



Big Data
WITHOUT A BIG DATABASE

Kate Matsudaira
popforms
@katemats

TWO KINDS OF DATA

nicknames	“user”, “transactional”	“reference”, “non-transactional”
<i>examples:</i>	<ul style="list-style-type: none">• <i>user accounts</i>• <i>shopping cart/orders</i>• <i>user messages</i>• <i>...</i>	<ul style="list-style-type: none">• <i>product/offer catalogs</i>• <i>service catalogs</i>• <i>static geolocation data</i>• <i>dictionaries</i>• <i>...</i>
created/modified by:	users	business (you)
<i>sensitivity to staleness:</i>	<i>high</i>	<i>low</i>
plan for growth:	hard	easy
<i>access optimization:</i>	<i>read/write</i>	<i>mostly read</i>

TWO KINDS OF DATA

nicknames	“user”, “transactional”	“reference”, “non-transactional”
<i>examples:</i>	<ul style="list-style-type: none"> • <i>user accounts</i> • <i>shopping cart/orders</i> • <i>user messages</i> • <i>...</i> 	<ul style="list-style-type: none"> • <i>product/offer catalogs</i> • <i>service catalogs</i> • <i>static geolocation data</i> • <i>dictionaries</i> • <i>...</i>
created/modified by:	users	business (you)
<i>sensitivity to staleness:</i>	<i>high</i>	<i>low</i>
plan for growth:	hard	easy
<i>access optimization:</i>	<i>read/write</i>	<i>mostly read</i>

Your Amazon.com

Product Recommendations Health & Personal Care Shoes & Accessories Beauty Home & Kitchen Daily Recommendations

Electronics

Grid of 15 electronics products including headphones, earbuds, and speakers with prices and ratings.

Health & Personal Care

Grid of 15 health and personal care products including massage chairs, foot massagers, and ergonomic chairs.

Shoes & Accessories

Grid of 15 shoes and accessories including men's dress shoes, casual shoes, and sandals.

Beauty

Grid of 15 beauty products including facial cleansers, toners, and moisturizers.

Home & Kitchen

Grid of 15 home and kitchen appliances including blenders, coffee makers, and air fryers.

Baby

Grid of 15 baby products including strollers, car seats, and baby monitors.

Home Improvement

Grid of 15 home improvement products including tools, paint, and storage solutions.

TRAVEL FLIGHTS HOTELS

Compare to sponsored sites

Portland Hotels


Wed, 7/18 - Fri, 7/20 · 1 adult · 1 room [Change search](#)

Rate Under \$100 \$101 - \$200 \$201 - \$300 \$301+

Class 1 star 2 star 3 star 4 star 5 star [More filters](#)

1 - 25 of 182 RESULTS Sort by Popularity Class Distance Price

- | | | | |
|---|--|--|---------------|
| 1 | | Hotel Deluxe
(503) 219-2094 · 4 star · 1.02 mi · | \$279 |
| Book with: \$279 www.hoteldeluxe... · Get rate Expedia | | | |
| 2 | | Hotel Lucia
(503) 225-1717 · 4 star · 0.39 mi · | Call for info |
| Book with: \$259 Orbitz · \$259 CheapTickets · More | | | |
| 3 | | The Heathman Hotel
(503) 241-4100 · 4 star · 0.87 mi · | \$259 |
| Book with: \$259 Orbitz · \$259 CheapTickets · More | | | |
| 4 | | Governor Hotel
(503) 224-3400 · 4 star · 0.69 mi · | \$238 |
| Book with: \$238 Kayak · \$238 Orbitz · \$238 CheapTickets · More | | | |
| 5 | | The Nines, Portland
(503) 222-9996 · 5 star · 0.56 mi · | Call for info |
| Book with: \$152 Kayak · \$152 Hotels.com · \$152 getaroom · More | | | |
| 6 | | Hilton Vancouver Washington
(360) 993-4500 · 3 star · 11.37 mi · | \$152 |
| Book with: \$329 Kimpton · \$329 ReserveTravel · \$339 Kayak · More | | | |
| 7 | | Hotel Vintage Plaza, a Kimpton Hotel
(503) 228-1212 · 4 star · 0.41 mi · | \$329 |
| Book with: \$329 Kimpton · \$329 ReserveTravel · \$339 Kayak · More | | | |





Compare to sponsored sites



Portland Hotels

Wed, 7/18 - Fri, 7/20 · 1 adult · 1 room

Change search

Rate Under \$100 \$101 - \$200 \$201 - \$300 \$301+ More filters
Class 1 star 2 star 3 star 4 star 5 star

1 - 25 of 182 RESULTS Sort by Popularity Class Distance Price

- Hotel Deluxe**
 (503) 219-2094 · 4 star · 1.02 mi · ★★★★★
 Photos · Amenities · Details · Website · 31 reviews
 Book with: \$279 Expedia · Get rate Expedia **\$279**
- Hotel Lucia**
 (503) 225-1717 · 4 star · 0.39 mi · ★★★★★
 Photos · Amenities · Details · Website · 15 reviews
 Call for rate
- The Heathman Hotel**
 (503) 241-4100 · 4 star · 0.17 mi · ★★★★★
 Photos · Amenities · Details · Website · 12 reviews
 Book with: \$259 Orbitz · \$259 CheapTickets **\$259**
- Governor Hotel**
 (503) 224-3400 · 4 star · 0.69 mi · ★★★★★
 Photos · Amenities · Details · Website · 24 reviews
 Book with: \$238 Kayak · \$238 Orbitz · \$238 CheapTickets · More **\$238**
- The Nines, Portland**
 (503) 222-9996 · 5 star · 0.56 mi · ★★★★★
 Photos · Amenities · Details · Website · 32 reviews
 Call for info
- Hilton Vancouver Washington**
 (360) 993-4500 · 3 star · 11.37 mi · ★★★★★
 Photos · Amenities · Details · Website · 16 reviews
 Book with: \$152 Kayak · \$152 Hotels.com · \$152 getaroom · More **\$152**
- Hotel Vintage Plaza, a Kimpton Hotel**
 (503) 228-1212 · 4 star · 0.41 mi · ★★★★★
 Photos · Amenities · Details · Website · 37 reviews **\$329**

reference data






user data

TRAVEL FLIGHTS **HOTELS**

Compare to sponsored sites



Portland Hotels

Wed, 7/18 - Fri, 7/20 · 1 adult · 1 room

Change search

Rate Under \$100 \$101 - \$200 \$201 - \$300 \$301+ More filters
Class 1 star 2 star 3 star 4 star 5 star

1 - 25 of 182 RESULTS Sort by Popularity Class Distance Price

- Hotel Deluxe**
 (503) 219-2094 · 4 star · 1.02 mi · ★★★★★
 Photos · Amenities · Details · Website · 31 reviews
 Book with: \$279 Expedia · Get rate Expedia **\$279**
- Hotel Lucia**
 (503) 225-1717 · 4 star · 0.39 mi · ★★★★★
 Photos · Amenities · Details · Website · 15 reviews
 Call for rate
- The Heathman Hotel**
 (503) 241-4100 · 4 star · 0.17 mi · ★★★★★
 Photos · Amenities · Details · Website · 15 reviews
 Book with: \$259 Orbitz · \$259 CheapTickets **\$259**
- Governor Hotel**
 (503) 224-3400 · 4 star · 0.69 mi · ★★★★★
 Photos · Amenities · Details · Website · 24 reviews
 Book with: \$238 Kayak · \$238 Orbitz · \$238 CheapTickets · More **\$238**
- The Nines, Portland**
 (503) 222-9996 · 5 star · 0.56 mi · ★★★★★
 Photos · Amenities · Details · Website · 32 reviews
 Call for info
- Hilton Vancouver Washington**
 (360) 993-4500 · 3 star · 11.37 mi · ★★★★★
 Photos · Amenities · Details · Website · 16 reviews
 Book with: \$152 Kayak · \$152 Hotels.com · \$152 getaroom · More **\$152**
- Hotel Vintage Plaza, a Kimpton Hotel**
 (503) 228-1212 · 4 star · 0.41 mi · ★★★★★
 Photos · Amenities · Details · Website · 37 reviews **\$329**


reference data




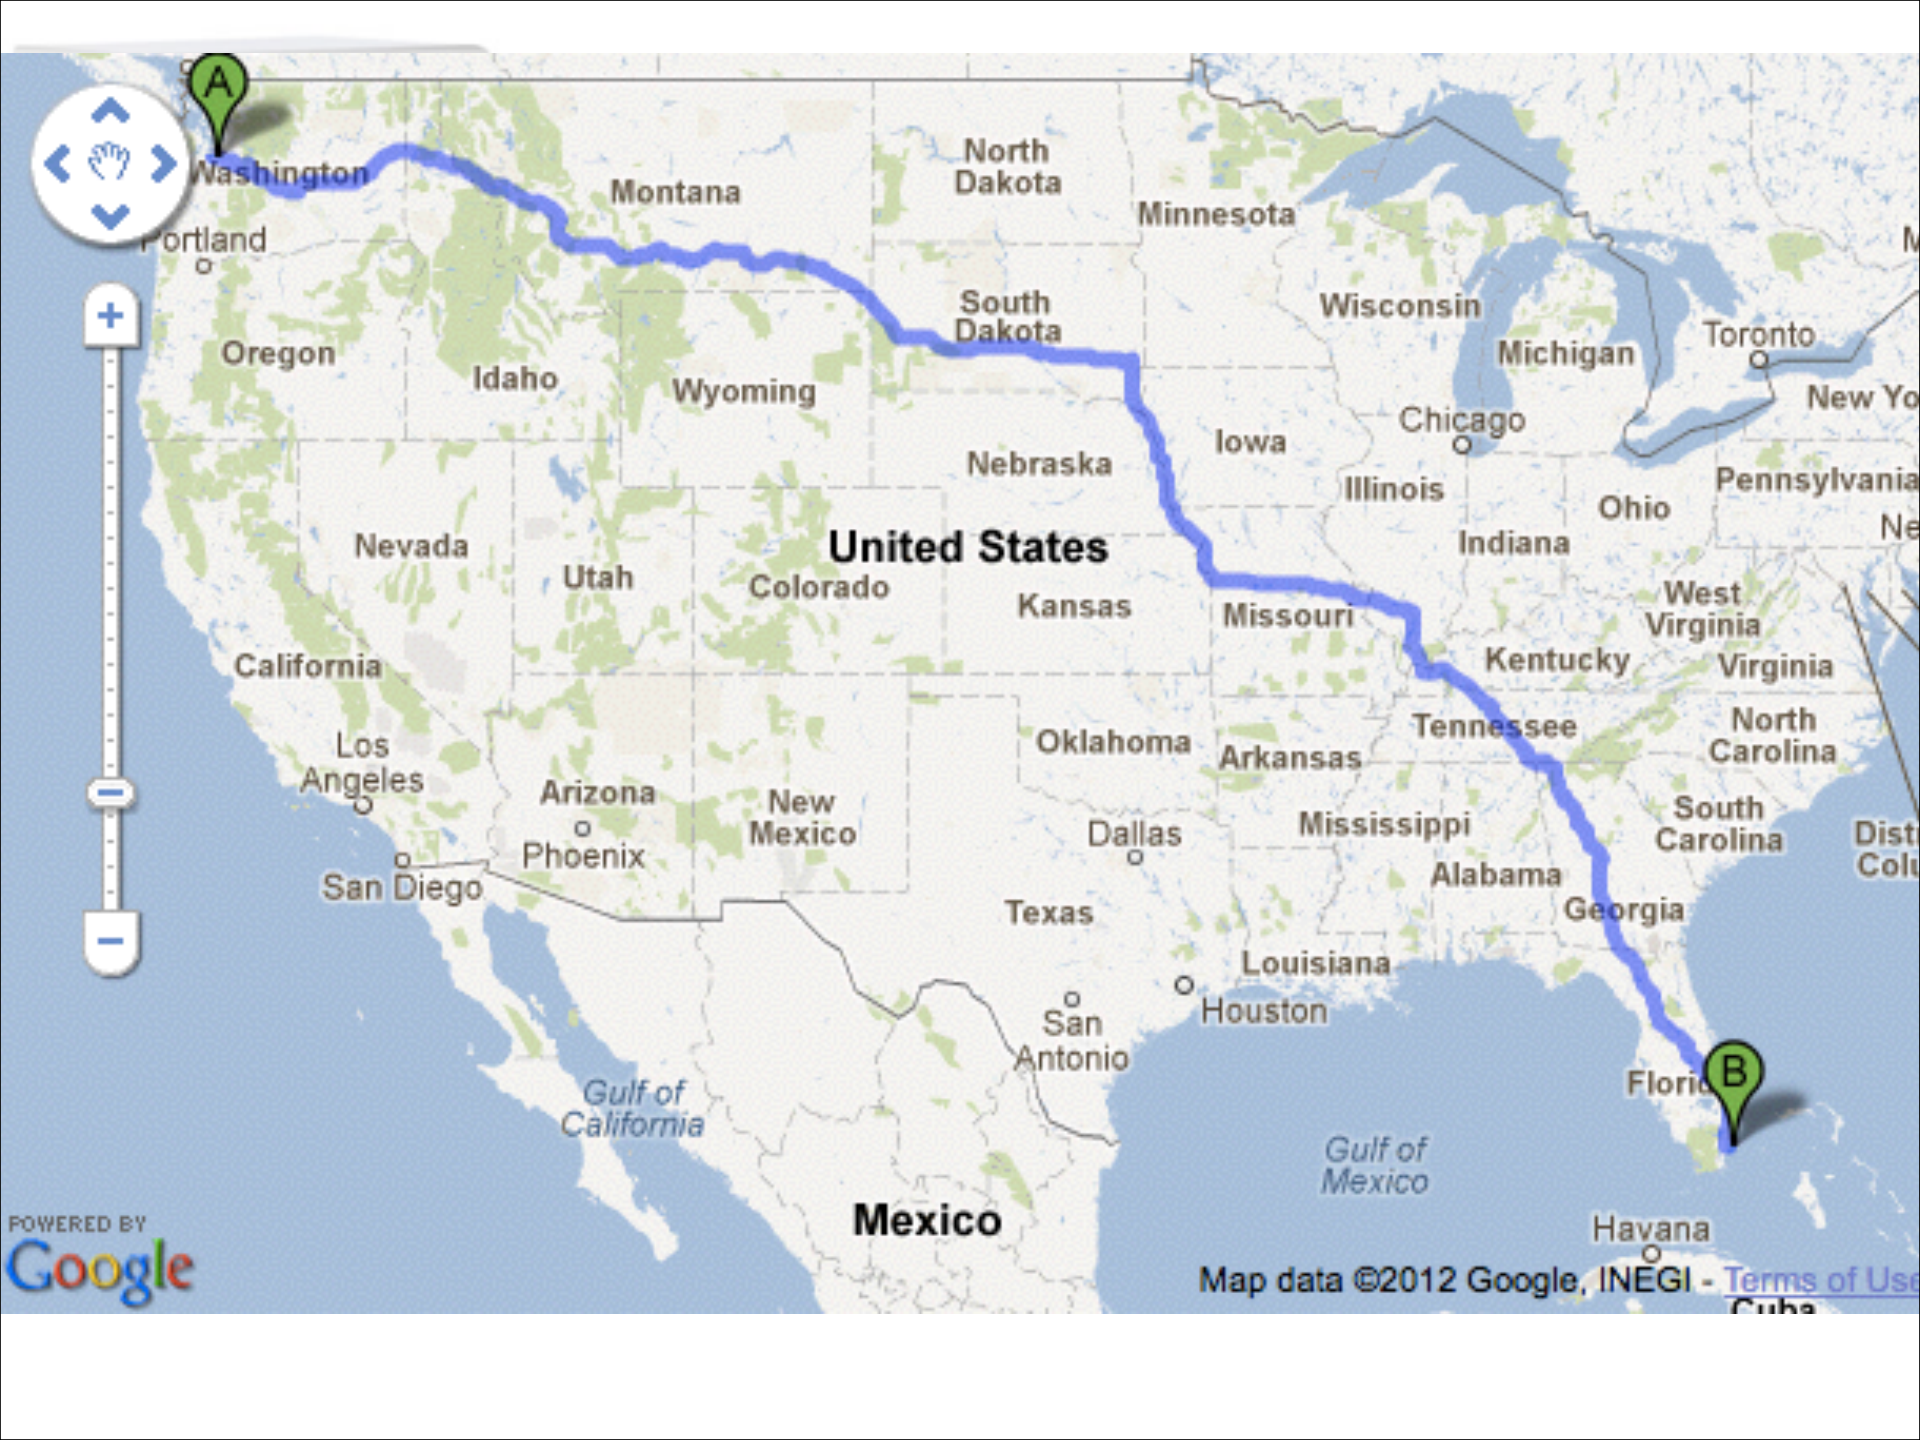
PERFORMANCE

REMINDER

main memory read	0.0001 ms (100 ns)
<i>network round trip</i>	<i>0.5 ms (500,000 ns)</i>
disk seek	10 ms (10,000,000 ns)







Washington

Montana

North Dakota

Minnesota

Portland

Oregon

Idaho

Wyoming

South Dakota

Wisconsin

Michigan

Toronto

New York

Nevada

Utah

Colorado

United States

Nebraska

Iowa

Chicago

Illinois

Ohio

Pennsylvania

California

Los Angeles

Arizona

Phoenix

New Mexico

Kansas

Missouri

Indiana

West Virginia

Virginia

North Carolina

Oklahoma

Arkansas

Tennessee

South Carolina

San Diego

Texas

San Antonio

Louisiana

Houston

Alabama

Georgia

Gulf of California

Mexico

Gulf of Mexico


Florida

Havana

POWERED BY
Google

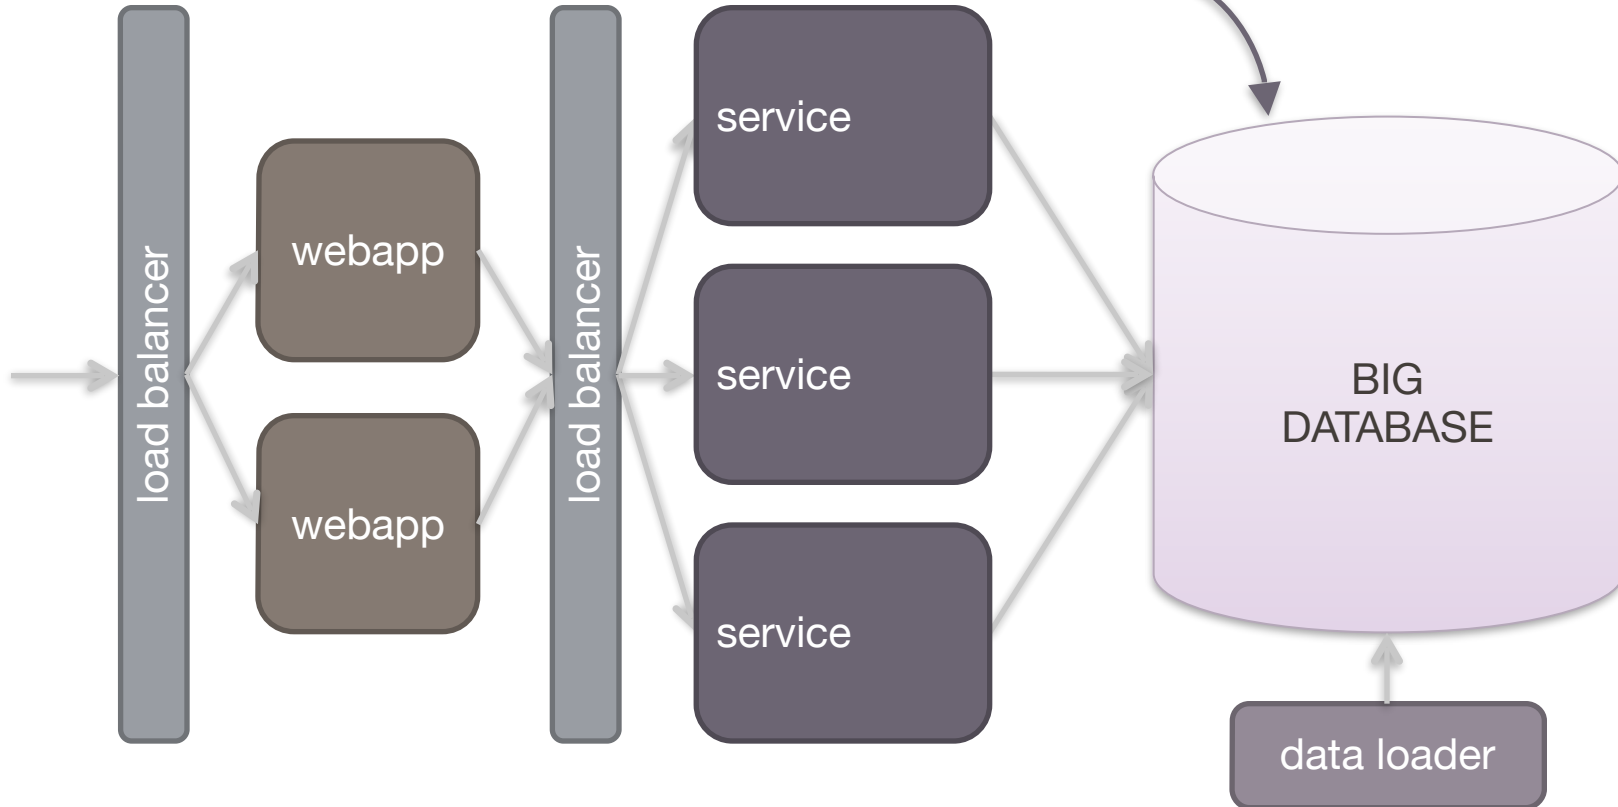
Map data ©2012 Google, INEGI - [Terms of Use](#)
Cuba

