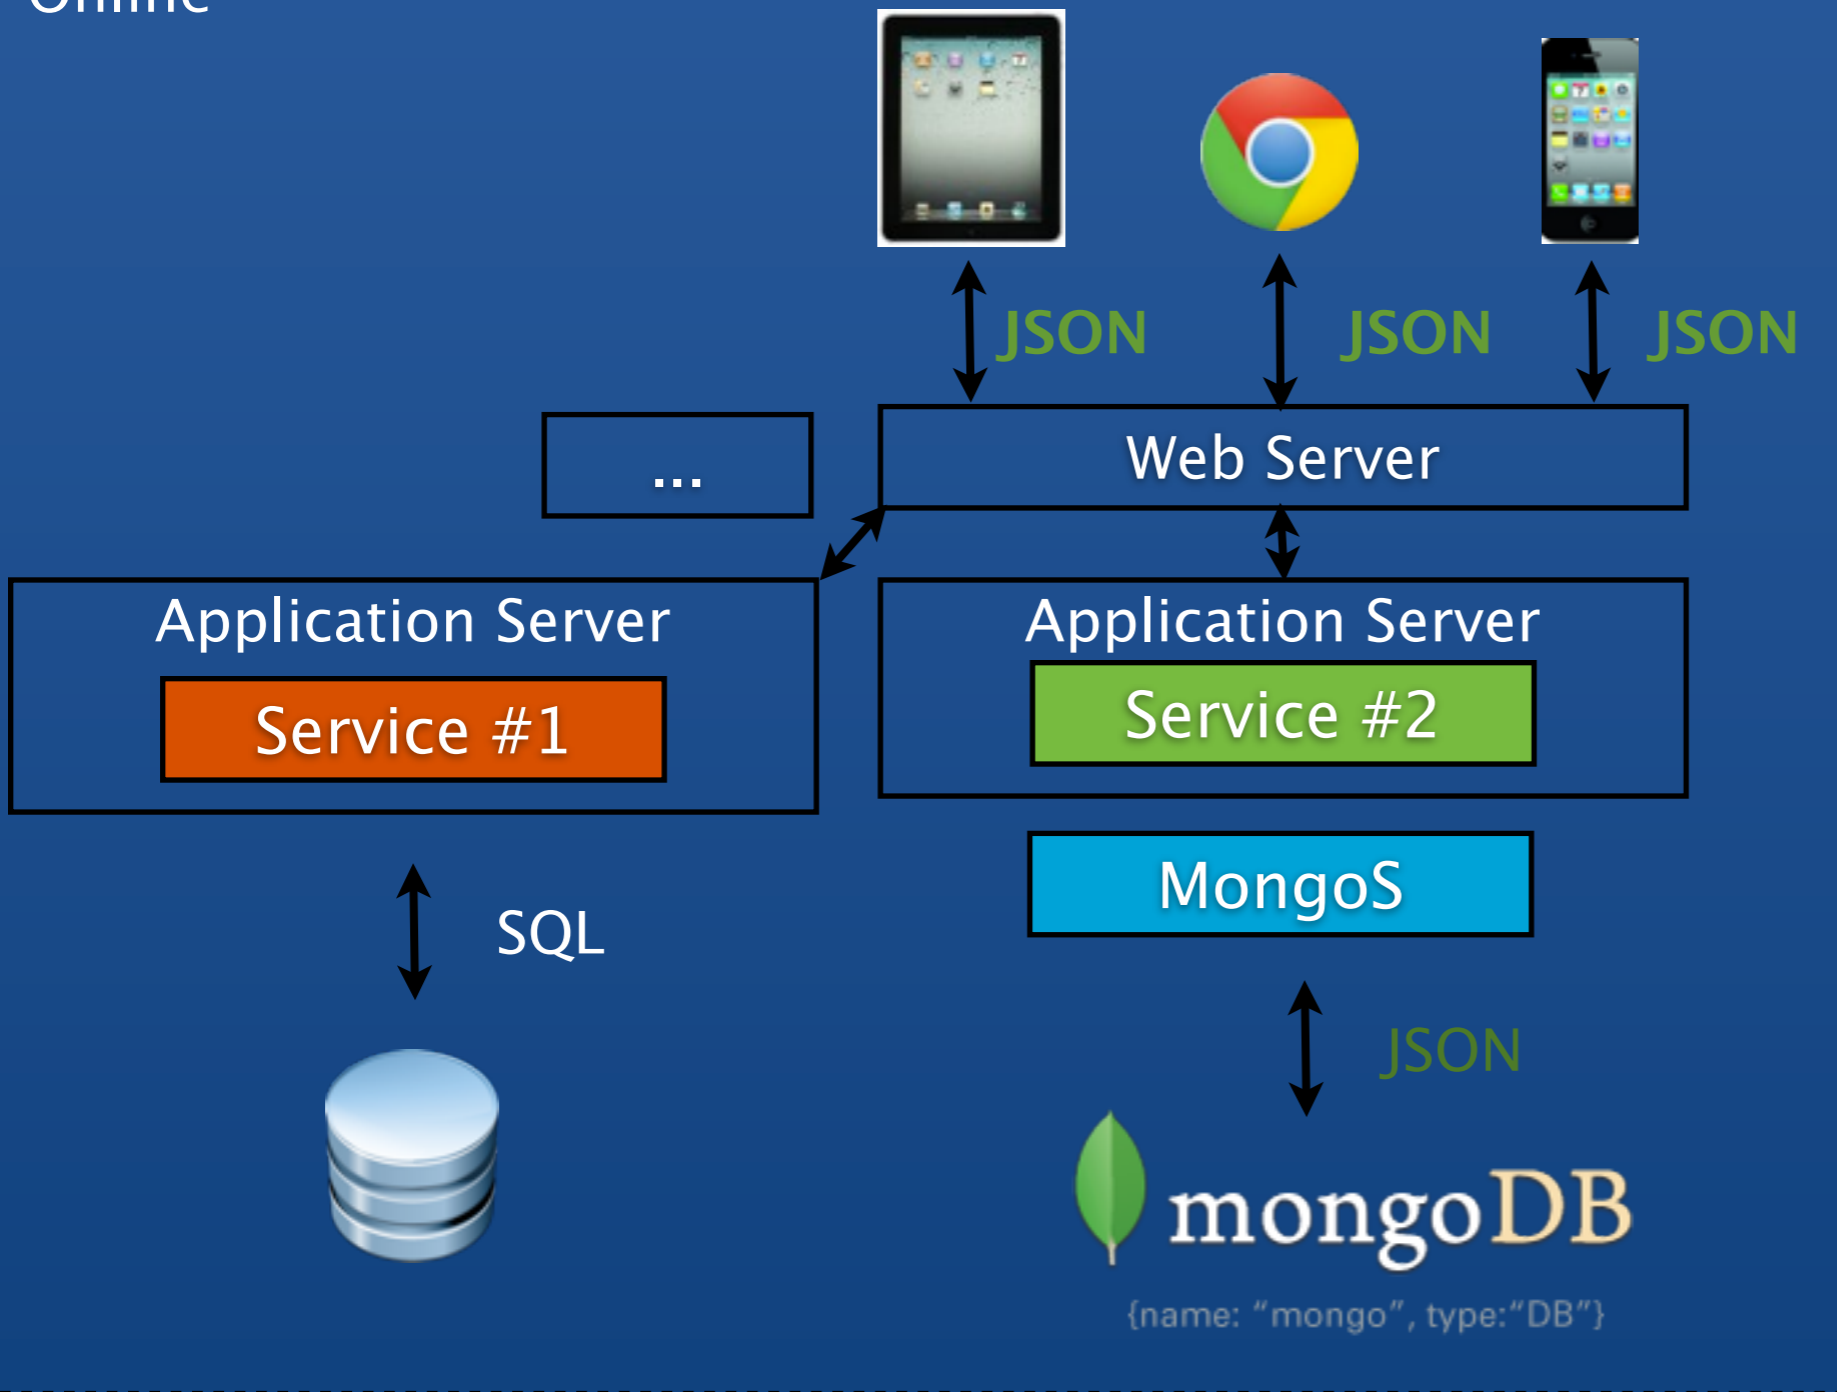
MongoDB Hadoop Adapter is the same code!

New Hybrid World



New Hybrid World

Online



Offline



 download at mongodb.org

We're Hiring !

Chris Harris
Email : charris@10gen.com
Twitter : [cj_harris5](https://twitter.com/cj_harris5)

conferences, appearances
<http://www.10gen.com/events>