THE BEGINNING

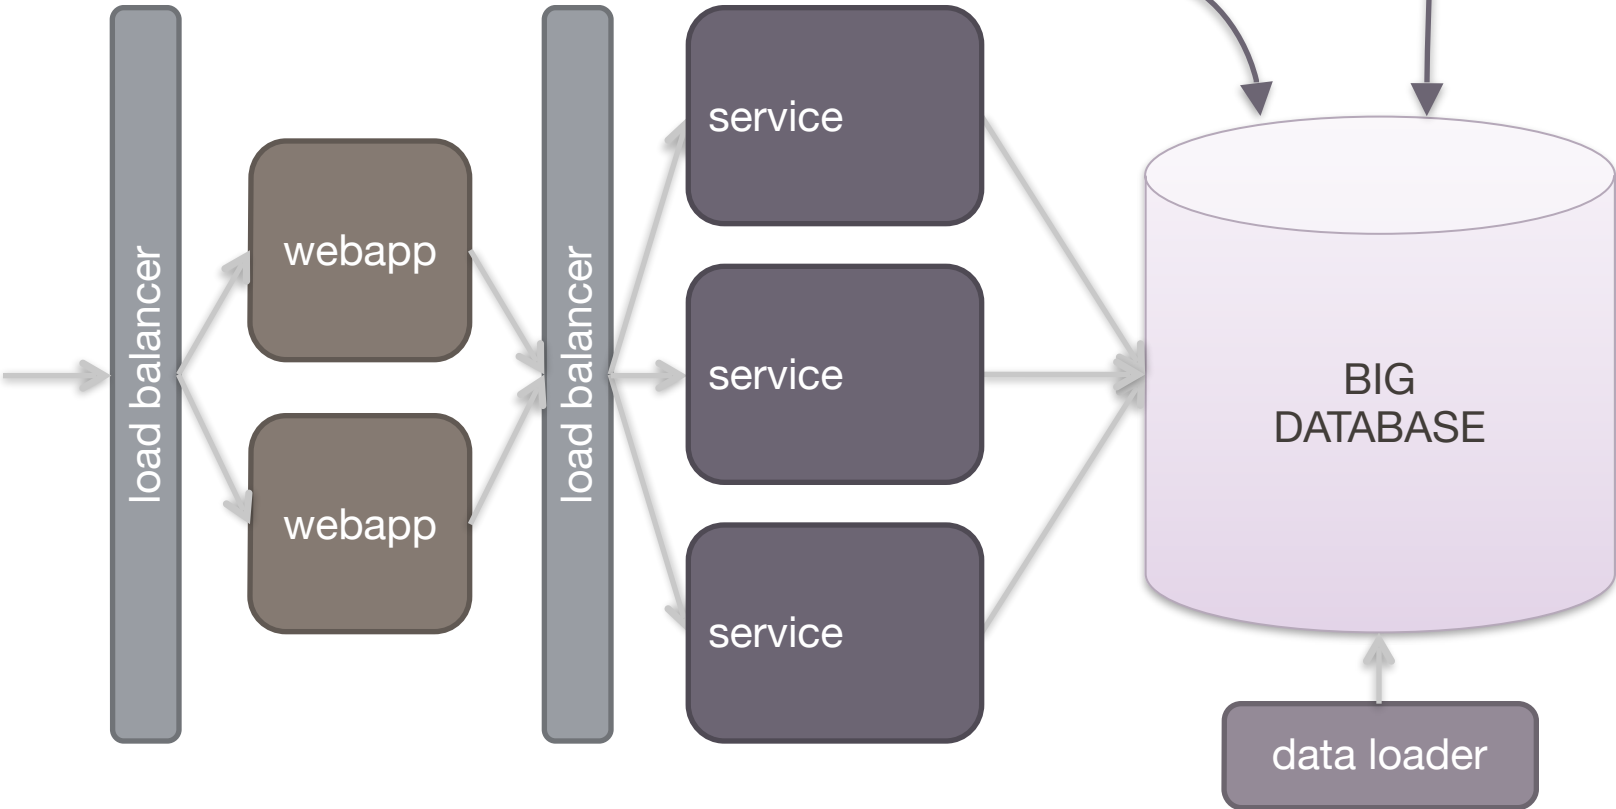


THE BEGINNING

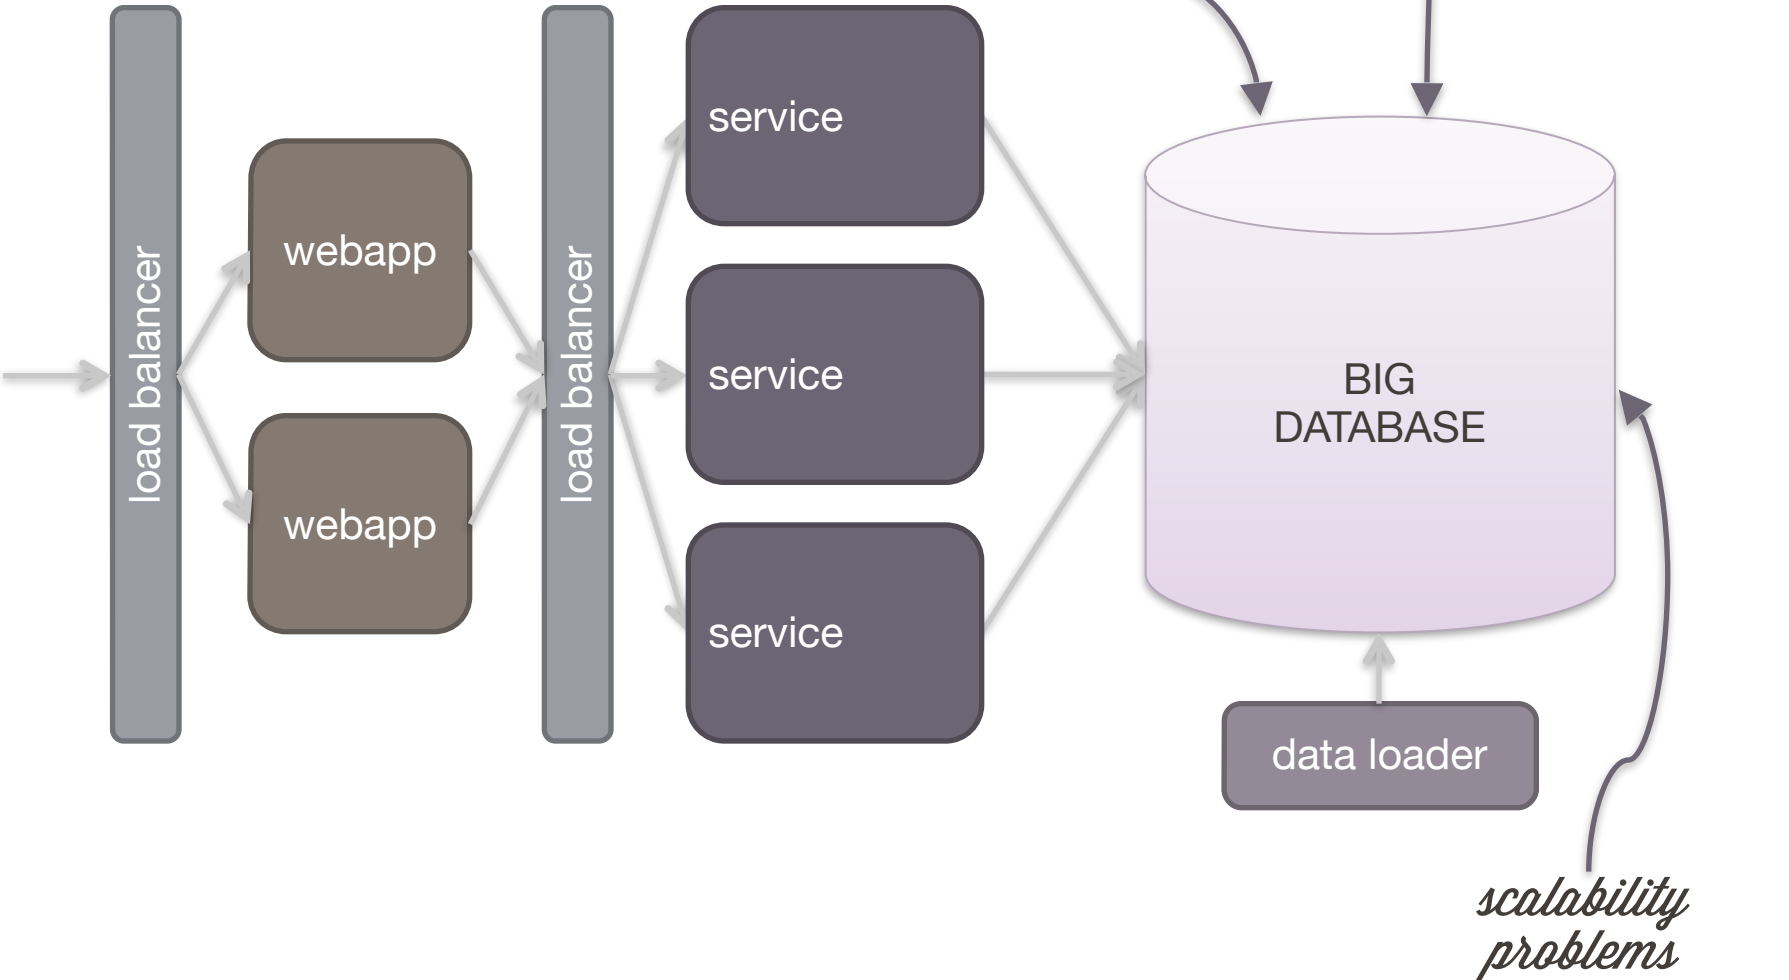
*availability
problems*




THE BEGINNING



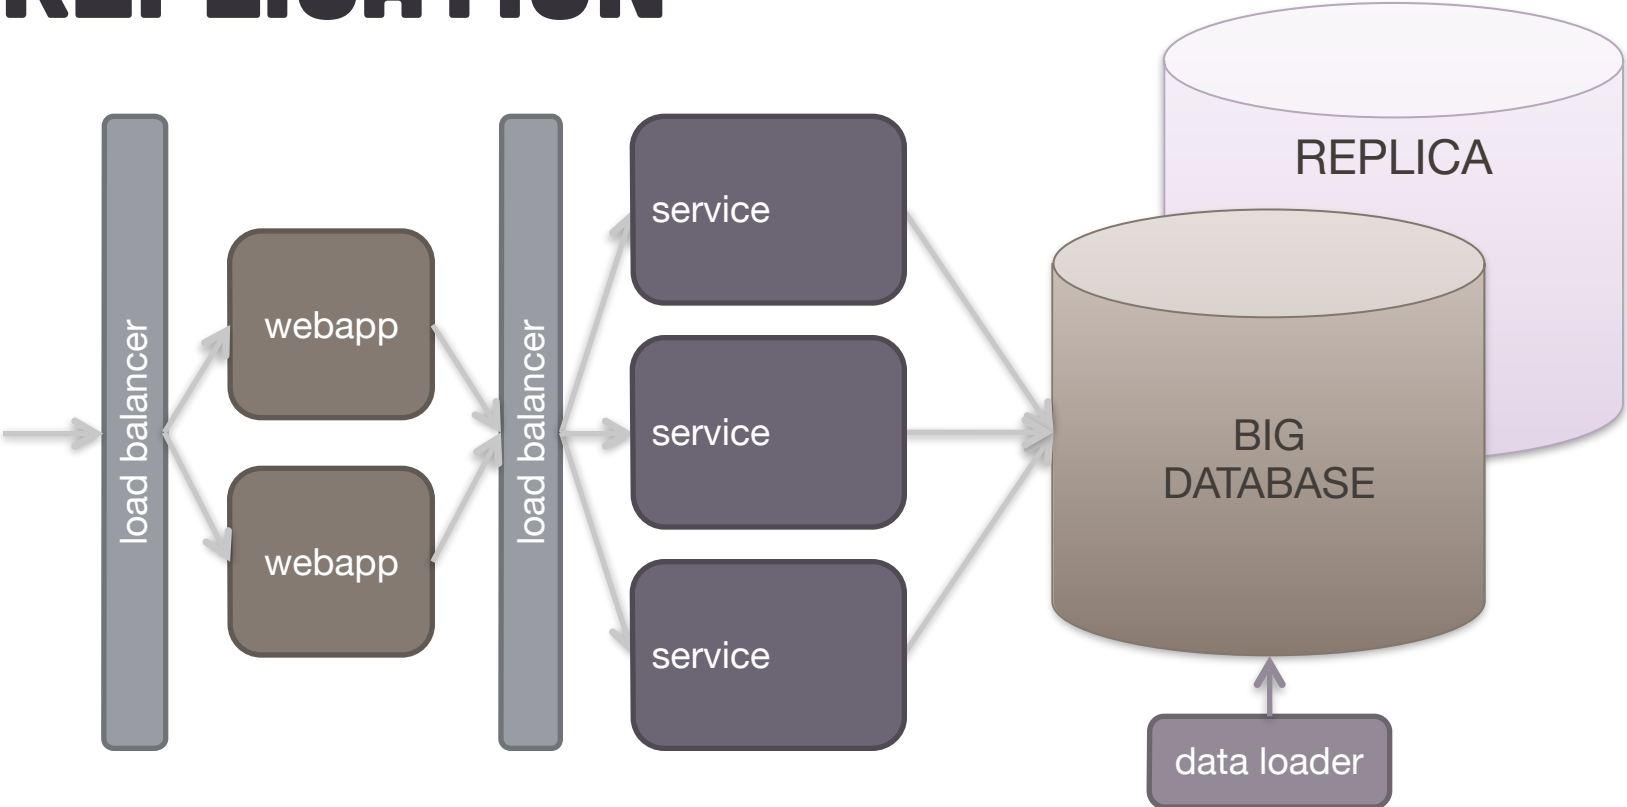
THE BEGINNING



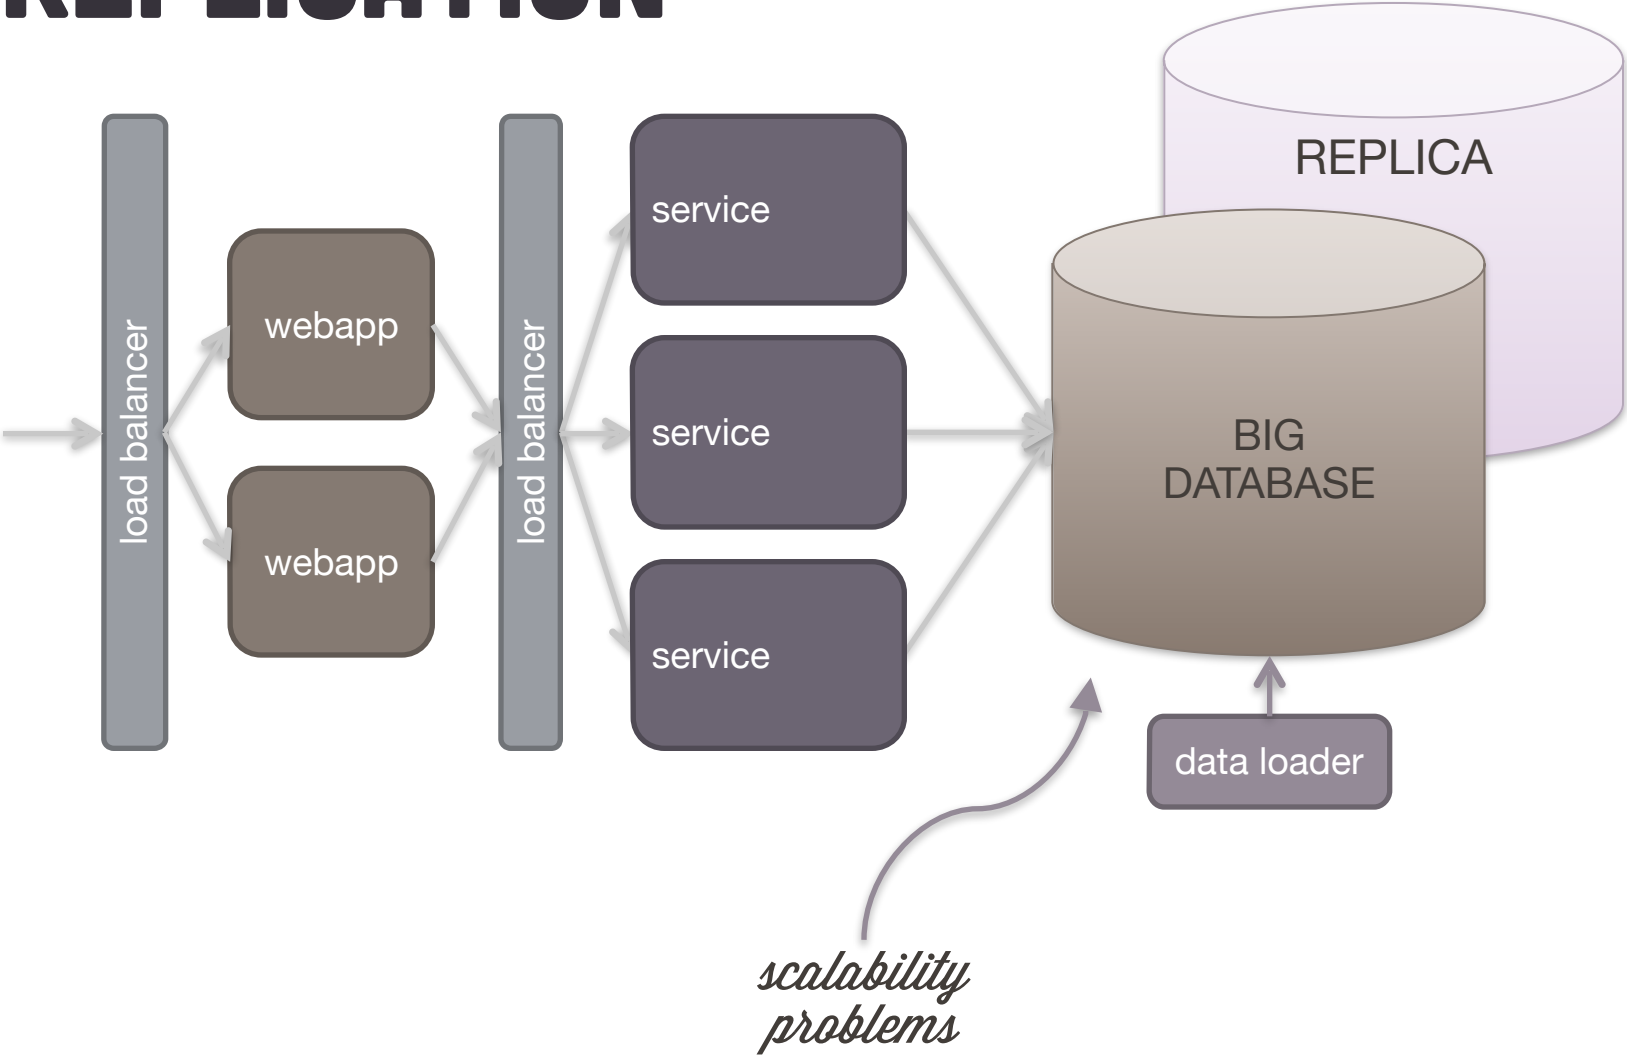
REPLICATION



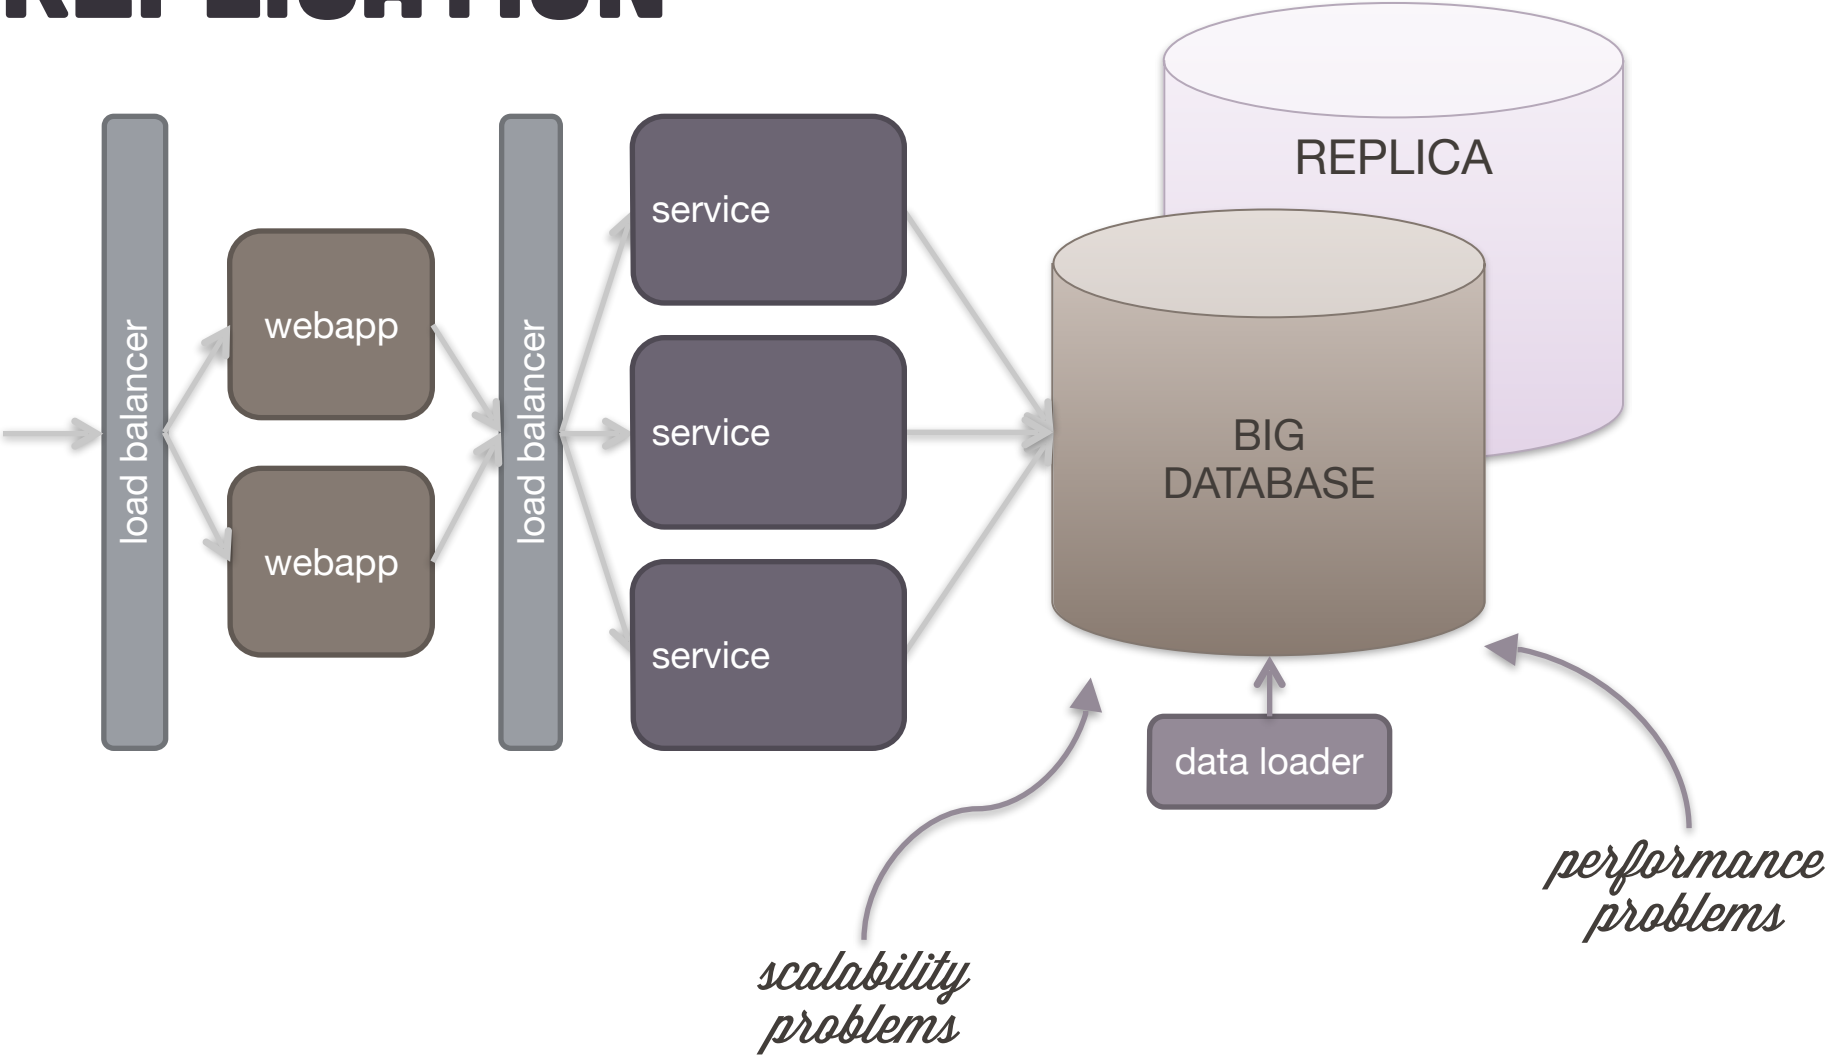
REPLICATION



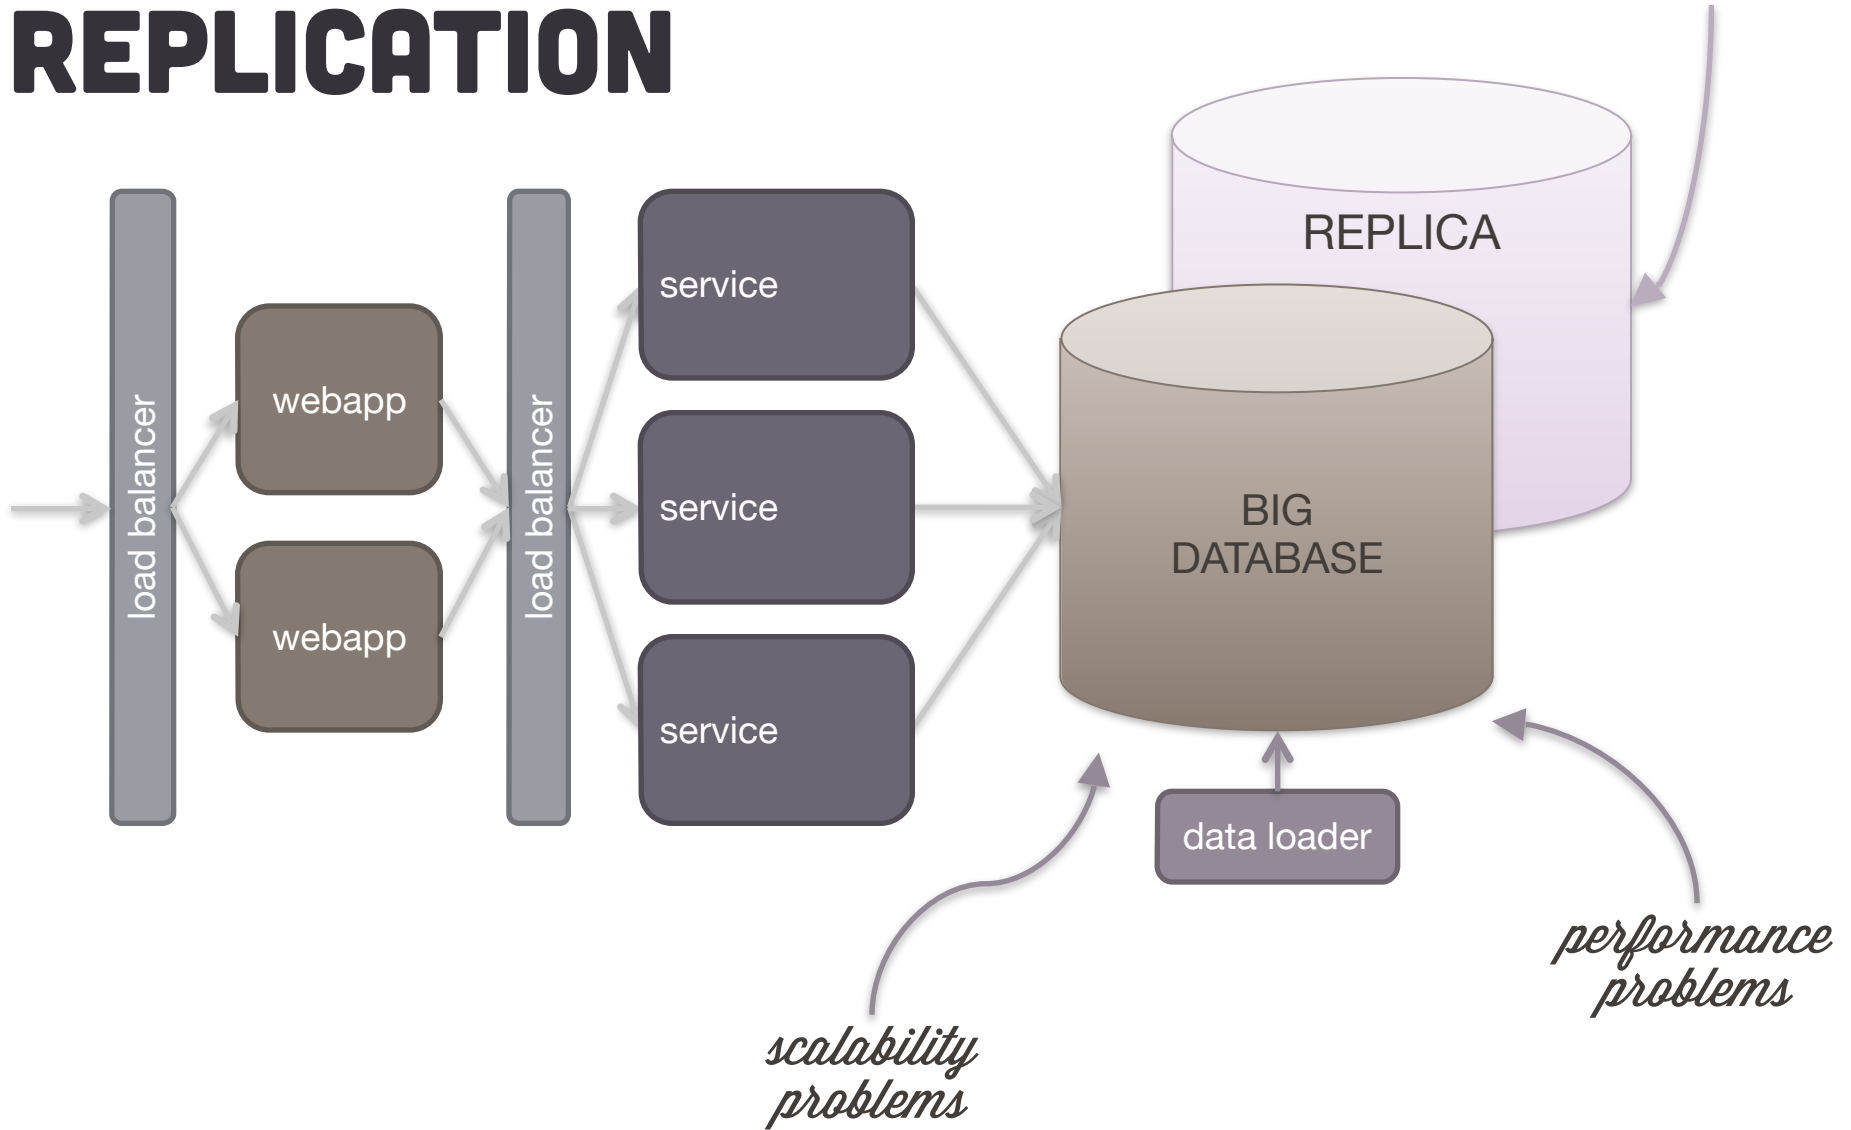
REPLICATION



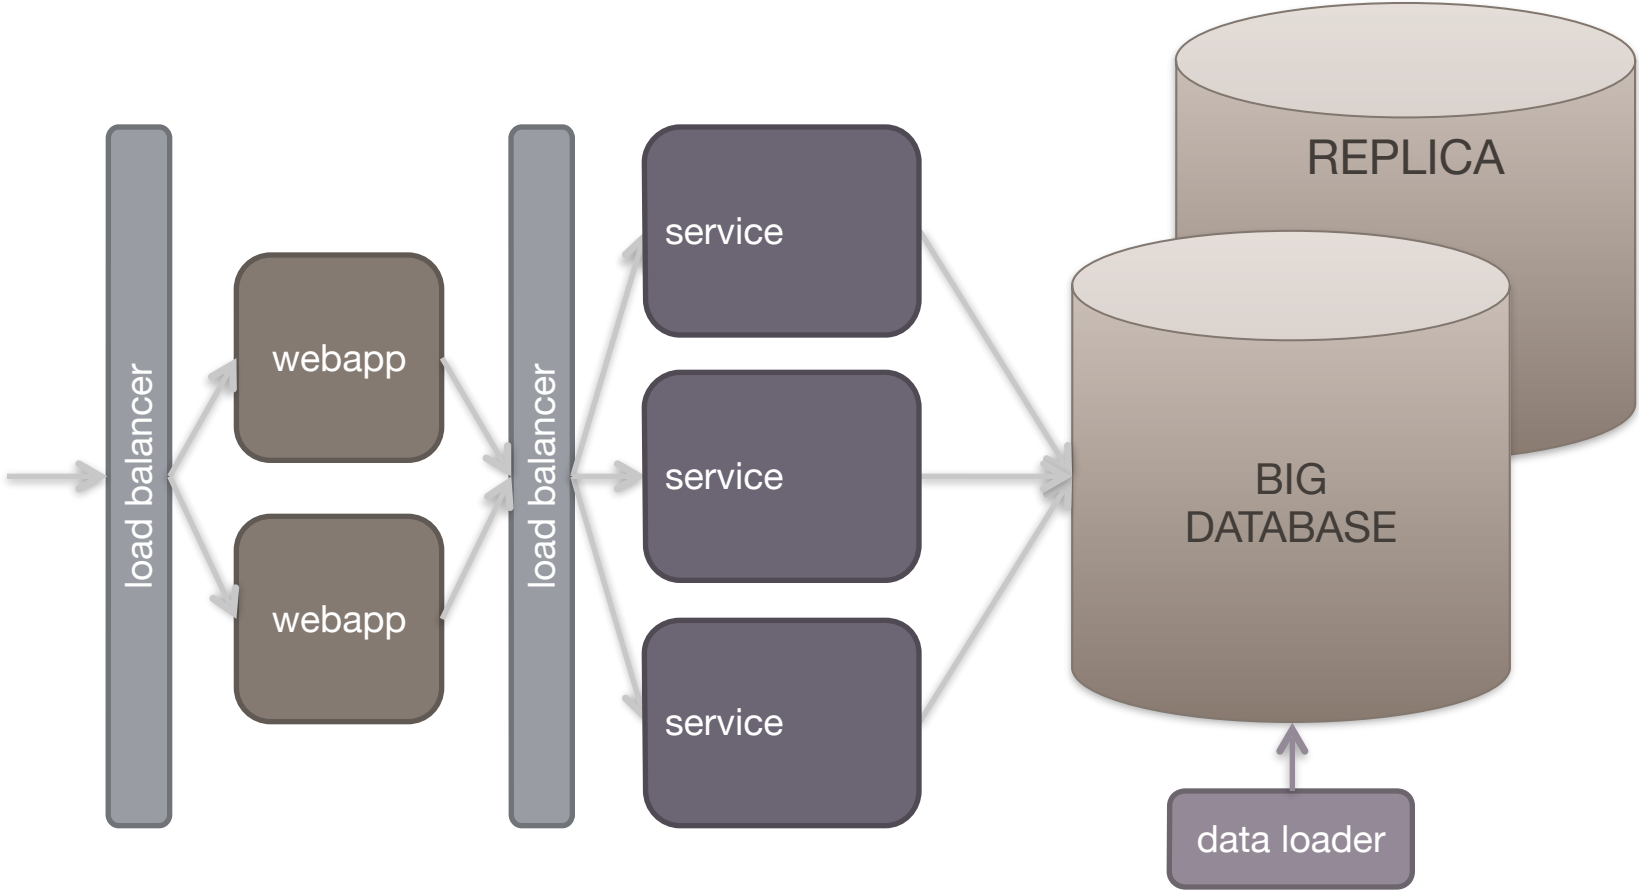
REPLICATION



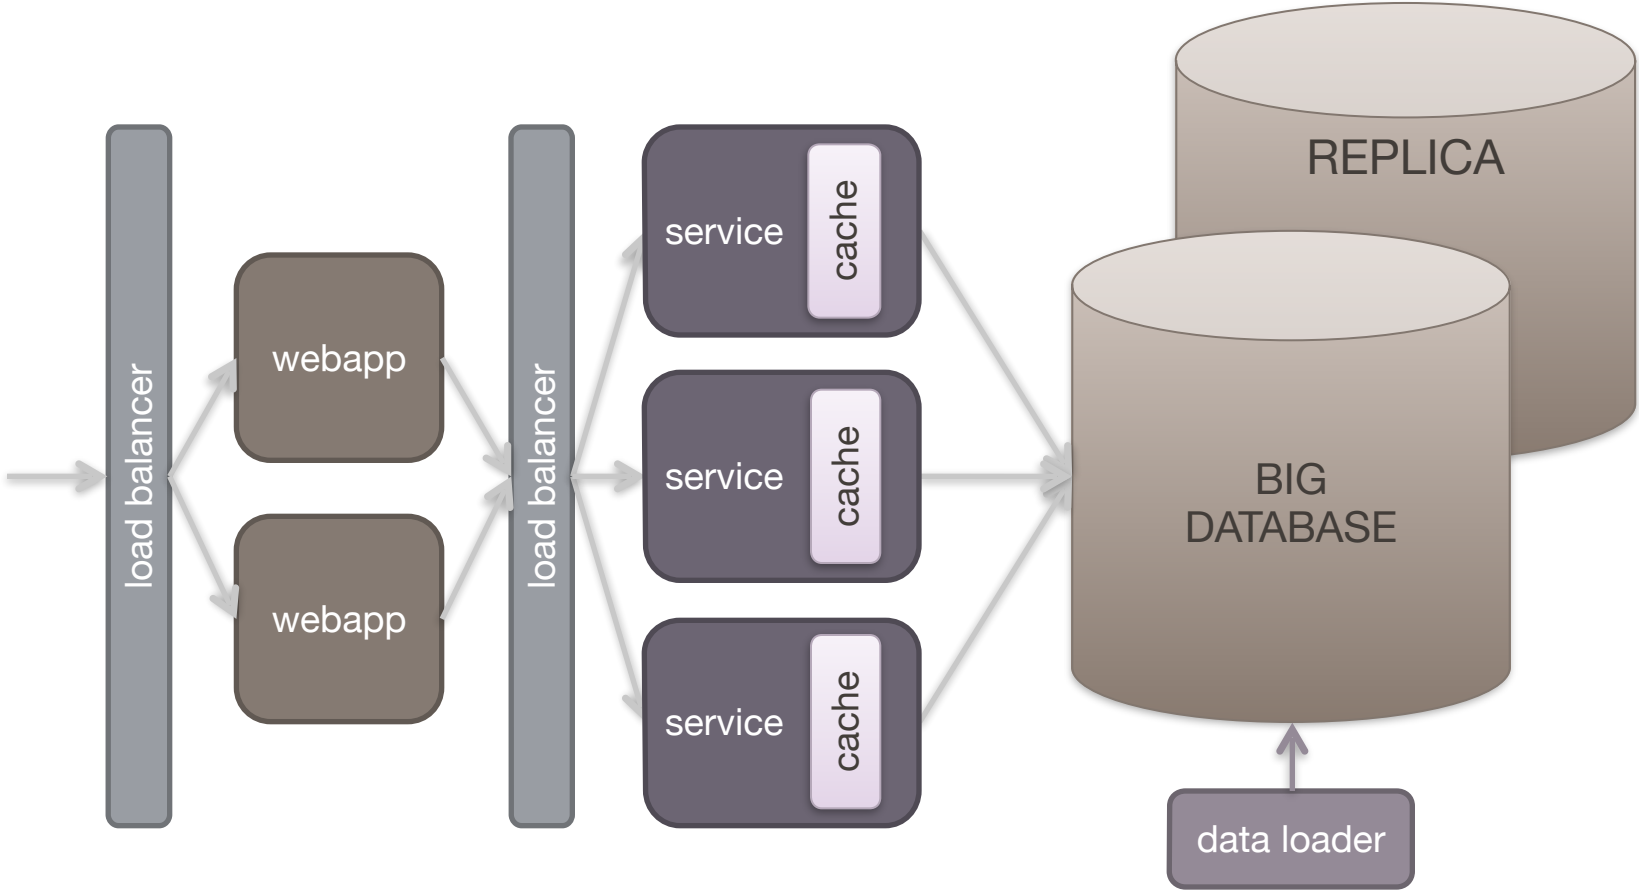
REPLICATION



LOCAL CACHING

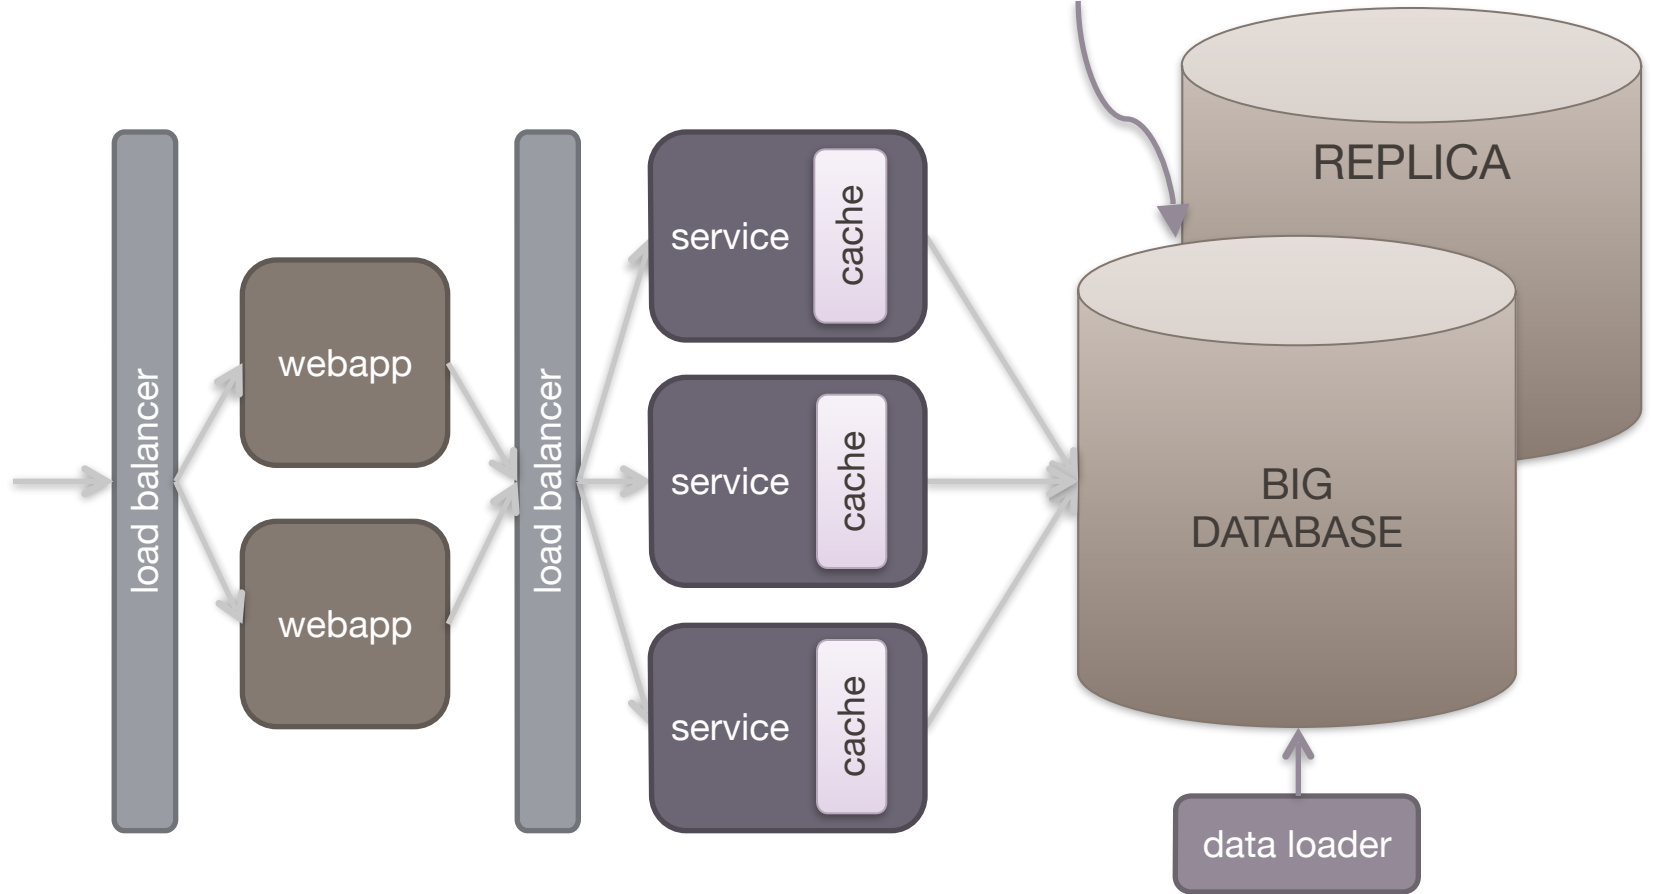


LOCAL CACHING




LOCAL CACHING

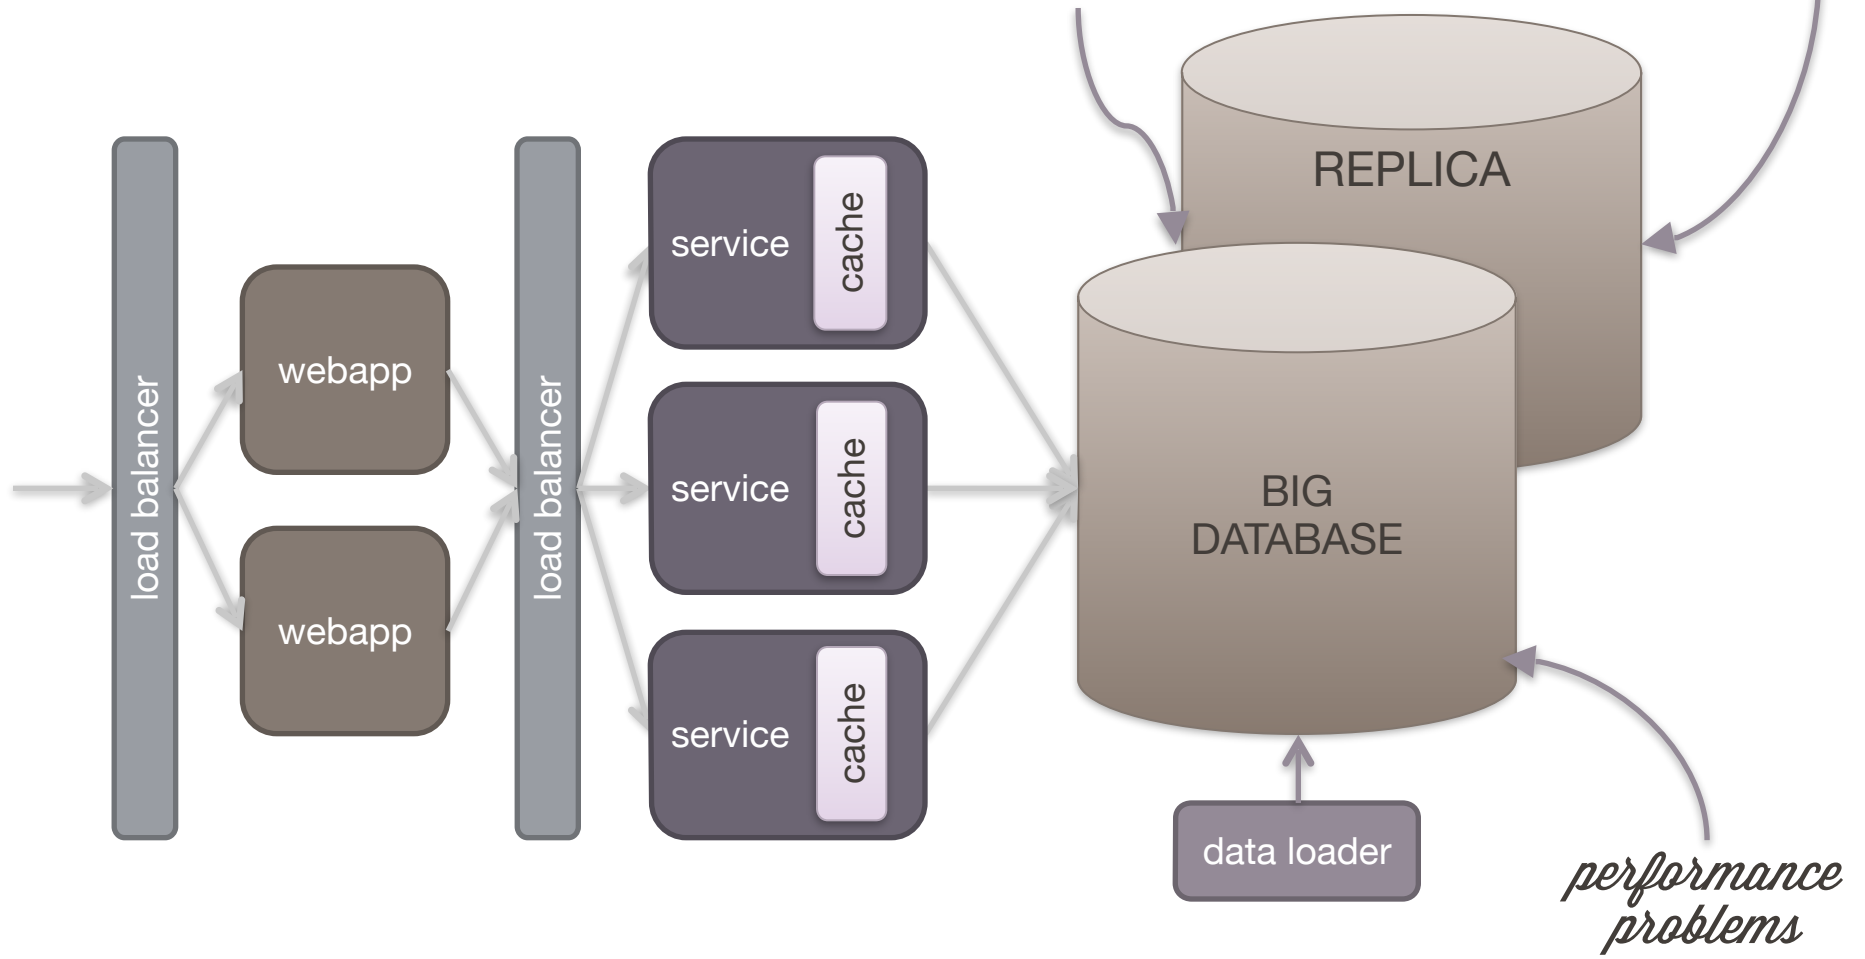
scalability problems



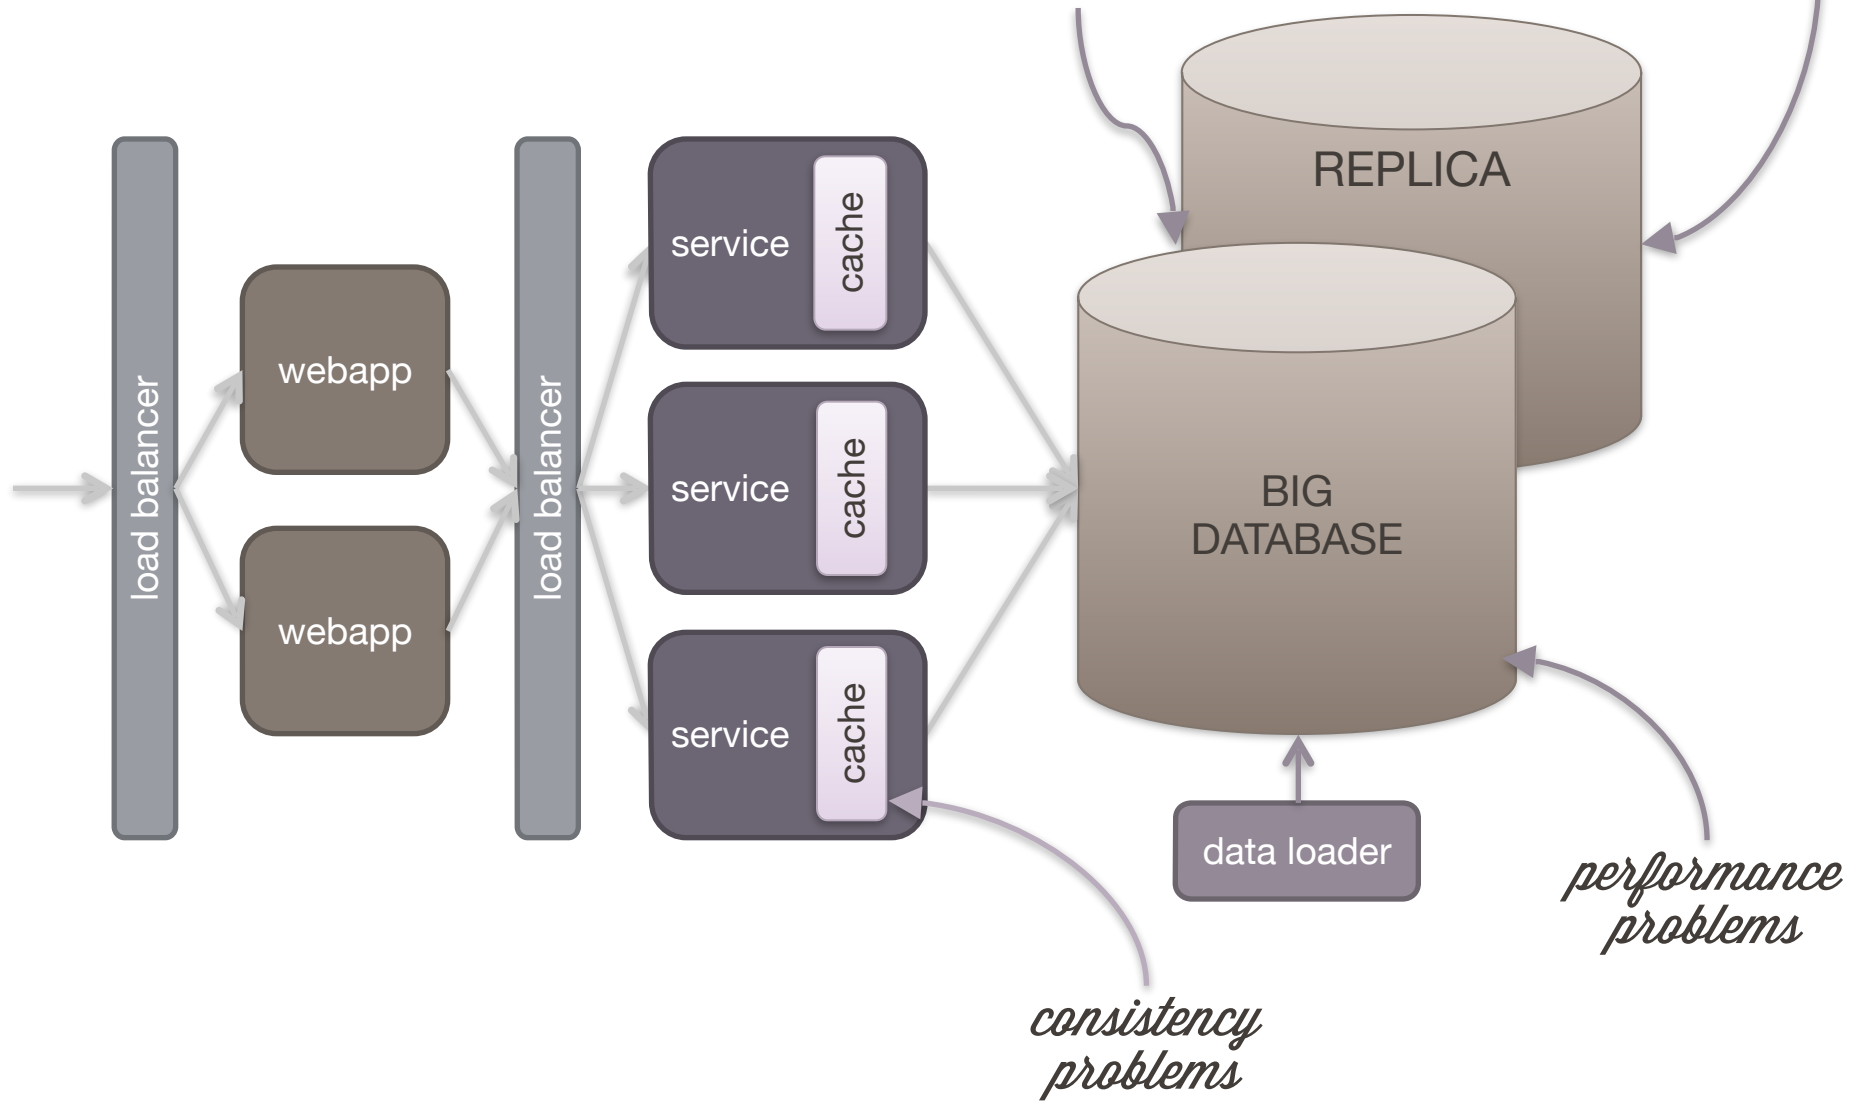
LOCAL CACHING



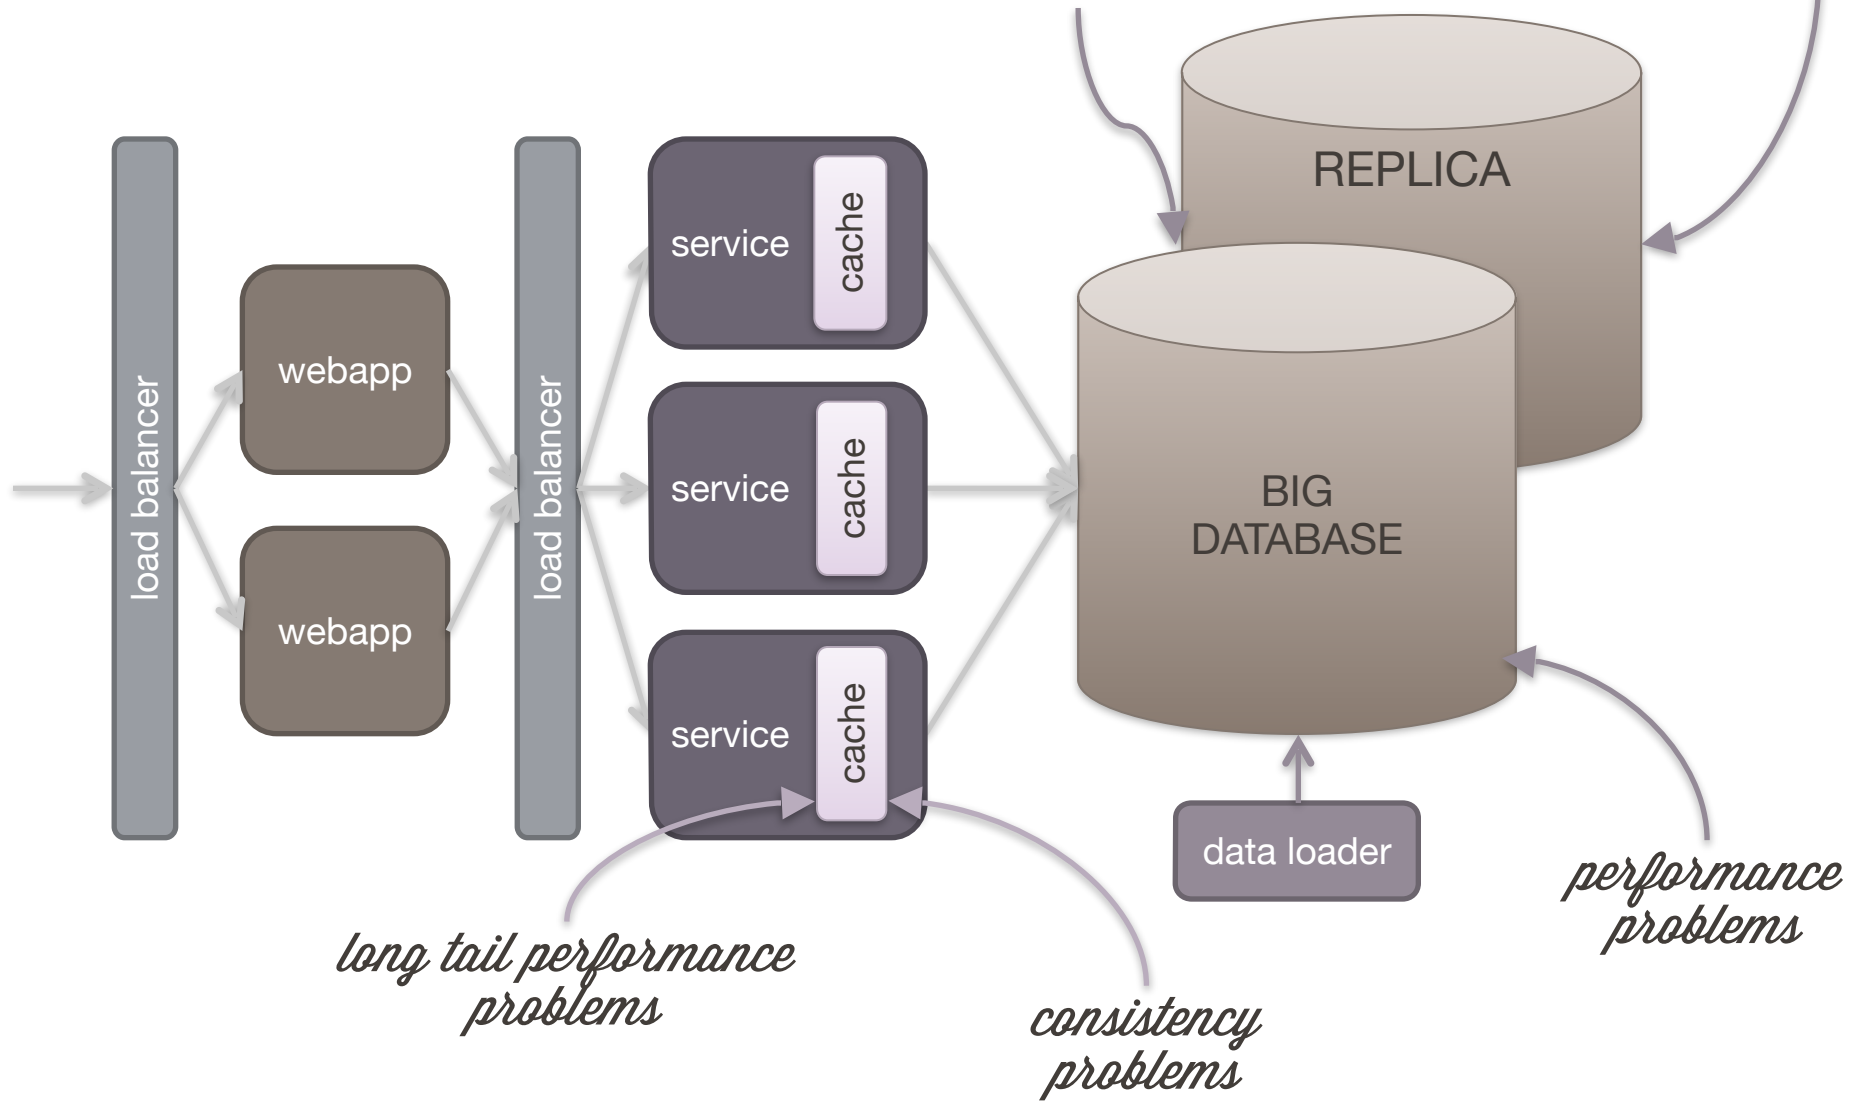
LOCAL CACHING



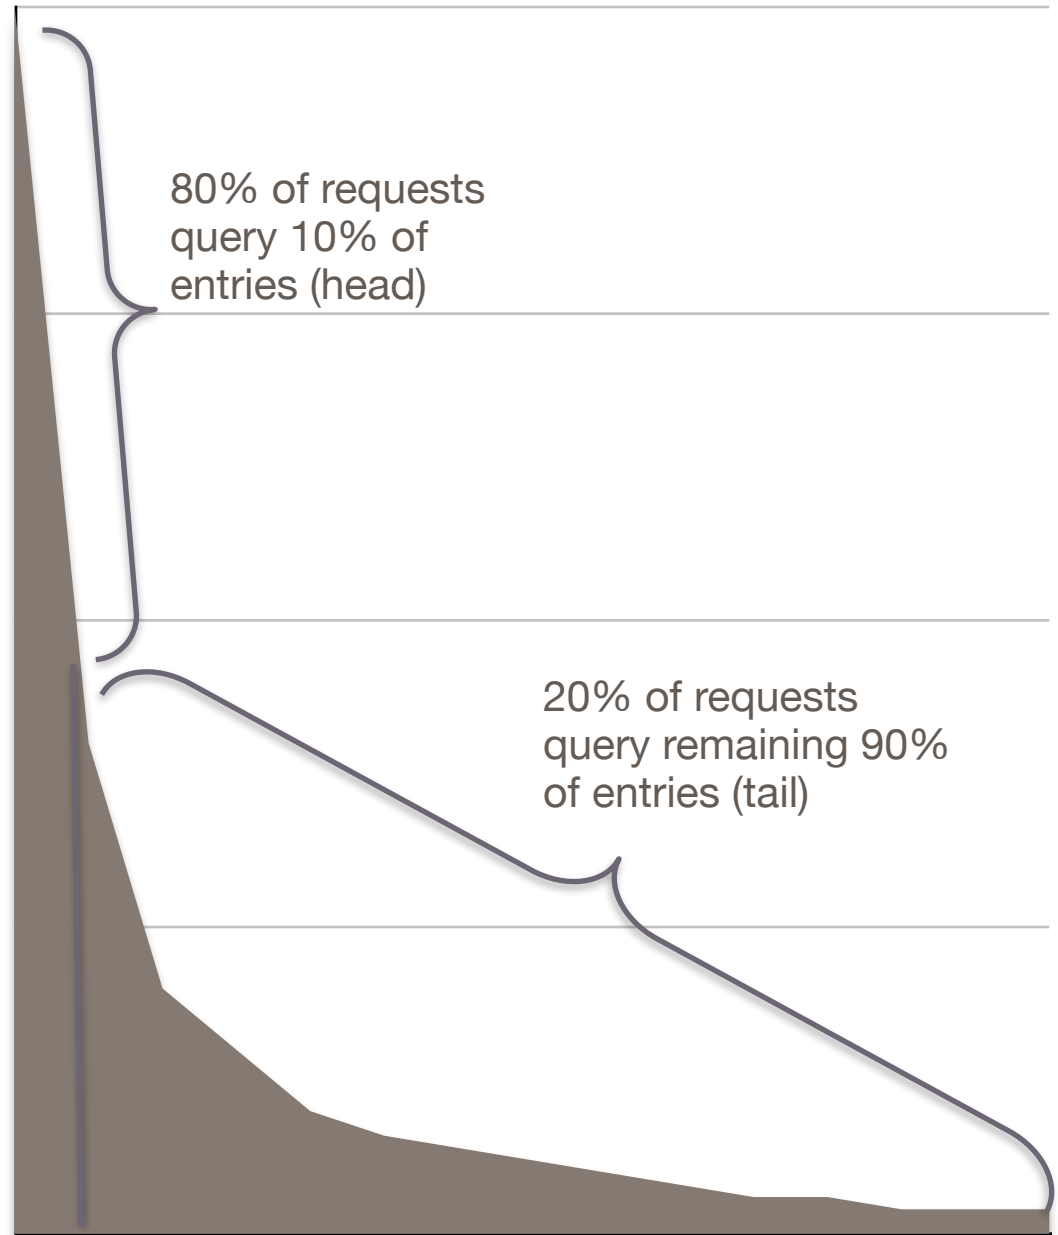
LOCAL CACHING



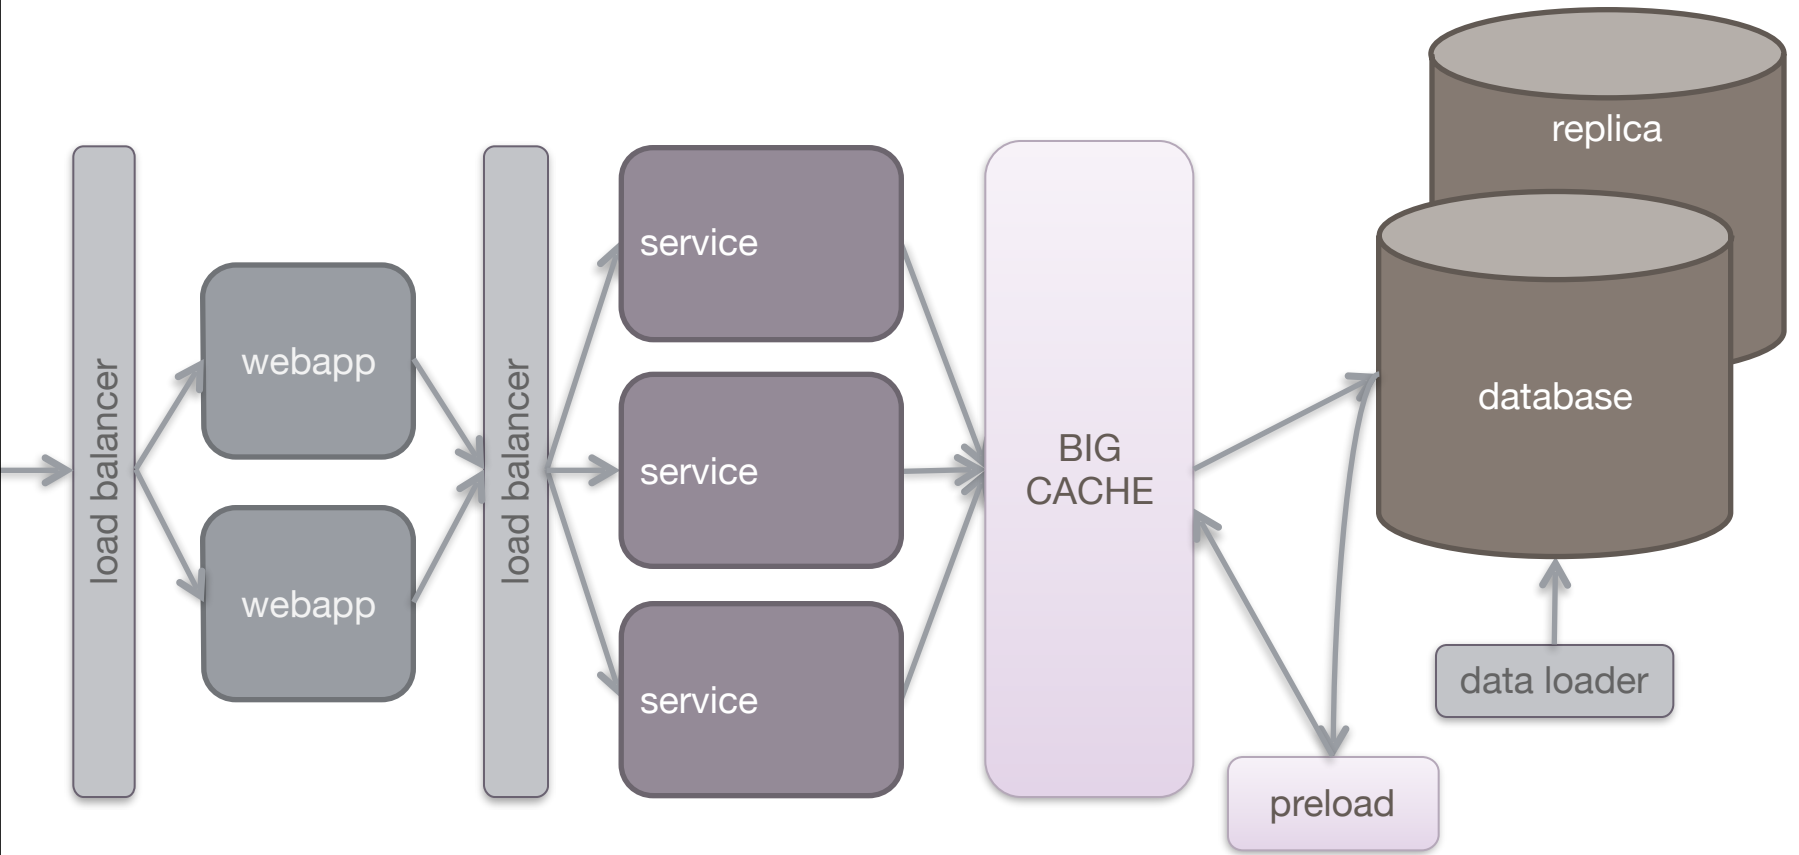
LOCAL CACHING



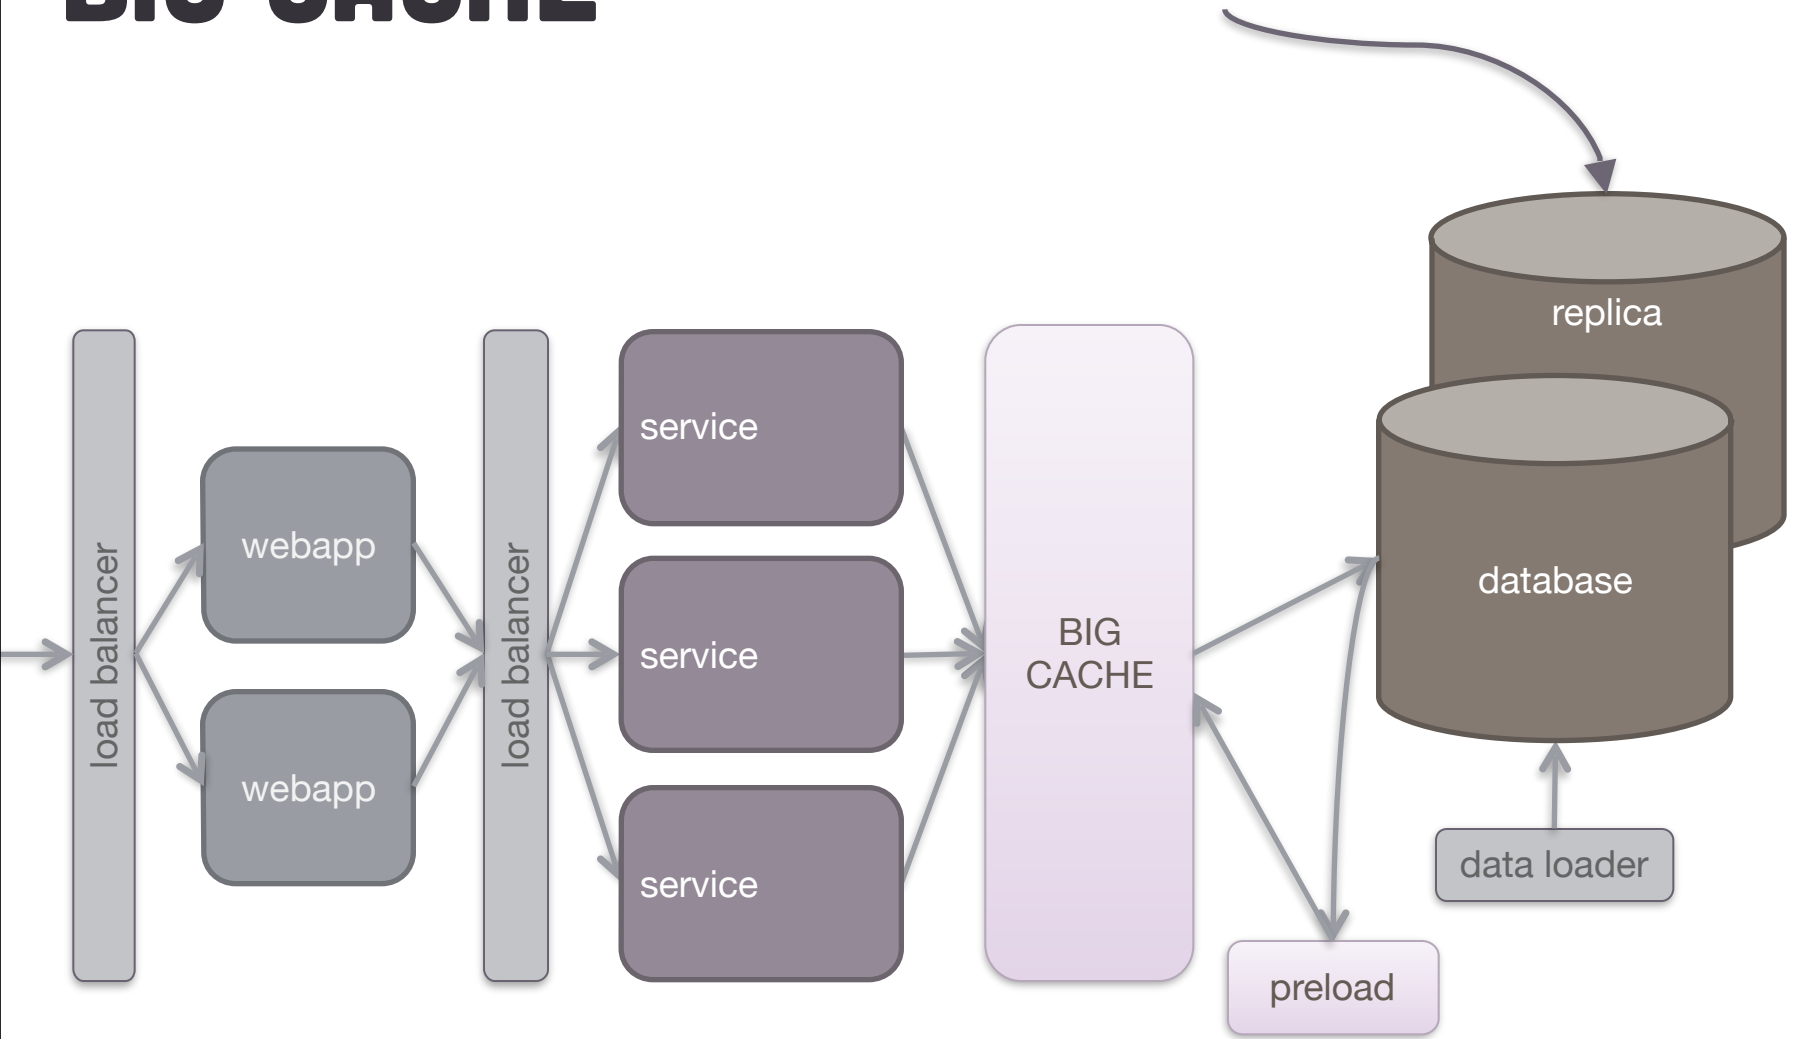
THE LONG TAIL PROBLEM



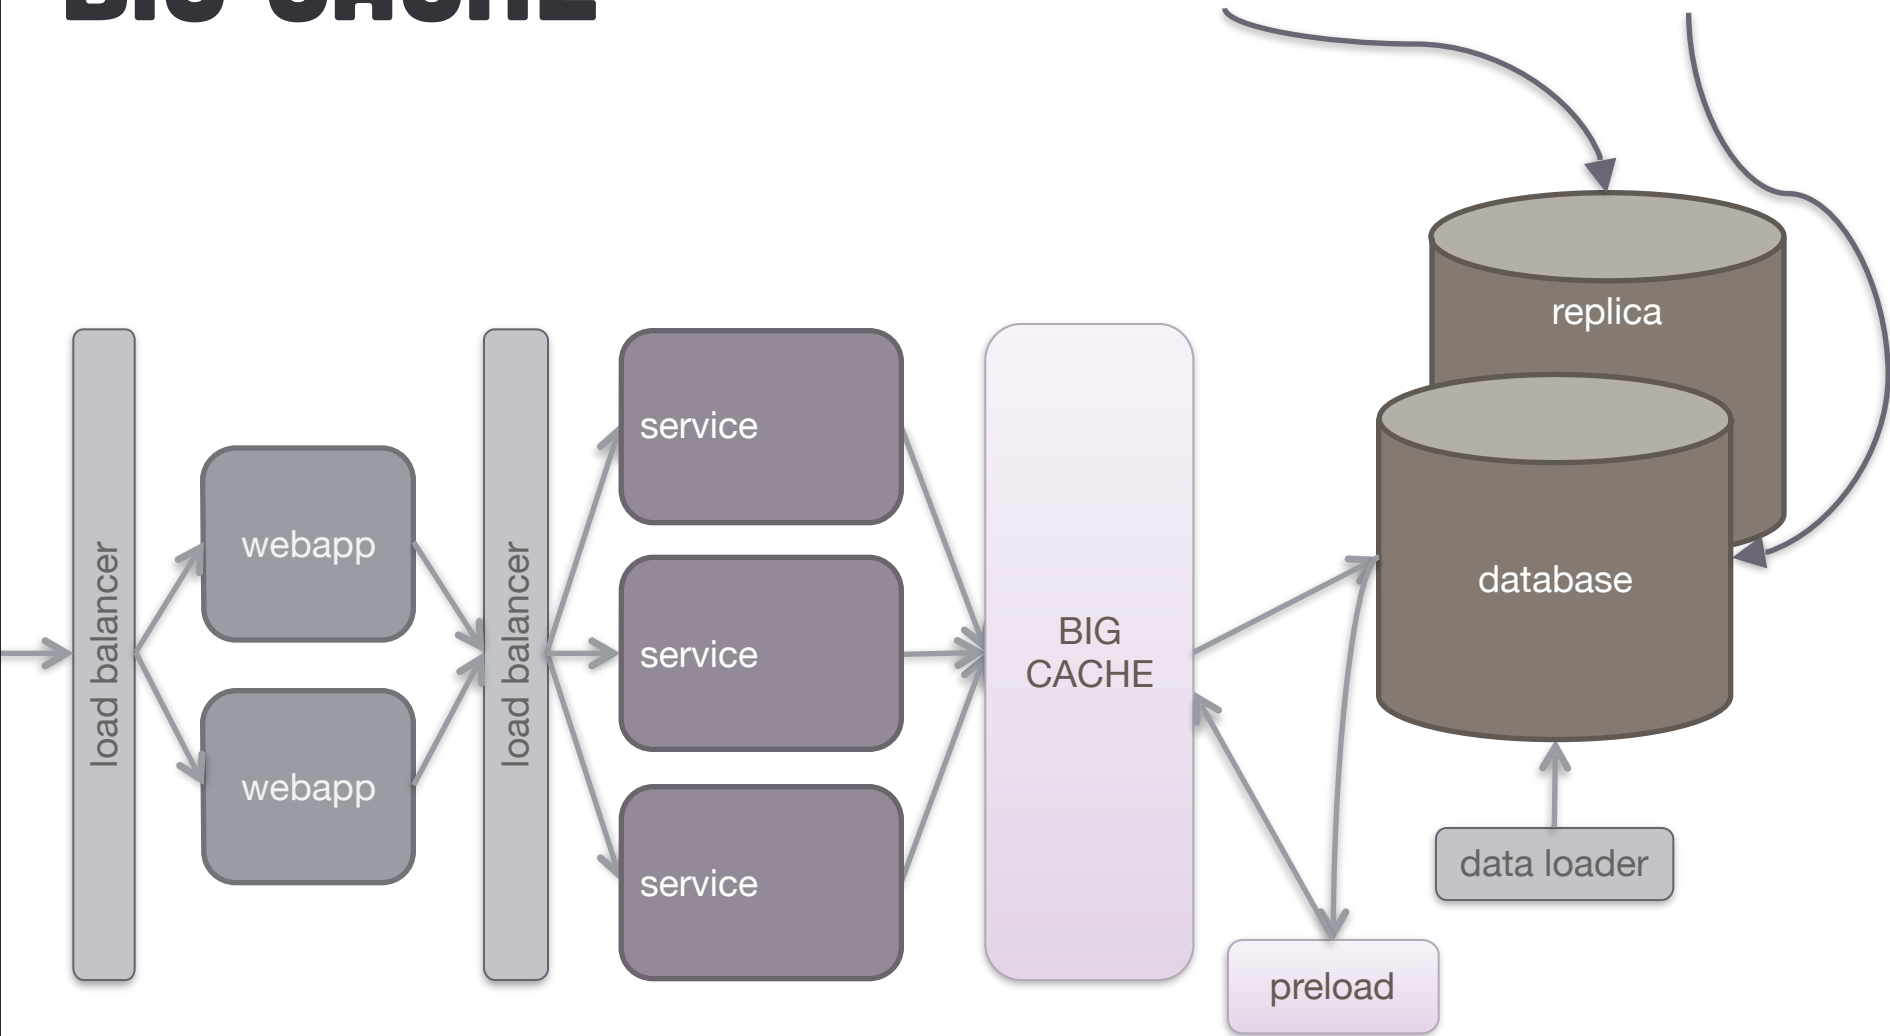
BIG CACHE



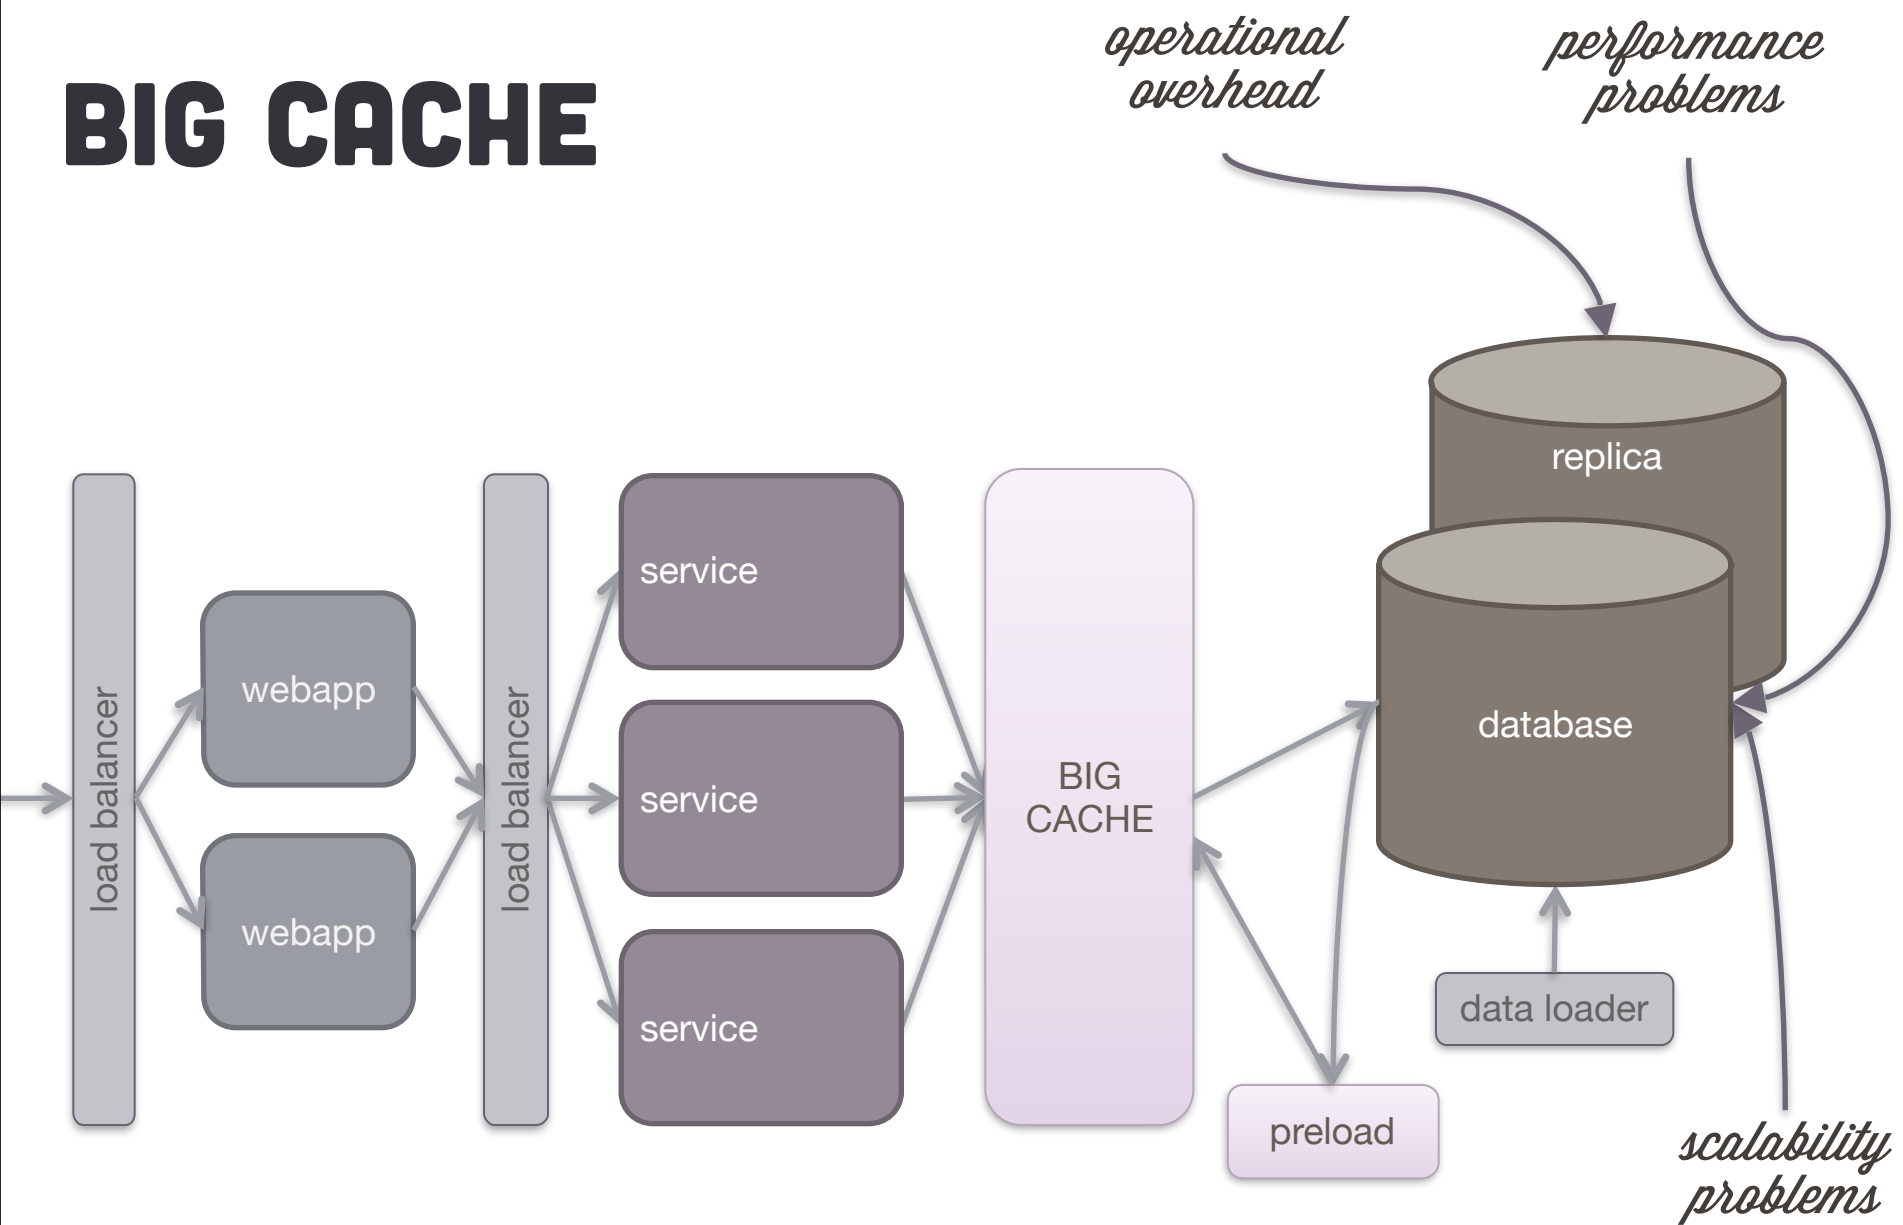
BIG CACHE



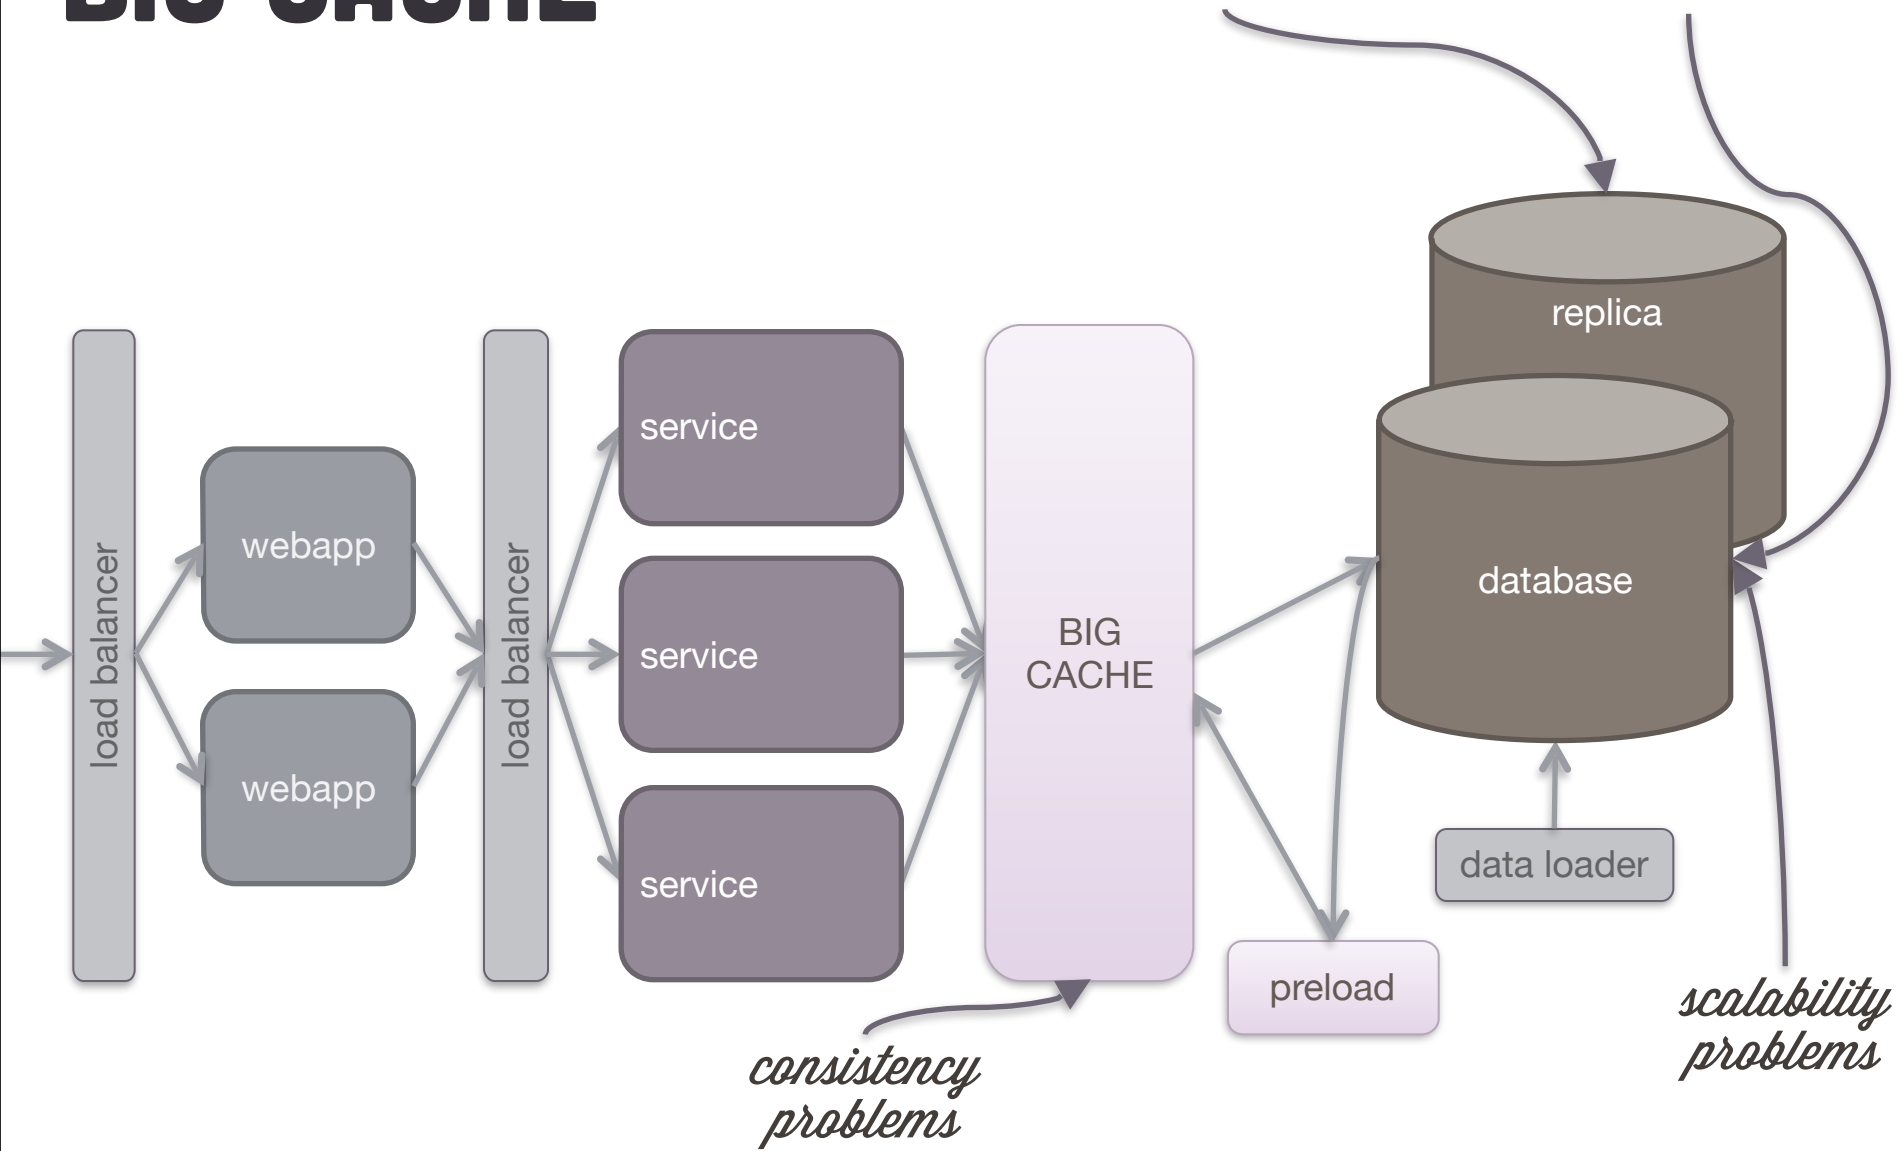
BIG CACHE



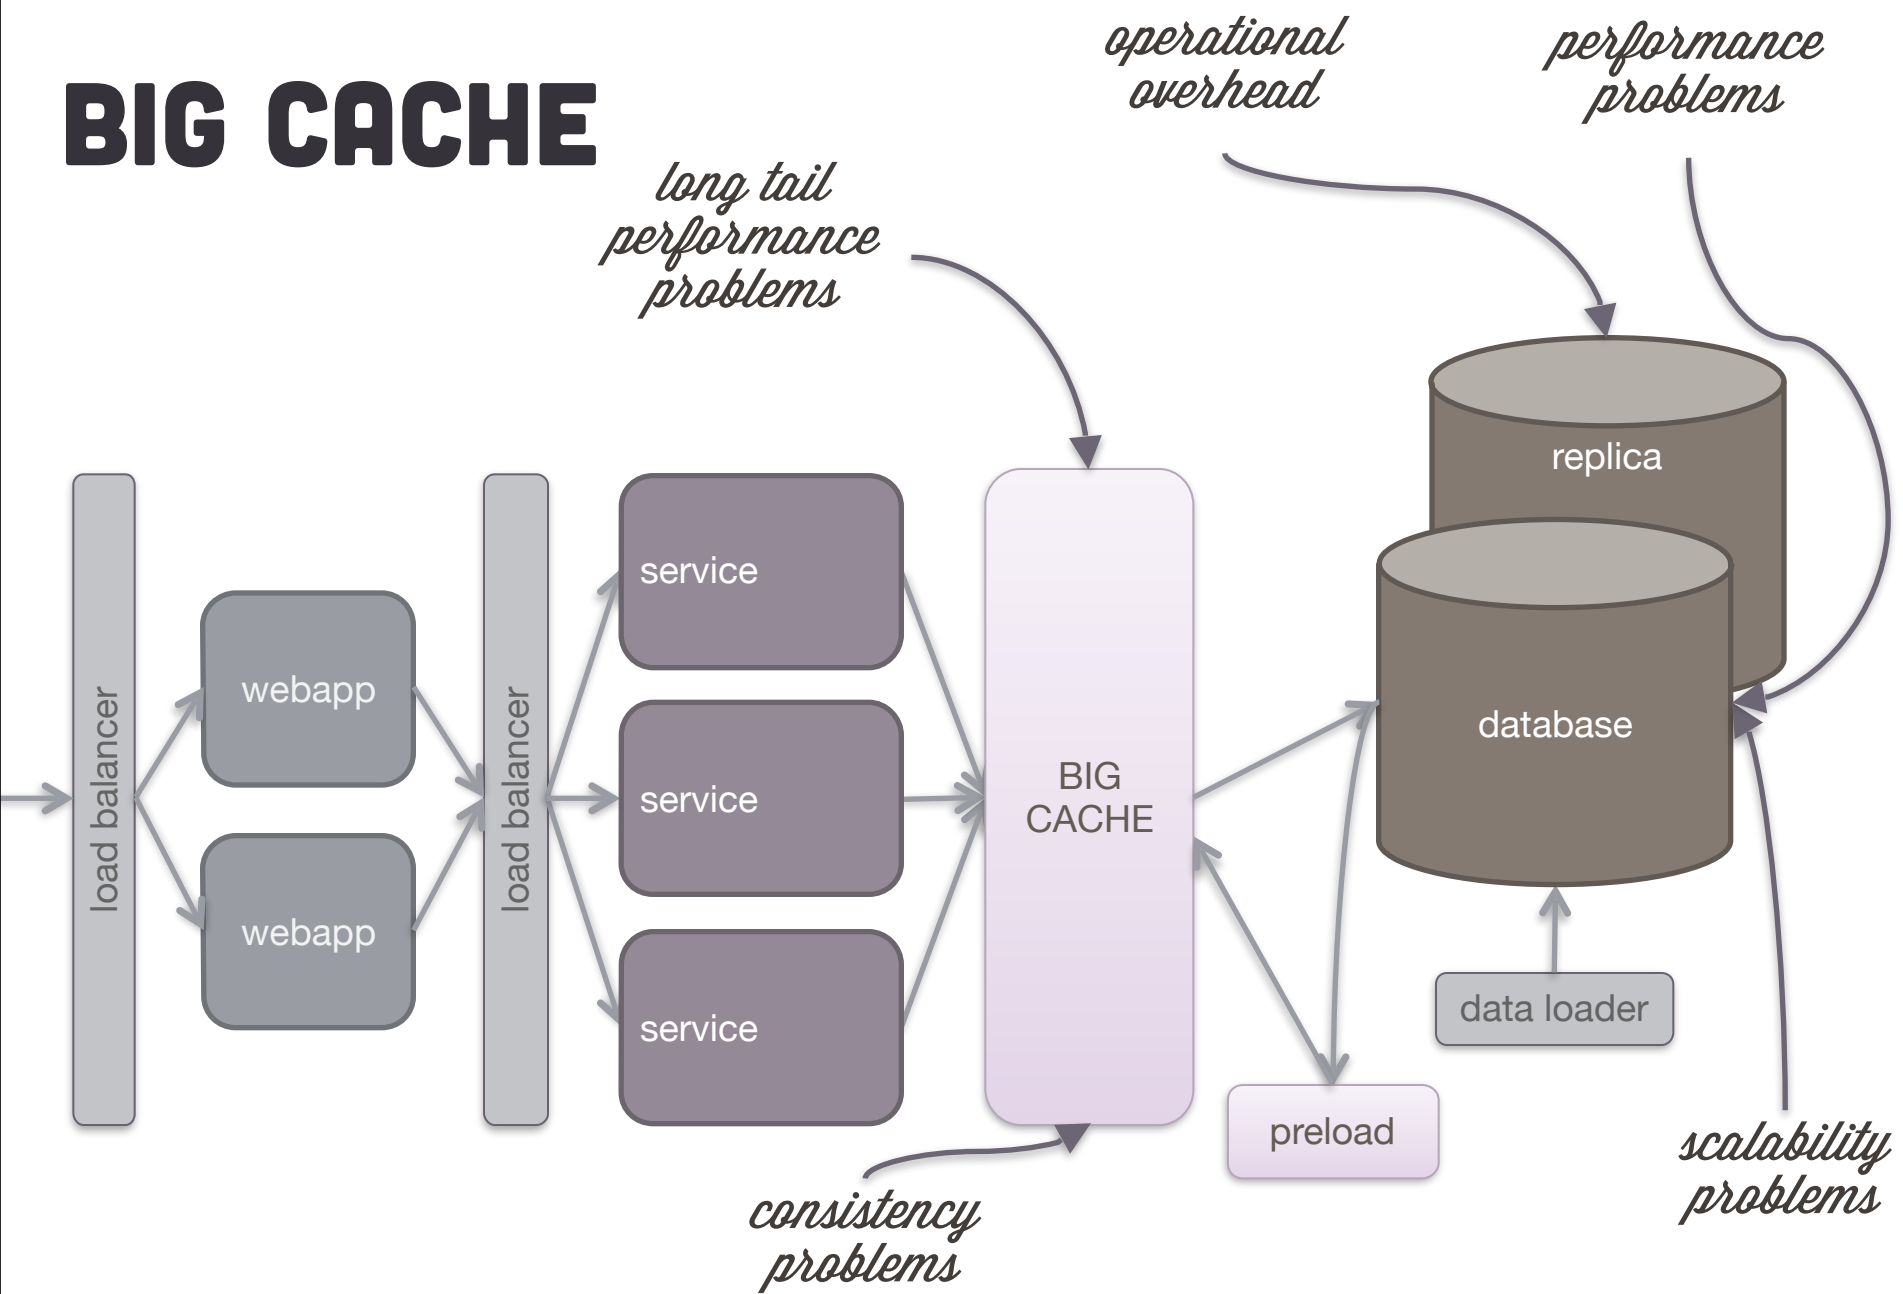
BIG CACHE



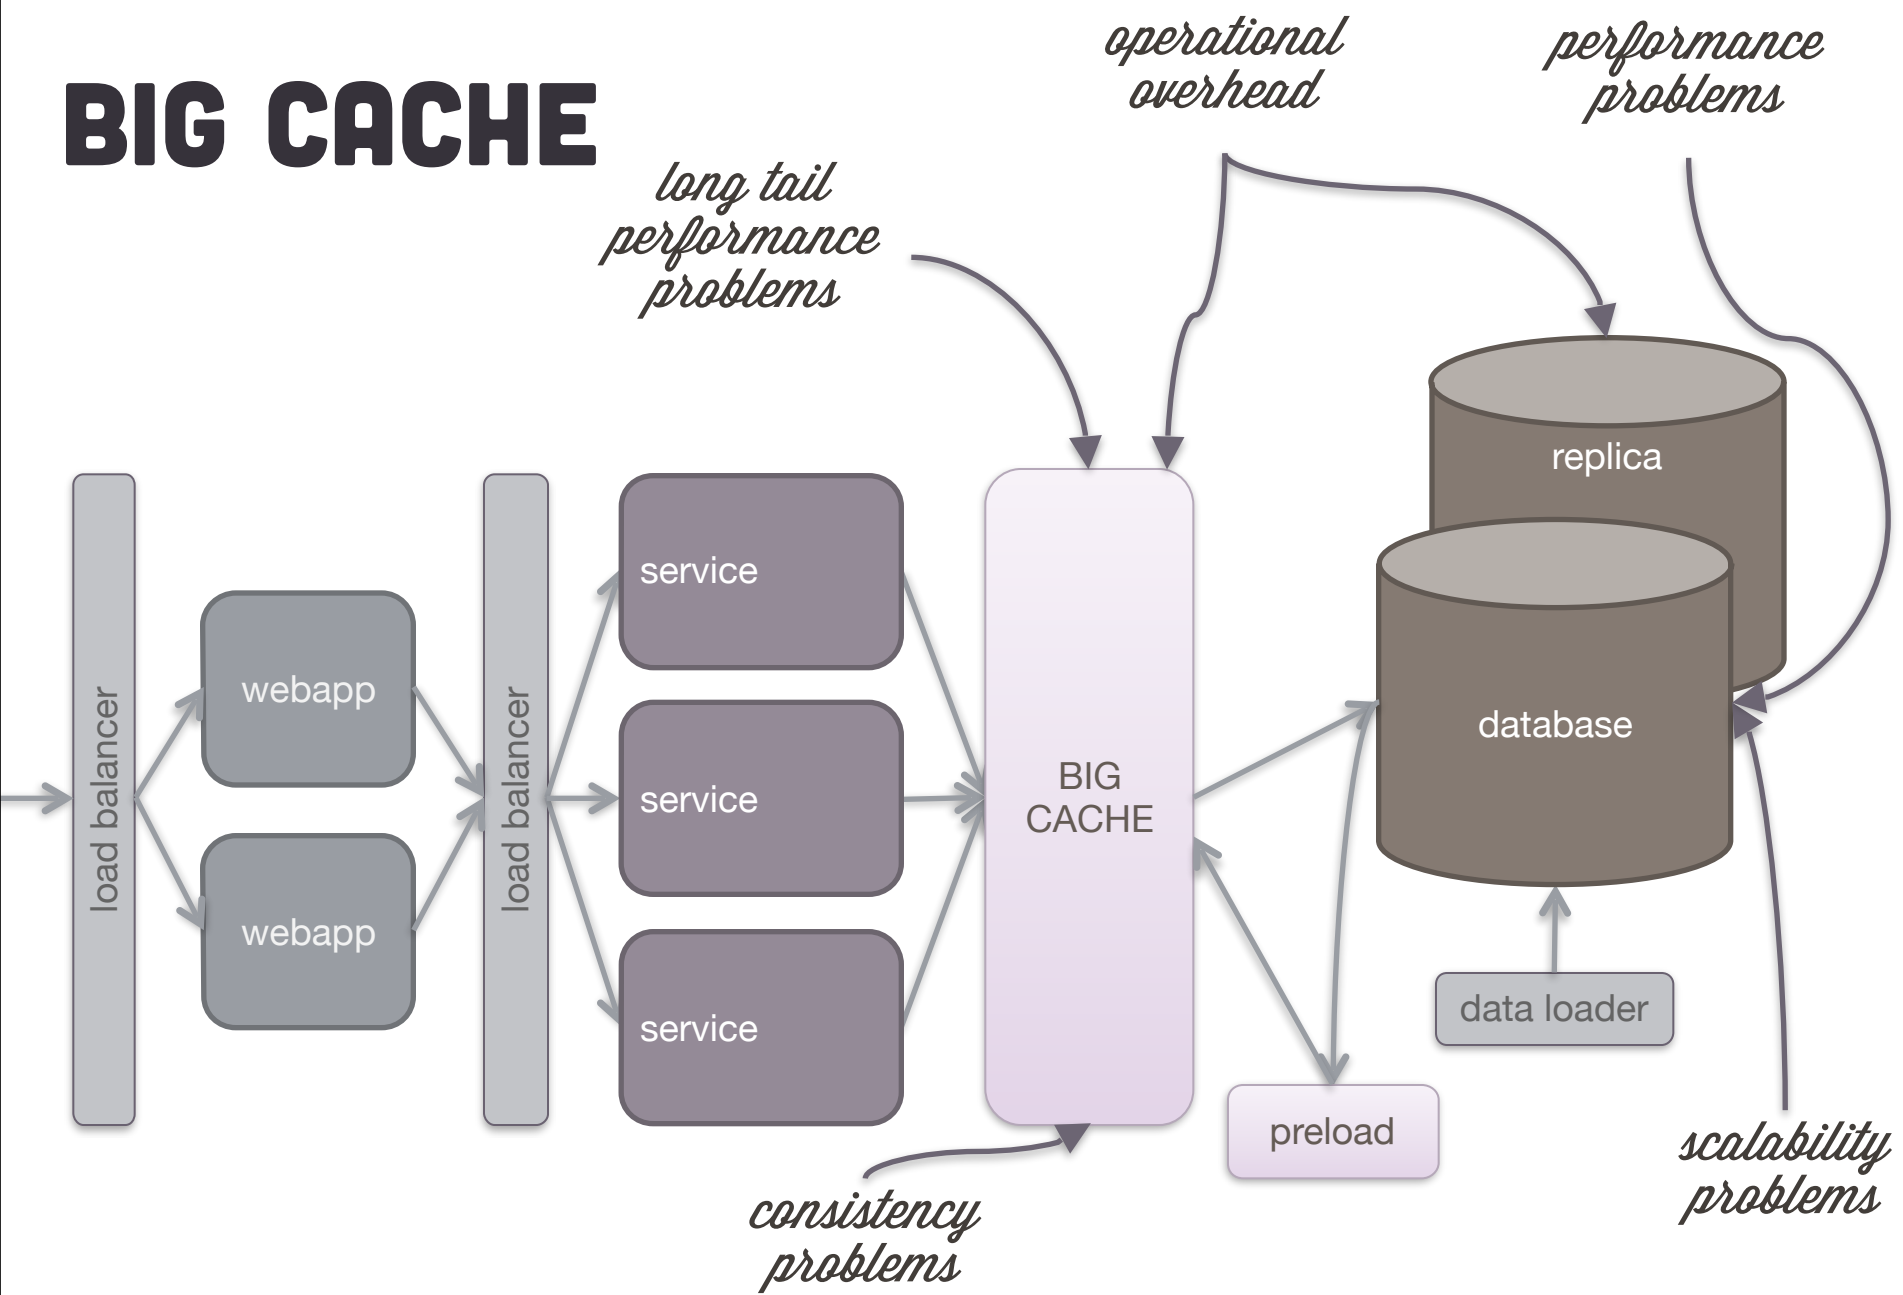
BIG CACHE



BIG CACHE



BIG CACHE



A pug dog is sitting on a green plastic chair that is floating in a body of water. The dog is looking towards the camera. In the background, there is a white boat on the water, a person standing on a dock, and a tree on the right side. The water is calm and reflects the sky and the boat.

BIG CACHE TECHNOLOGIES

BIG CACHE TECHNOLOGIES

memcached(b)




BIG CACHE TECHNOLOGIES

memcached(b)

Do I look like I
need a cache?

ElastiCache (AWS)



BIG CACHE TECHNOLOGIES

memcached(b)

Do I look like I
need a cache?

ElastiCache (AWS)

Oracle Coherence



New


Ikea Kitteh Storage Solutions



Targeted
generic data/
use cases.

New

Ikea Kitteh Storage Solutions




Dynamically
assign keys to
the “nodes”

Targeted
generic data/
use cases.

New

Ikea Kitteh Storage Solutions




Dynamically
assign keys to
the “nodes”

Targeted
generic data/
use cases.

Scales
horizontally

New

Ikea Kitteh Storage Solutions



Dynamically
assign keys to
the “nodes”


Targeted
generic data/
use cases.

Dynamically
rebalances data

Scales
horizontally

New

Ikea Kitteh Storage Solutions



Dynamically
assign keys to
the “nodes”

Targeted
generic data/
use cases.


Dynamically
rebalances data

Scales
horizontally

Poor
performance
on cold starts

New

Ikea Kitten Storage Solutions



Dynamically
assign keys to
the “nodes”

Targeted
generic data/
use cases.

Dynamically
rebalances data

Scales
horizontally

No assumptions
about loading/
updating data

Poor
performance
on cold starts

Ikea Kitteh Storage Solutions

BIG CACHE TECHNOLOGIES

- Additional hardware
 - Additional configuration
 - Additional monitoring
-
- Extra network hop
 - Slow scanning
 - Additional deserialization

BIG CACHE TECHNOLOGIES

- Additional hardware
- Additional configuration
- Additional monitoring



operational overhead

- Extra network hop
- Slow scanning
- Additional deserialization

BIG CACHE TECHNOLOGIES

- Additional hardware
- Additional configuration
- Additional monitoring




operational overhead

- Extra network hop
- Slow scanning
- Additional deserialization

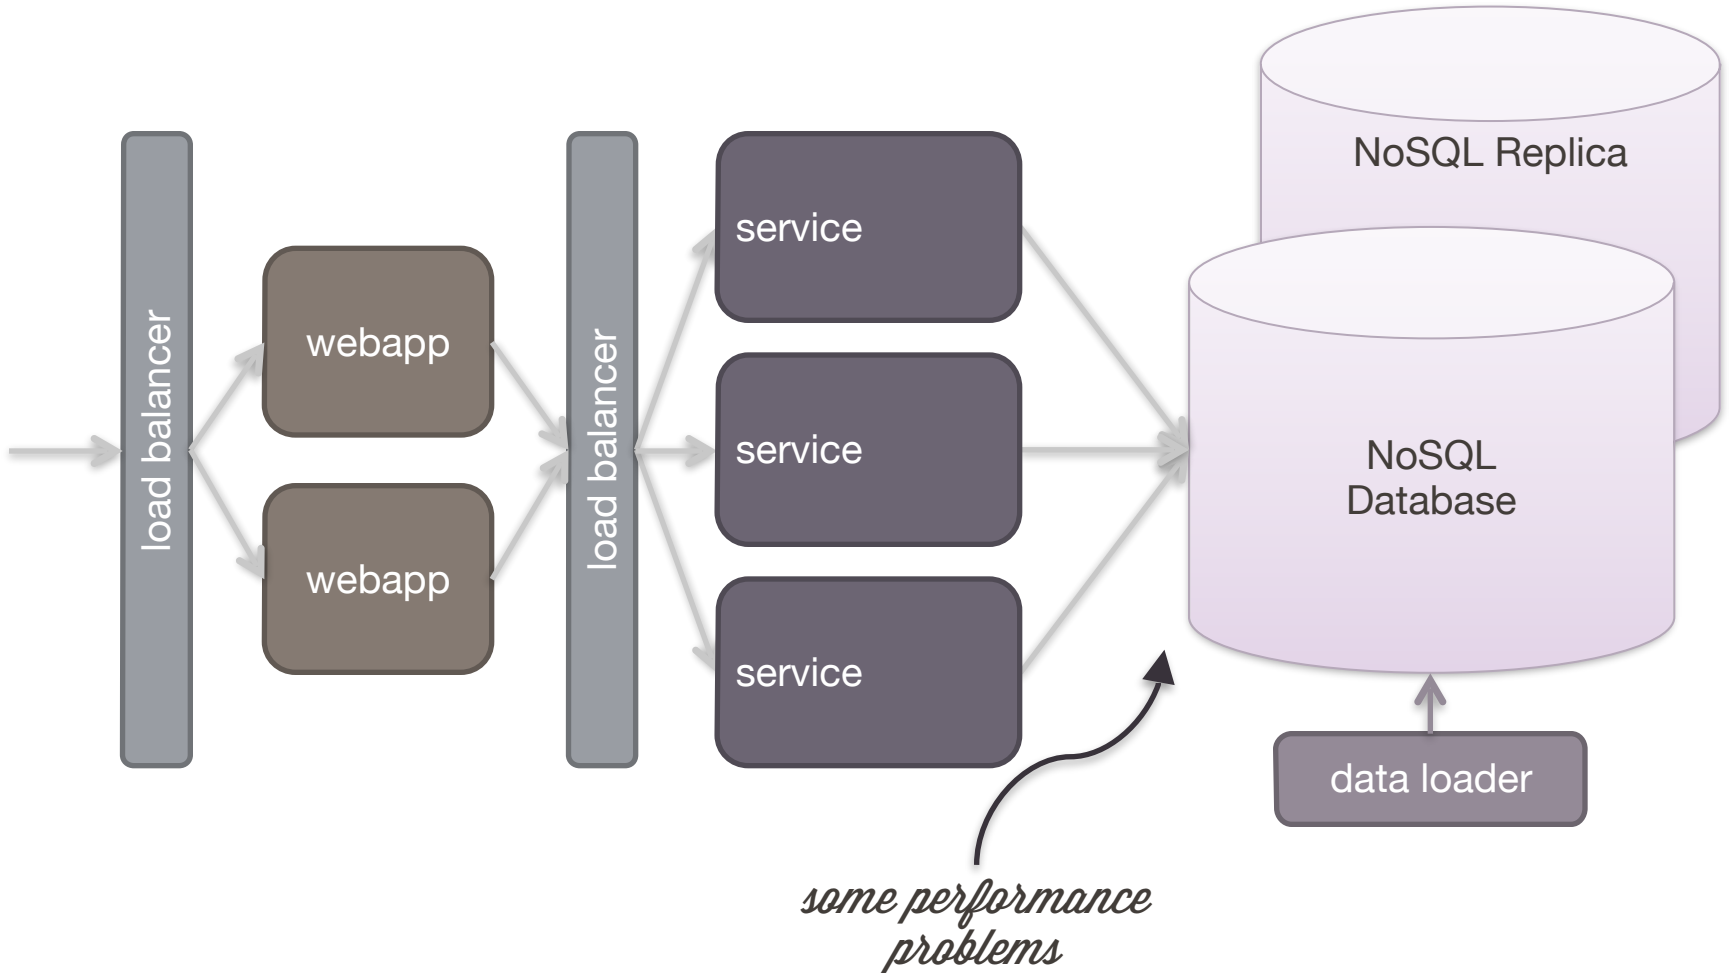


performance

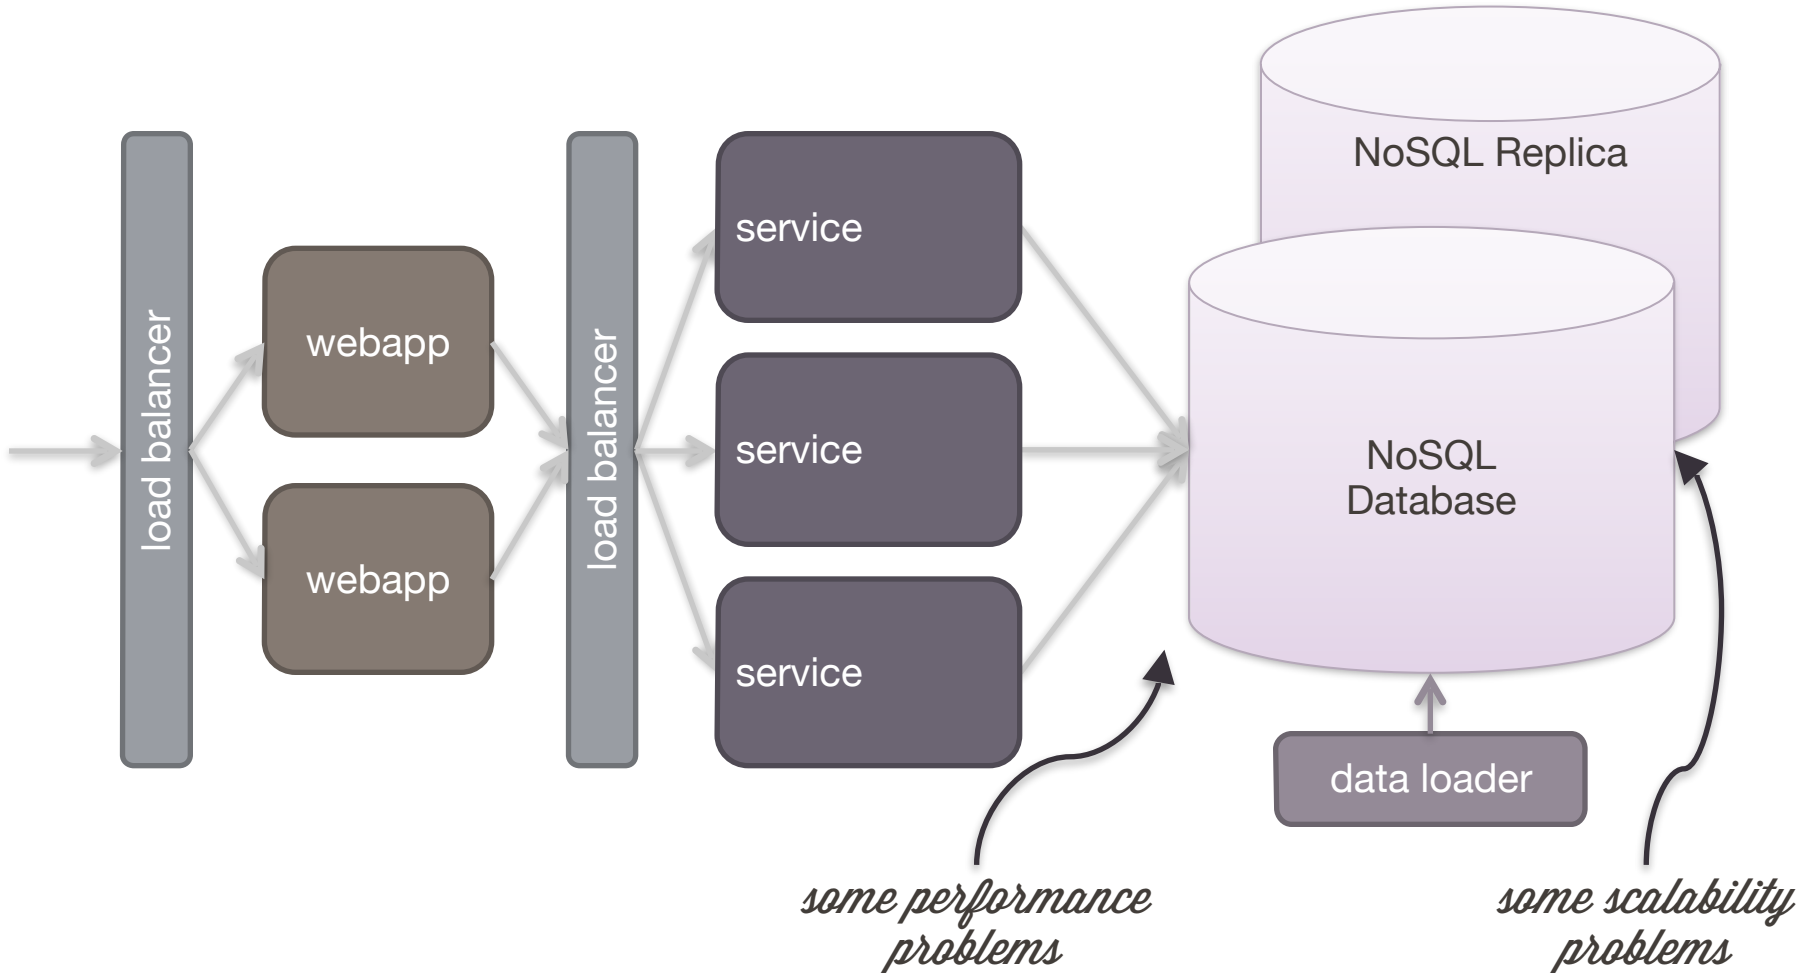
NOSQL TO THE RESCUE?



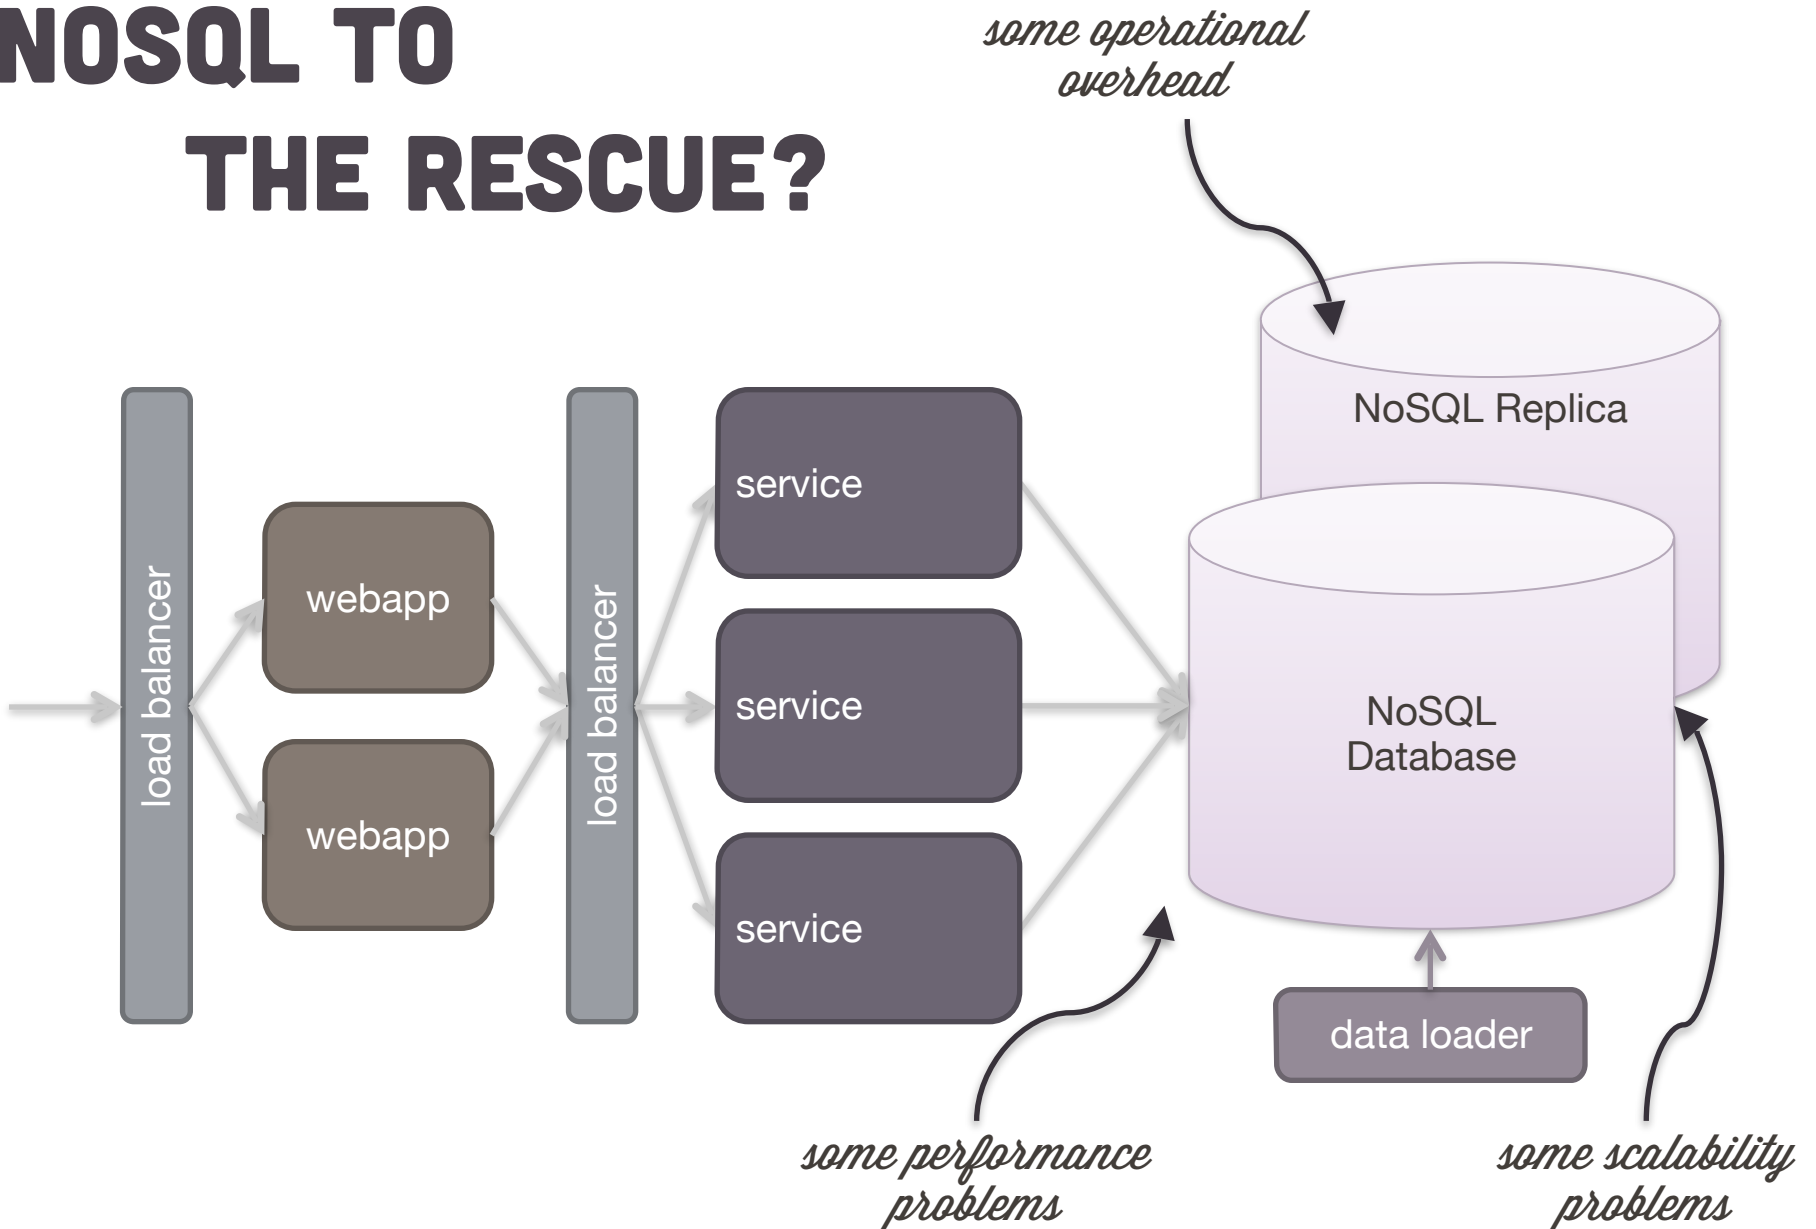
NOSQL TO THE RESCUE?



NOSQL TO THE RESCUE?



NOSQL TO THE RESCUE?




REMOTE STORE RETRIEVAL LATENCY

network

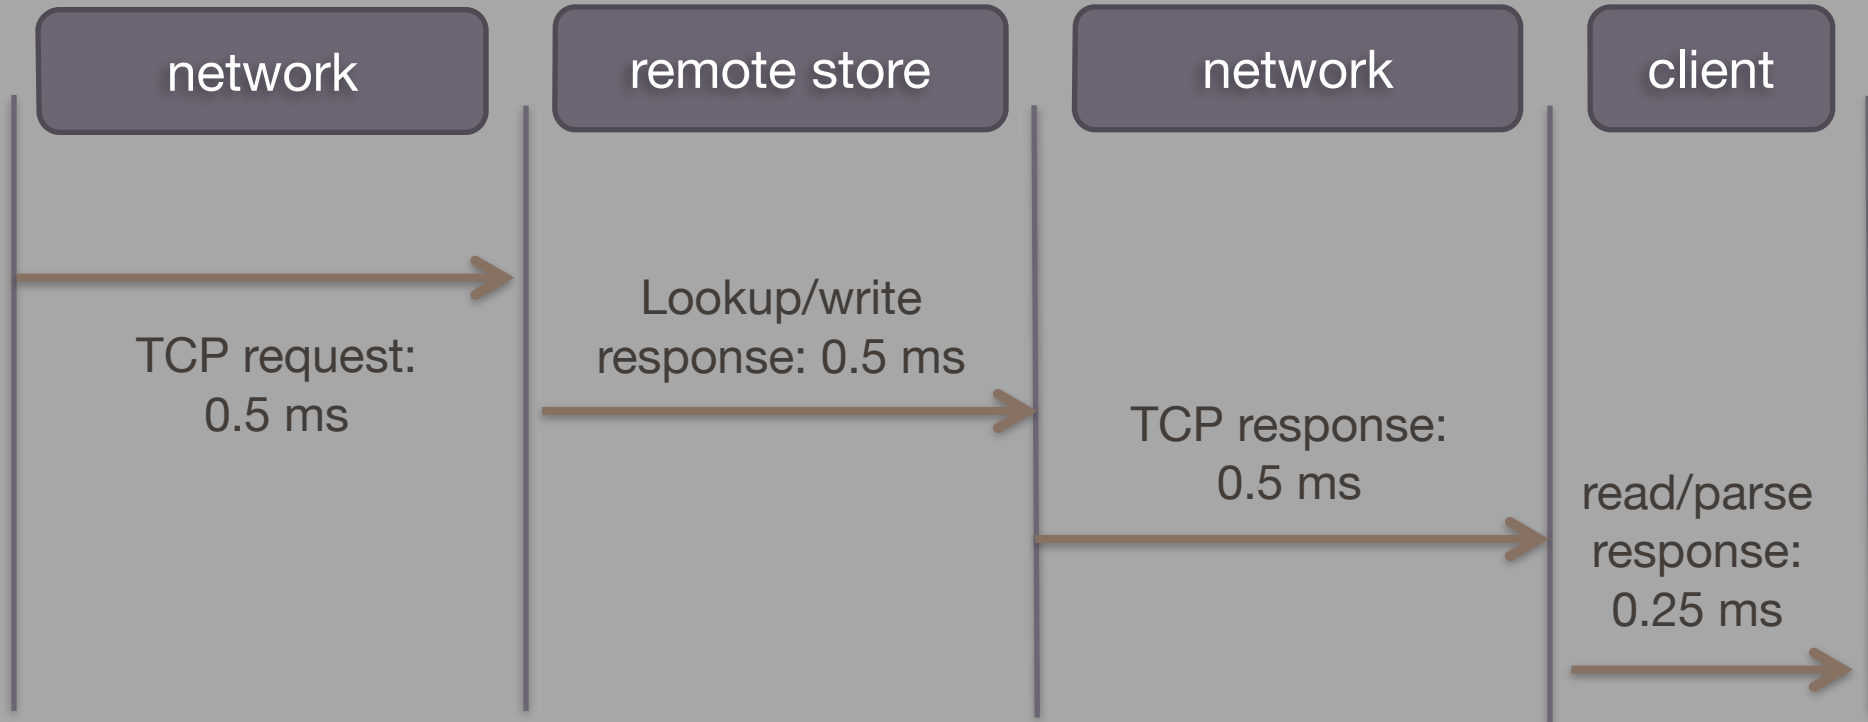
remote store

network

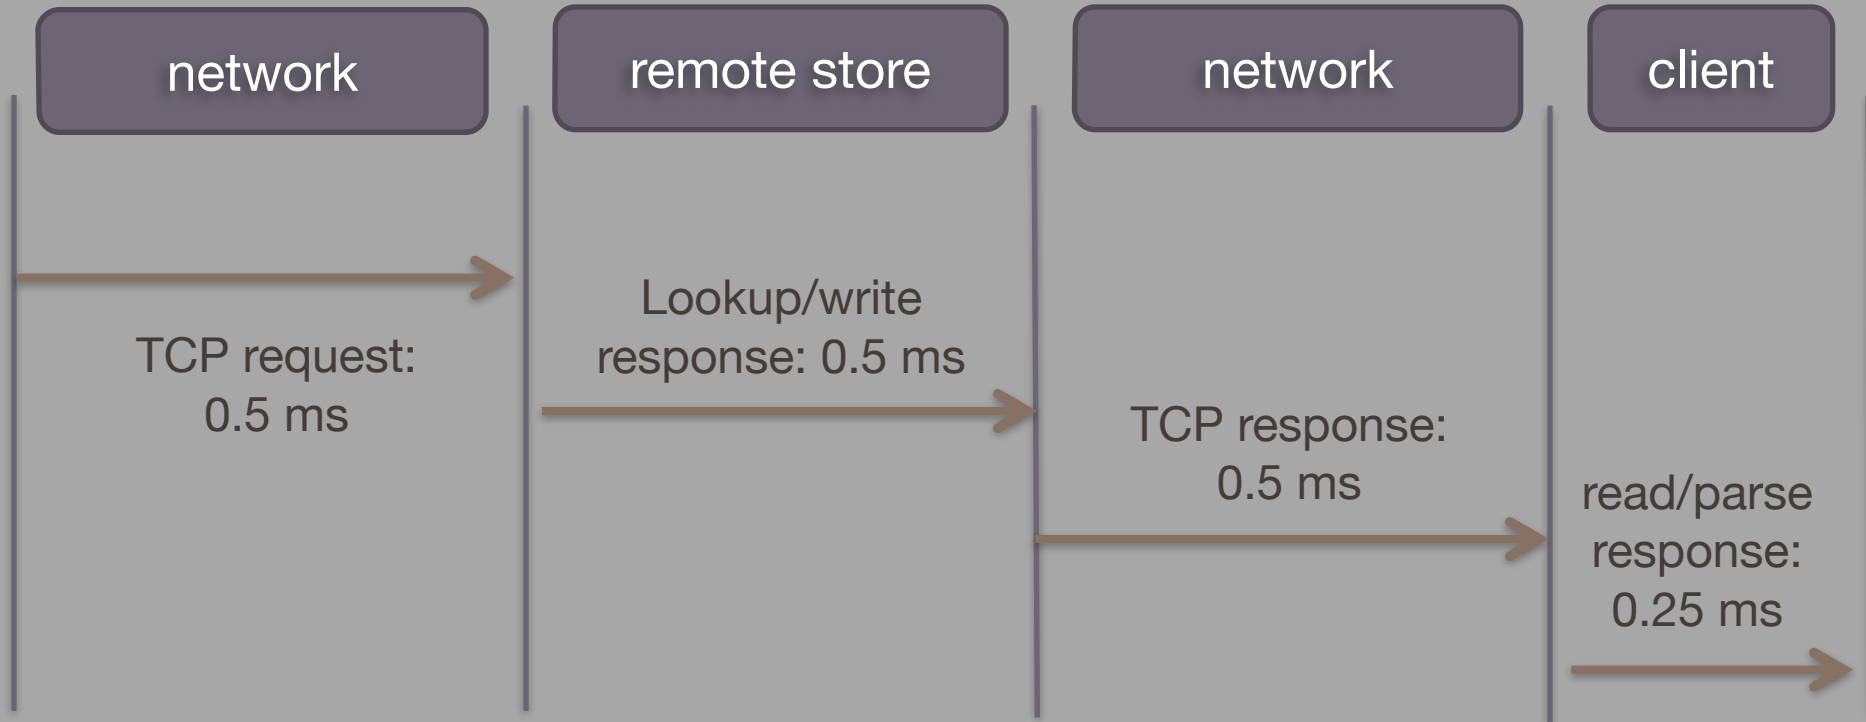
client



REMOTE STORE RETRIEVAL LATENCY



REMOTE STORE RETRIEVAL LATENCY



Total time to retrieve single value:

1.75 MS

TOTAL TIME TO RETRIEVE A SINGLE VALUE

from remote store: 1.75 ms
from memory: 0.001 ms
(10 main memory reads)

TOTAL TIME TO RETRIEVE A SINGLE VALUE

from remote store: 1.75 ms
from memory: 0.001 ms
(10 main memory reads)

SEQUENTIAL ACCESS OF 1 MILLION RANDOM KEYS

from remote store: 30 minutes
from memory: 1 second

THE TRUTH ABOUT DATABASES

“What I'm going to call as the *hot data cliff*: As the size of your hot data set (data frequently read at sustained rates above disk I/O capacity) approaches available memory, write operation bursts that exceeds disk write I/O capacity can create a *trashing death spiral* where hot disk pages that MongoDB desperately needs are evicted from disk cache by the OS as it consumes more buffer space to hold the writes in memory.”

MONGODB

“Redis is an in-memory but persistent on disk database, so it represents a different trade off where very high write and read speed is achieved with the limitation of data sets that can't be larger than memory.”

REDIS

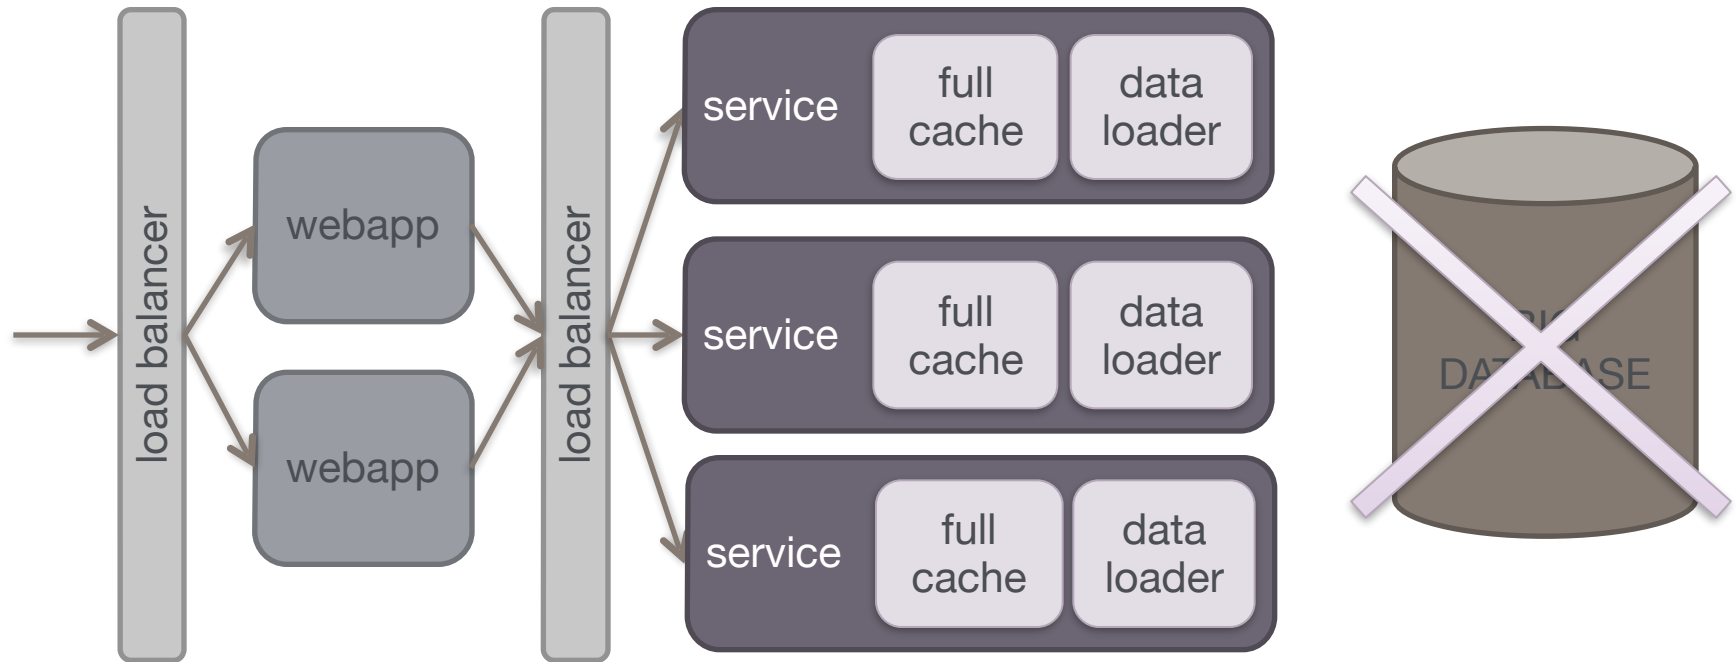
**THEY ARE FAST IF EVERYTHING
FITS INTO MEMORY.**



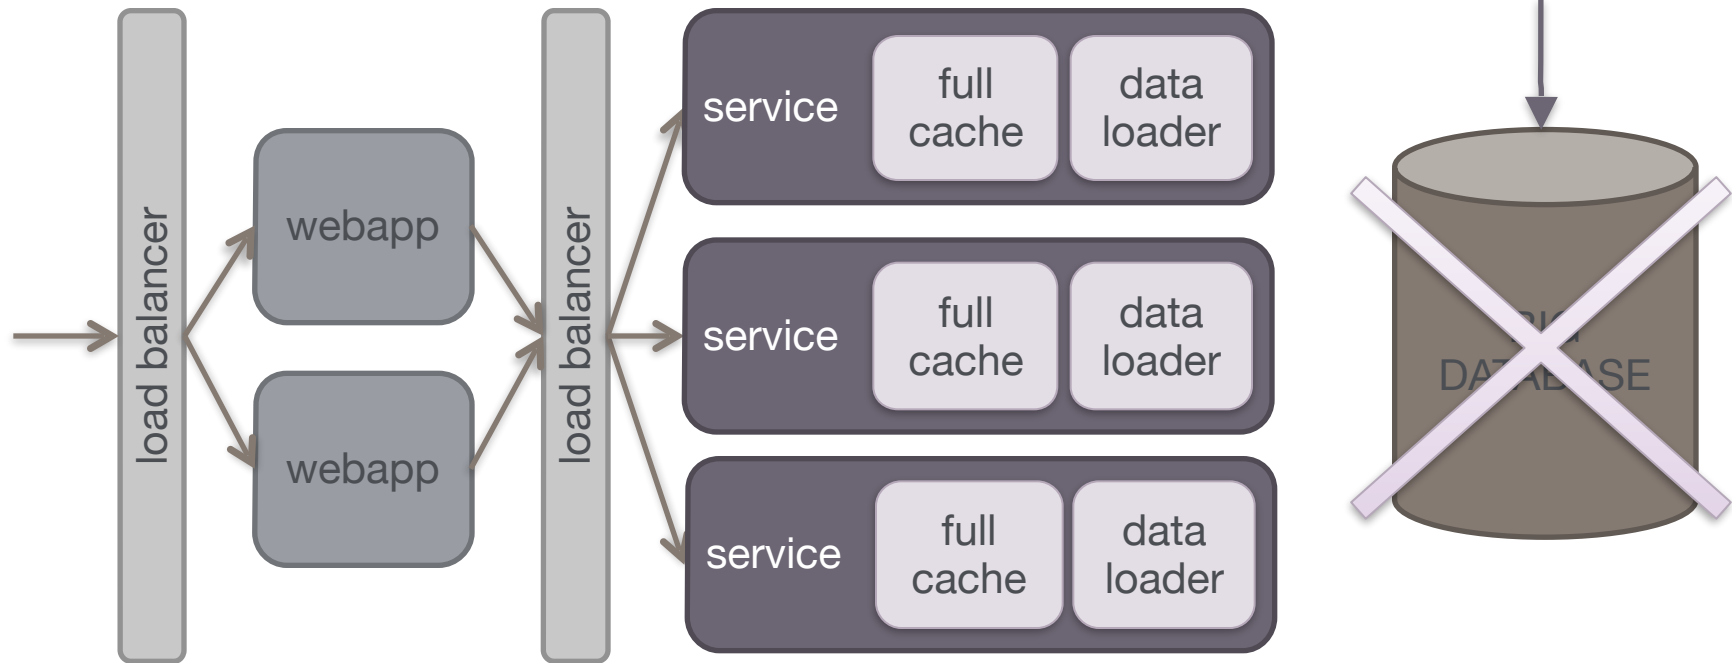
Ur computr needs

More rams


CAN YOU KEEP IT IN MEMORY YOURSELF?



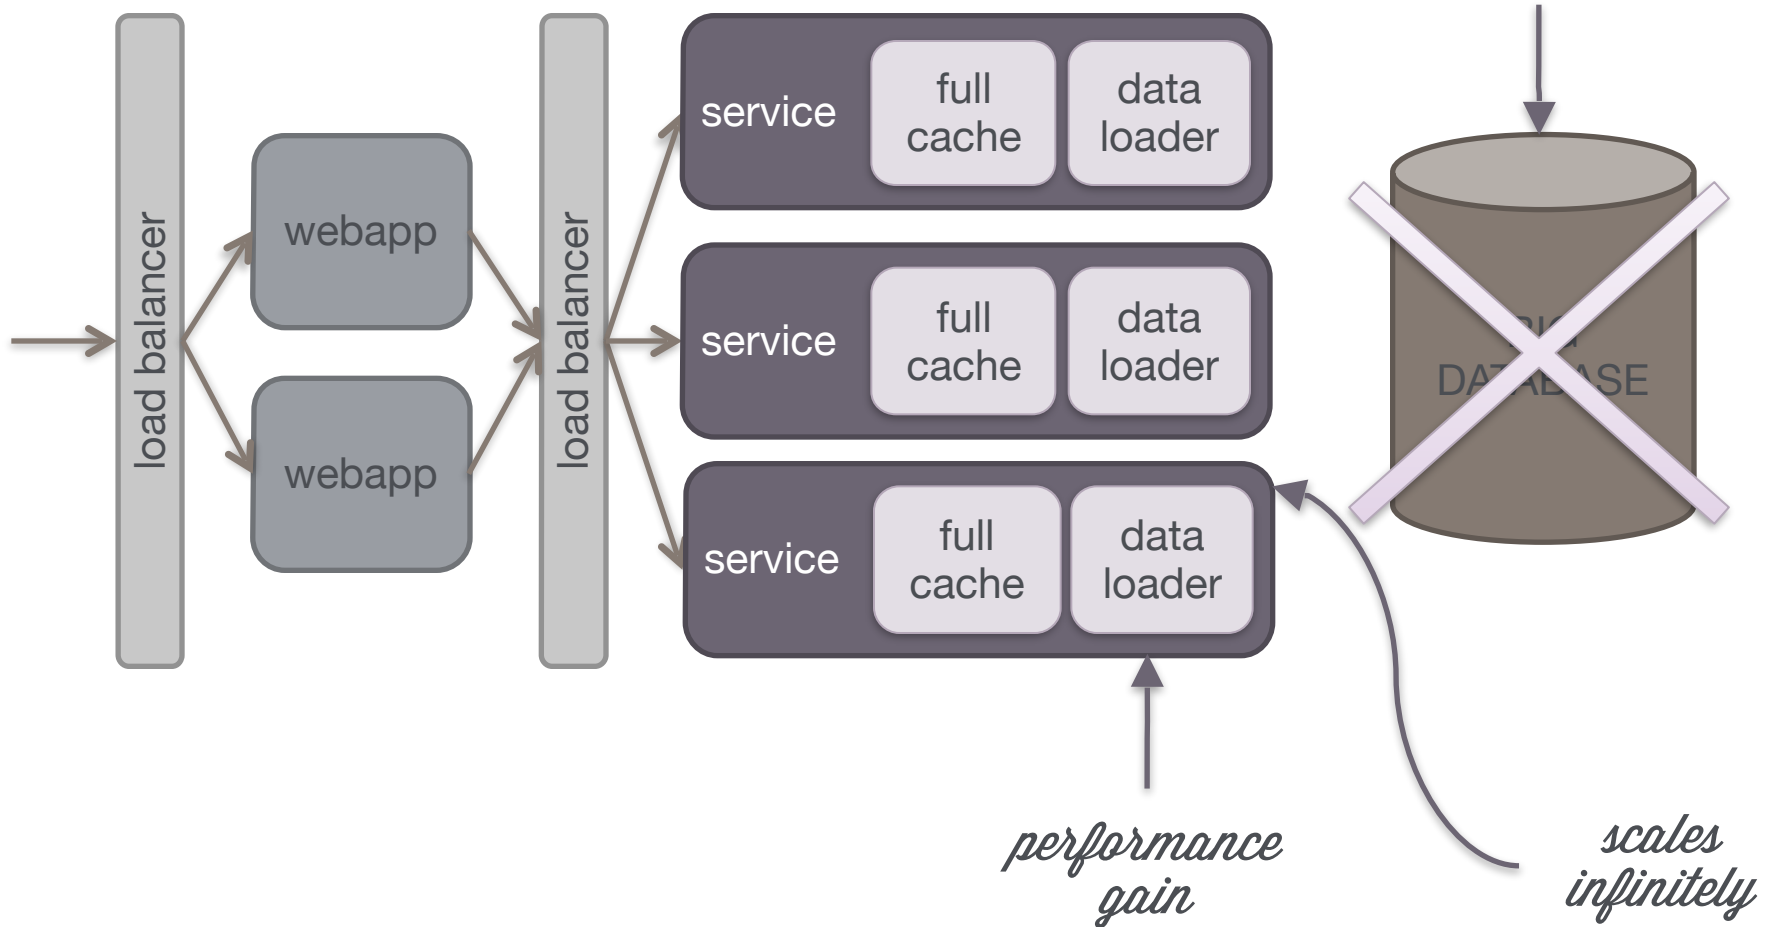
CAN YOU KEEP IT IN MEMORY YOURSELF?



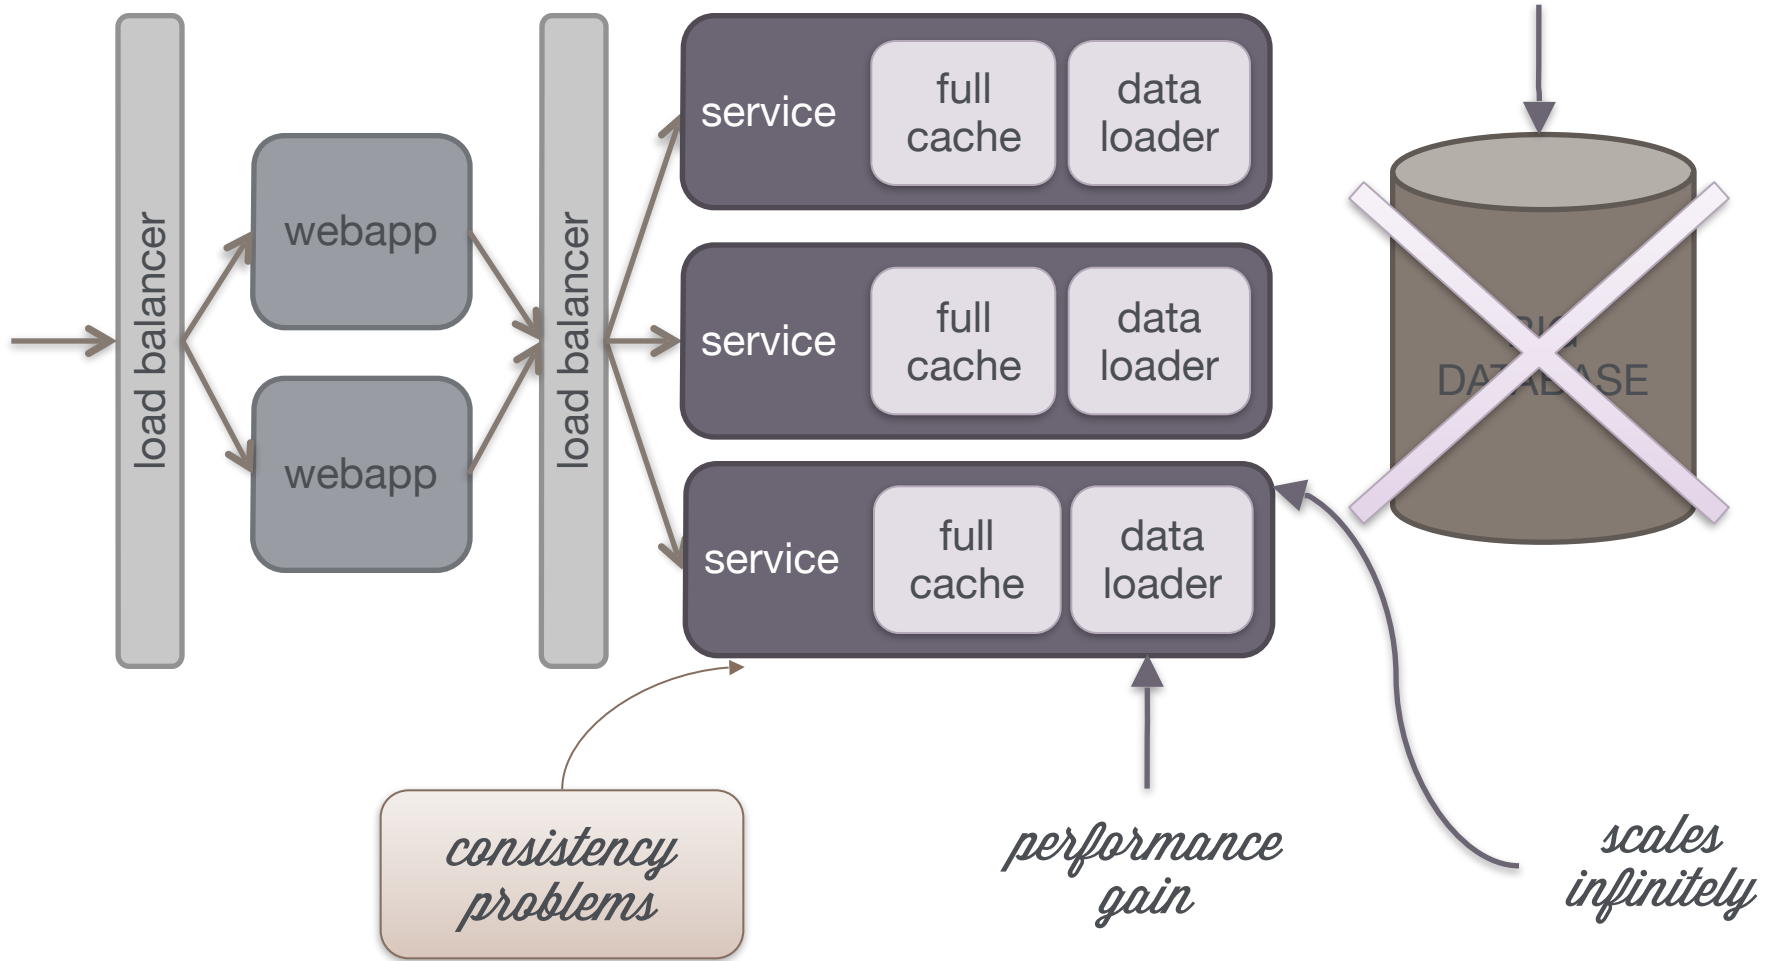
CAN YOU KEEP IT IN MEMORY YOURSELF?



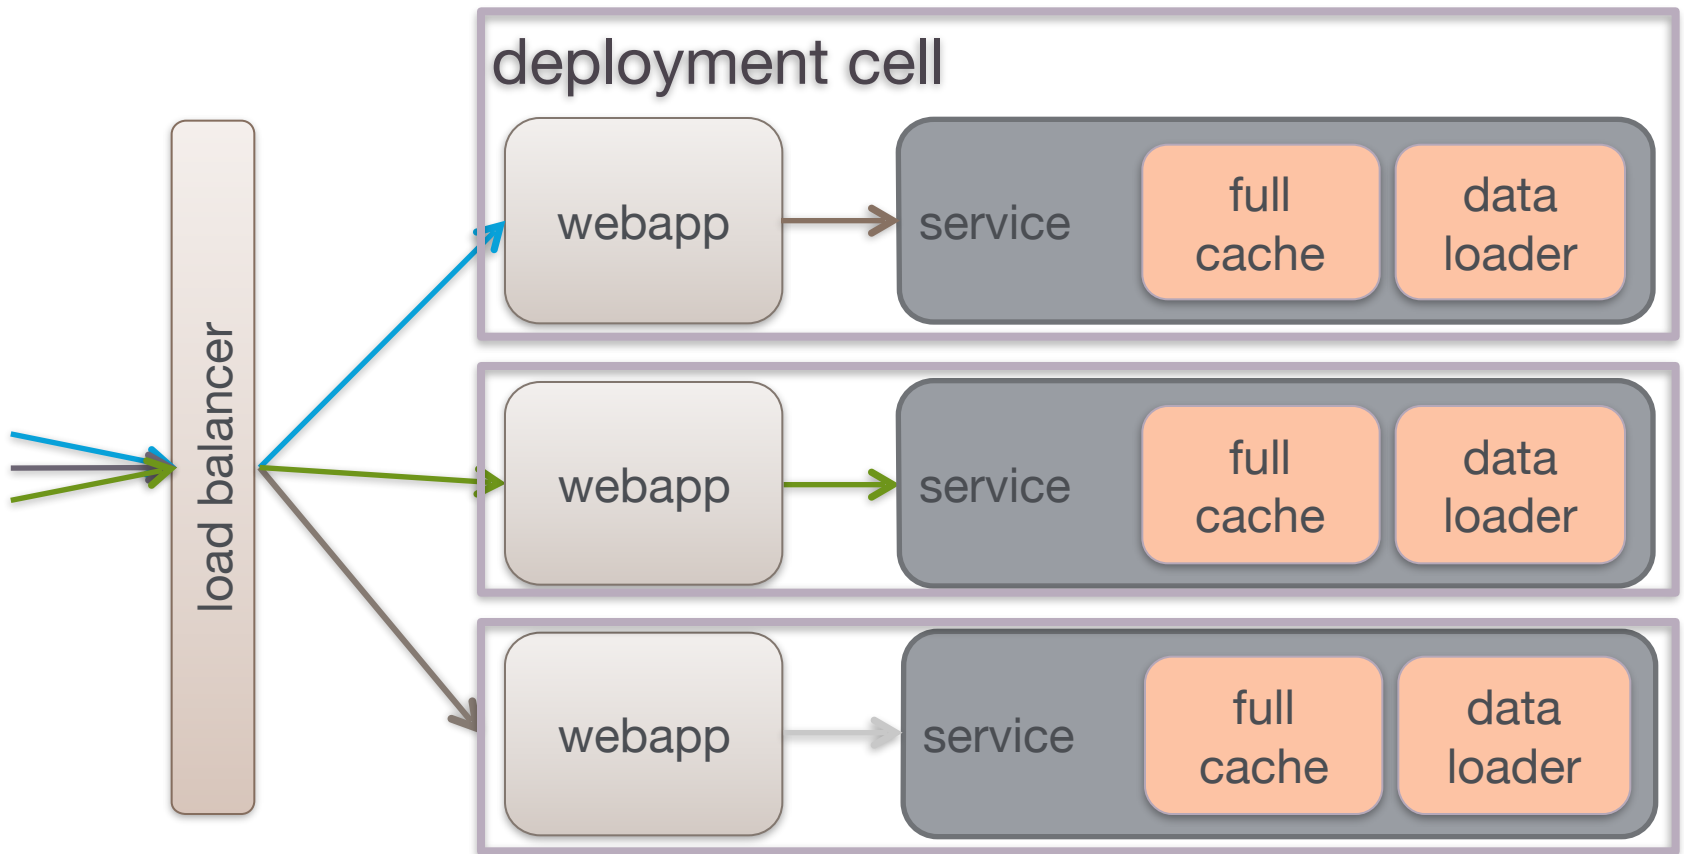
CAN YOU KEEP IT IN MEMORY YOURSELF?



CAN YOU KEEP IT IN MEMORY YOURSELF?

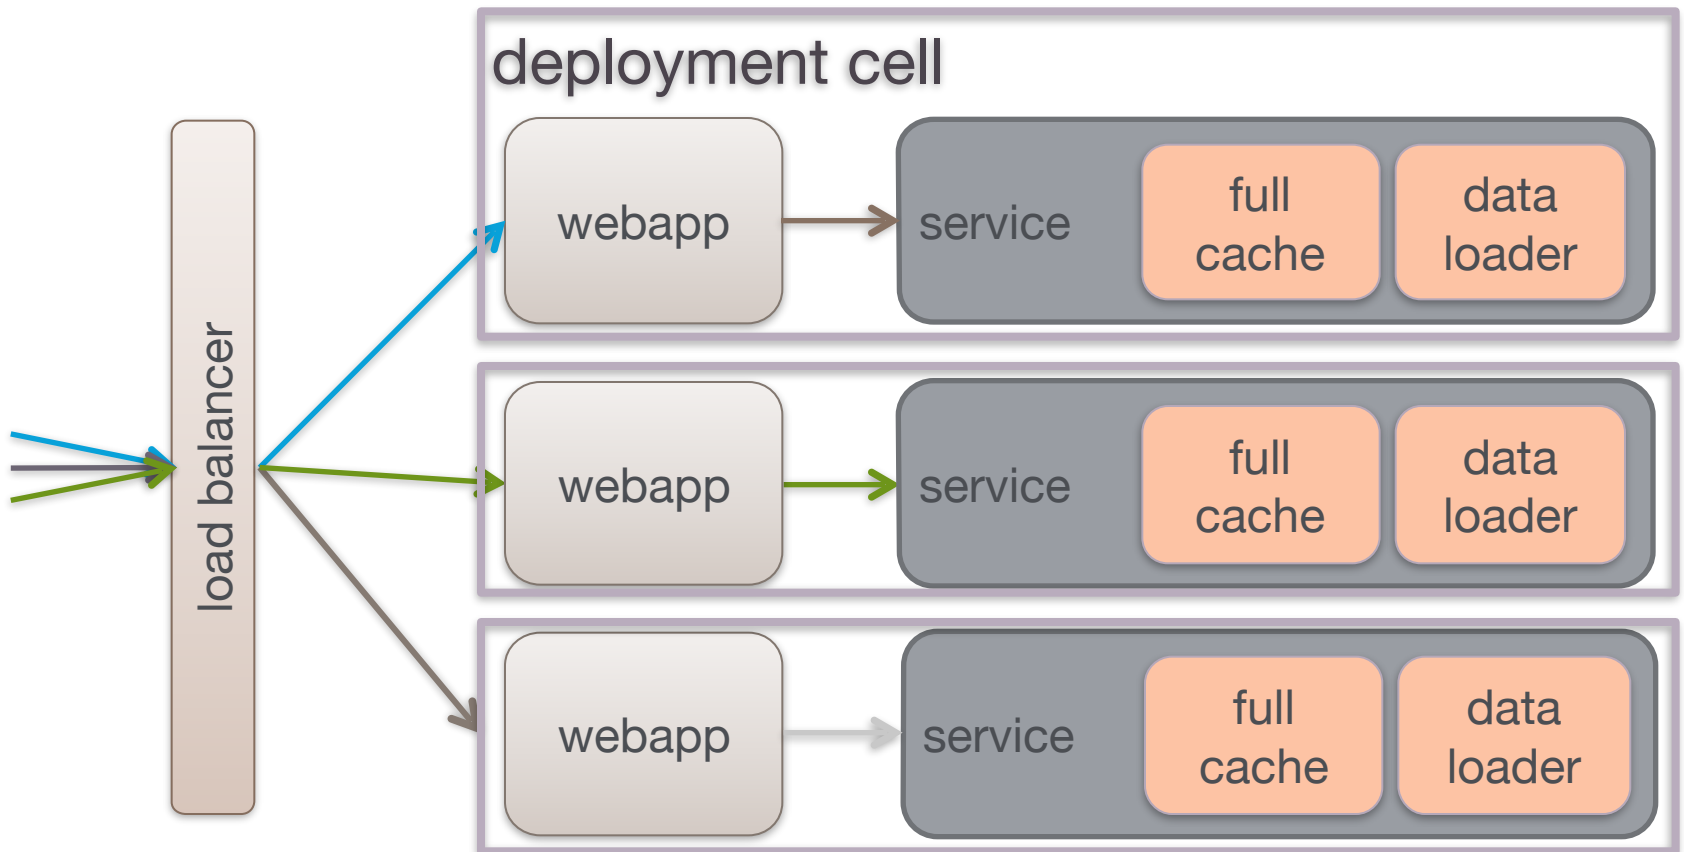


FIXING CONSISTENCY



FIXING CONSISTENCY

1. Deployment "Cells"
2. Sticky user sessions



HOW DO YOU FIT ALL OF THAT DATA INTO MEMORY?



"Programmers waste enormous amounts of time thinking about, or worrying about, the speed of noncritical parts of their programs, and these attempts at efficiency actually have a strong negative impact when debugging and maintenance are considered.

We *should* forget about small efficiencies, say about 97% of the time: **premature optimization is the root of all evil.** *Yet we should not pass up our opportunities in that critical 3%."*

Donald Knuth

**HOW DO YOU FIT ALL
THAT DATA IN MEMORY?**



THE ANSWER




DOMAIN MODEL DESIGN

“Domain Layer (or Model Layer):

Responsible for representing concepts of the business, information about the business situation, and business rules. State that reflects the business situation is controlled and used here, even though the technical details of storing it are delegated to the infrastructure. This layer is the heart of business software.”

Eric Evans, Domain-Driven Design, 2003



DOMAIN MODEL DESIGN GUIDELINES



DOMAIN MODEL DESIGN GUIDELINES

#1 Keep it immutable



DOMAIN MODEL DESIGN GUIDELINES

#1 Keep it immutable



#2 Use independent hierarchies

DOMAIN MODEL DESIGN GUIDELINES

#1 Keep it immutable

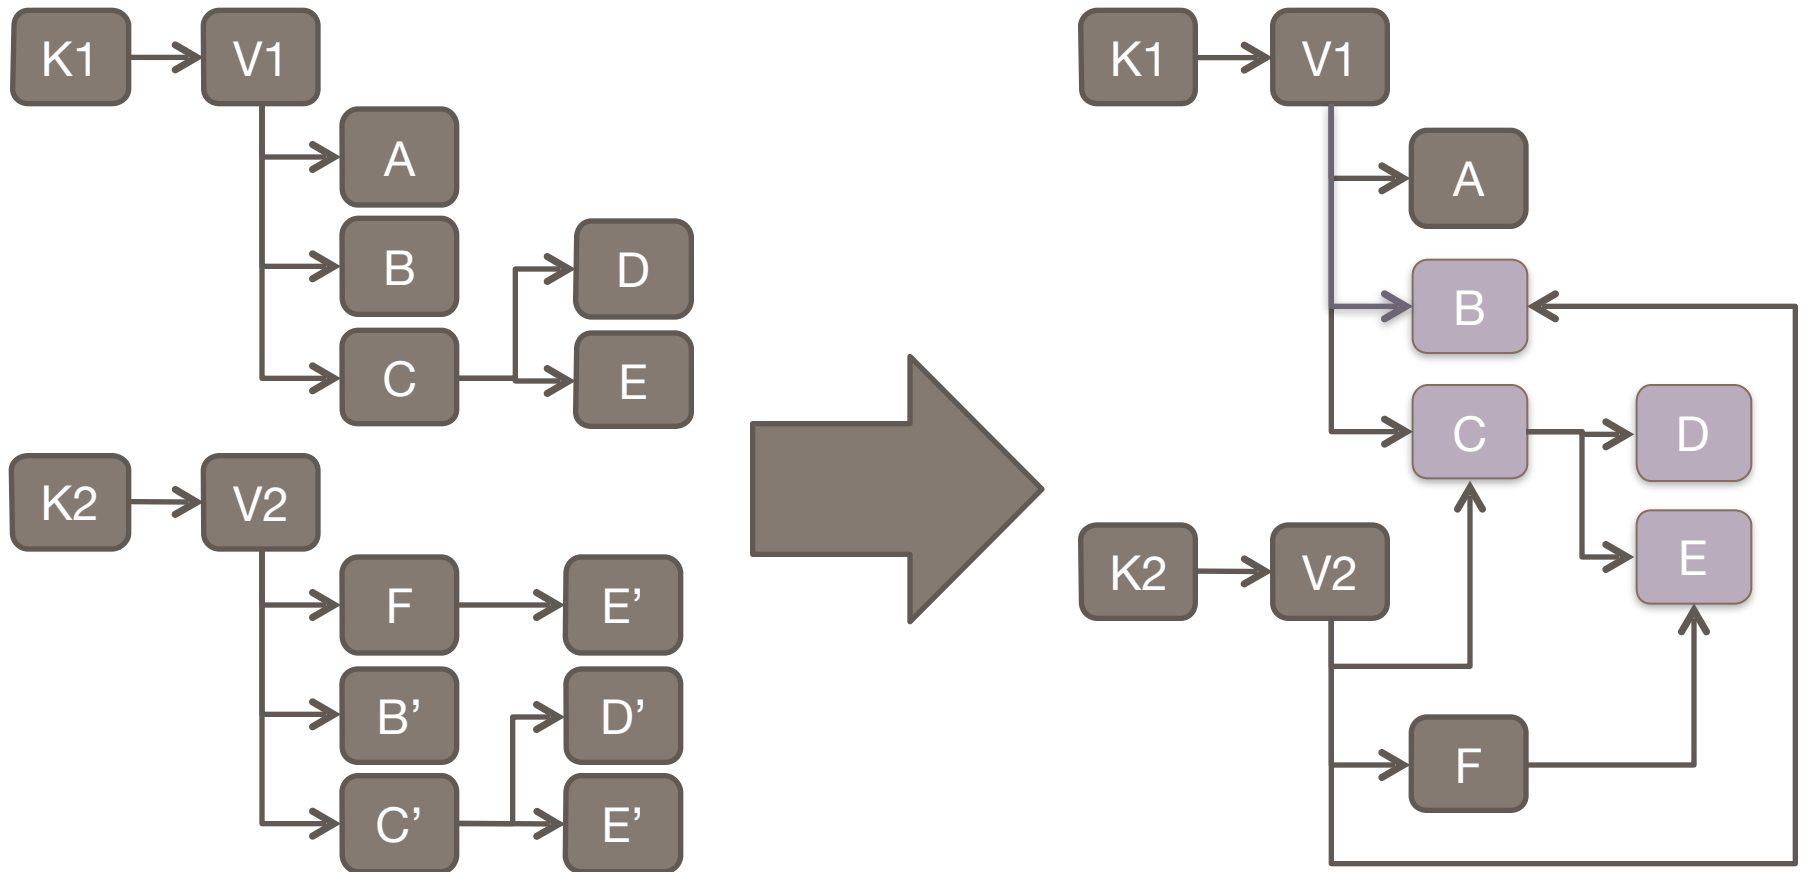
#3 Optimize Data



#2 Use independent hierarchies

Help! I am in the trunk!

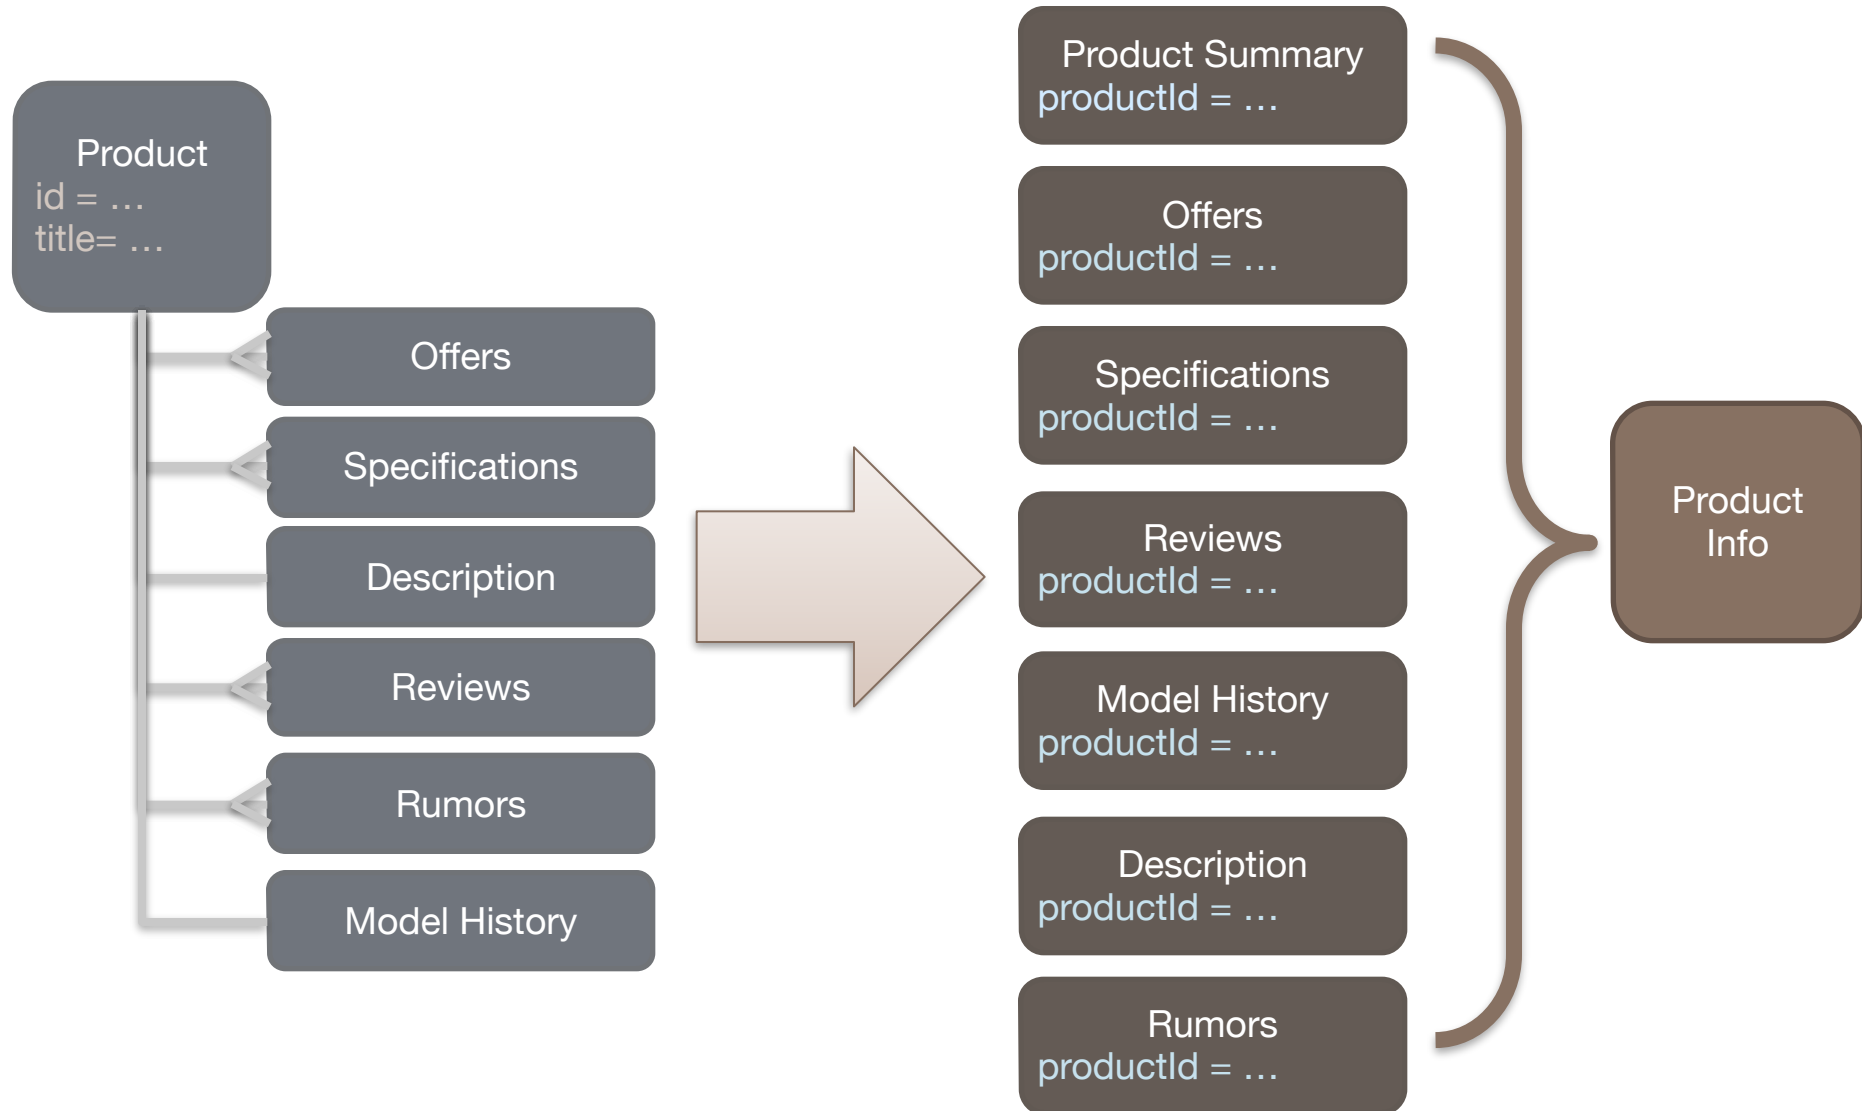
intern() YOUR IMMUTABLES




```
private final Map<Class<?>, Map<Object, WeakReference<Object>>> cache =  
    new ConcurrentHashMap<Class<?>, Map<Object,  
WeakReference<Object>>>();
```

```
public <T> T intern(T o) {  
    if (o == null)  
        return null;  
    Class<?> c = o.getClass();  
    Map<Object, WeakReference<Object>> m = cache.get(c);  
    if (m == null)  
        cache.put(c, m = synchronizedMap(new WeakHashMap<Object,  
WeakReference<Object>>()));  
    WeakReference<Object> r = m.get(o);  
    @SuppressWarnings("unchecked")  
    T v = (r == null) ? null : (T) r.get();  
    if (v == null) {  
        v = o;  
        m.put(v, new WeakReference<Object>(v));  
    }  
    return v;  
}
```

USE INDEPENDENT HIERARCHIES



COLLECTION

OPTIMIZATION



Leverage PRIMITIVE KEYS/VALUES

collection with 10,000 elements [0 .. 9,999]	size in memory
<i>java.util.ArrayList<Integer></i>	200K
<i>java.util.HashSet<Integer></i>	546K
<i>gnu.trove.list.array.TIntArrayList</i>	40K
<i>gnu.trove.set.hash.TIntHashSet</i>	102K

[Trove](#) (“High Performance Collections for Java”)

Optimize

SMALL IMMUTABLE COLLECTIONS

Collections with small number of entries (up to ~20):

```
class ImmutableMap<K, V> implements Map<K,V>, Serializable {  
    ... }  
}
```

```
class MapN<K, V> extends ImmutableMap<K, V> {  
    final K k1, k2, ..., kN;  
    final V v1, v2, ..., vN;  
    @Override public boolean containsKey(Object key) {  
        if (eq(key, k1)) return true;  
        if (eq(key, k2)) return true;  
        ...  
        return false;  
    }  
    ...  
}
```

SPACE SAVINGS



`java.util.HashMap:`

128 bytes + 32 bytes per entry

compact immutable map:

24 bytes + 8 bytes per entry

NUMERIC DATA OPTIMIZATION



EXAMPLE: PRICE HISTORY

Problem:

- Store daily prices for 1M products, 2 offers per product
- Average price history length per product ~2 years

Total price points:

$$(1M + 2M) * 730 = \sim 2 \text{ billion}$$

PRICE
HISTORY
FIRST
ATTEMPT

TreeMap<Date, Double>

88 bytes per entry * 2 billion =
~180 GB

TYPICAL SHOPPING PRICE HISTORY

2 Year Silver Price in USD/oz

Last Close: 26.89

High: 48.58 Low: 17.60 ▲9.06 50.78%



0

20

60

70

90

100

120 121

days

TYPICAL SHOPPING PRICE HISTORY

2 Year Silver Price in USD/oz

Last Close: 26.89

High: 48.58 Low: 17.60 ▲9.06 50.78%



0

20

60

70


90

100



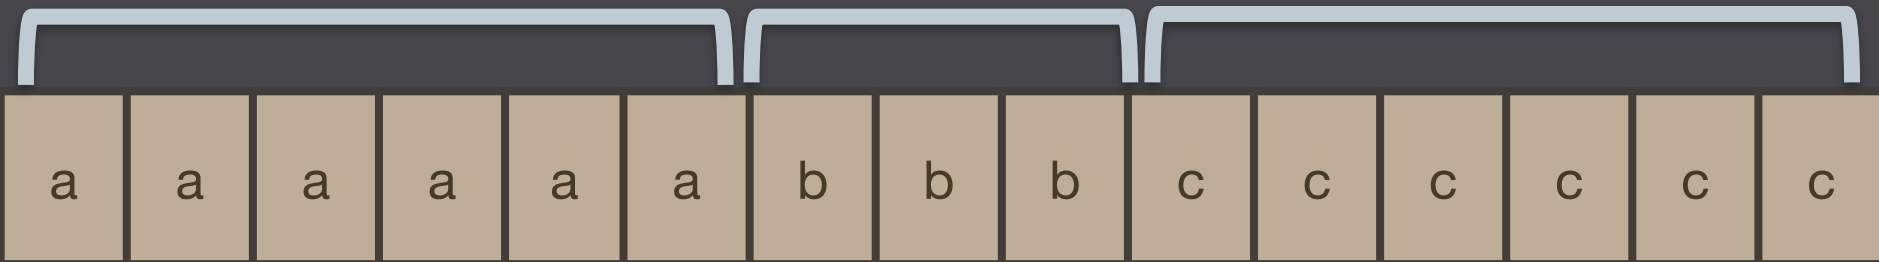
120 121

days

TYPICAL SHOPPING PRICE HISTORY



RUN LENGTH ENCODING



PRICE HISTORY OPTIMIZATION

- **positive:** price (adjusted to scale)
- **negative:** run length (precedes price)
- **zero:** unavailable

-20	100	-40	150	-10	140	-20	100	-10	0	-20	100	90	-9	80
-----	-----	-----	-----	-----	-----	-----	-----	-----	---	-----	-----	----	----	----

- Drop pennies
- Store prices in primitive short (use scale factor to represent prices greater than Short.MAX_VALUE)

Memory:

$15 * 2 + 16$ (array) + 24 (start date) + 4 (scale factor) = **74 bytes**

SPACE SAVINGS

Reduction compared to
TreeMap<Date, Double>:

155 times

Estimated memory for 2 billion price
points:

1.2 GB

SPACE SAVINGS

Reduction compared to
TreeMap<Date, Double>:

155 times

Estimated memory for 2 billion price
points:

1.2 GB << 180 GB

PRICE HISTORY MODEL

```
public class PriceHistory {  
  
    private final Date startDate; // or use org.joda.time.LocalDate  
    private final short[] encoded;  
    private final int scaleFactor;  
  
    public PriceHistory(SortedMap<Date, Double> prices) { ... } // encode  
    public SortedMap<Date, Double> getPricesByDate() { ... } // decode  
    public Date getStartDate() { return startDate; }  
  
    // Below computations implemented directly against encoded data  
    public Date getEndDate() { ... }  
    public Double getMinPrice() { ... }  
    public int getNumChanges(double minChangeAmt, double minChangePct,  
boolean abs) { ... }  
    public PriceHistory trim(Date startDate, Date endDate) { ... }  
    public PriceHistory interpolate() { ... }  
}
```


**KNOW YOUR
DATA**

COMPRESS TEXT



STRING COMPRESSION:

BYTE ARRAYS

- Use the minimum character set encoding

```
static Charset UTF8 = Charset.forName("UTF-8");
```

```
String s = "The quick brown fox jumps over the lazy dog"; // 42 chars, 136  
bytes
```

```
byte[] b a= "The quick brown fox jumps over the lazy dog".getBytes(UTF8); //  
64 bytes
```

```
String s1 = "Hello"; // 5 chars, 64 bytes
```

```
byte[] b1 = "Hello".getBytes(UTF8); // 24 bytes
```

```
byte[] toBytes(String s) { return s == null ? null : s.getBytes(UTF8); }
```

```
String toString(byte[] b) { return b == null ? null : new String(b, UTF8); }
```

STRING COMPRESSION: SHARED PREFIX

- Great for URLs

```
public class PrefixedString {  
    private PrefixedString prefix;  
    private byte[] suffix;  
  
    . . .  
  
    @Override public int hashCode() { ... }  
    @Override public boolean equals(Object o) { ... }  
}
```

STRING COMPRESSION: SHORT ALPHANUMERIC CASE-INSENSITIVE STRINGS

```
public abstract class AlphaNumericString {
    public static AlphaNumericString make(String s) {
        try { return new Numeric(Long.parseLong(s, Character.MAX_RADIX));
        } catch (NumberFormatException e) { return new Alpha(s.getBytes(UTF8)); }
    }
    protected abstract String value();
    @Override public String toString() {return value(); }
    private static class Numeric extends AlphaNumericString {
        long value;
        Numeric(long value) { this.value = value; }
        @Override protected String value() { return Long.toString(value, Character.MAX_RADIX); }
        @Override public int hashCode() { ... }
        @Override public boolean equals(Object o) { ... }
    }
    private static class Alpha extends AlphaNumericString {
        byte[] value;
        Alpha(byte[] value) {this.value = value; }
        @Override protected String value() { return new String(value, UTF8); }
        @Override public int hashCode() { ... }
        @Override public boolean equals(Object o) { ... }
    }
}
```



STRING COMPRESSION: LARGE STRINGS



Gzip

Become the master of your strings!

STRING COMPRESSION: LARGE STRINGS



bzip2

Gzip

Become the
master of your
strings!

STRING COMPRESSION: LARGE STRINGS



bzip2

Just convert to
byte[] first,
then compress

Gzip

Become the
master of your
strings!

STRING COMPRESSION: LARGE STRINGS

JVM TUNING

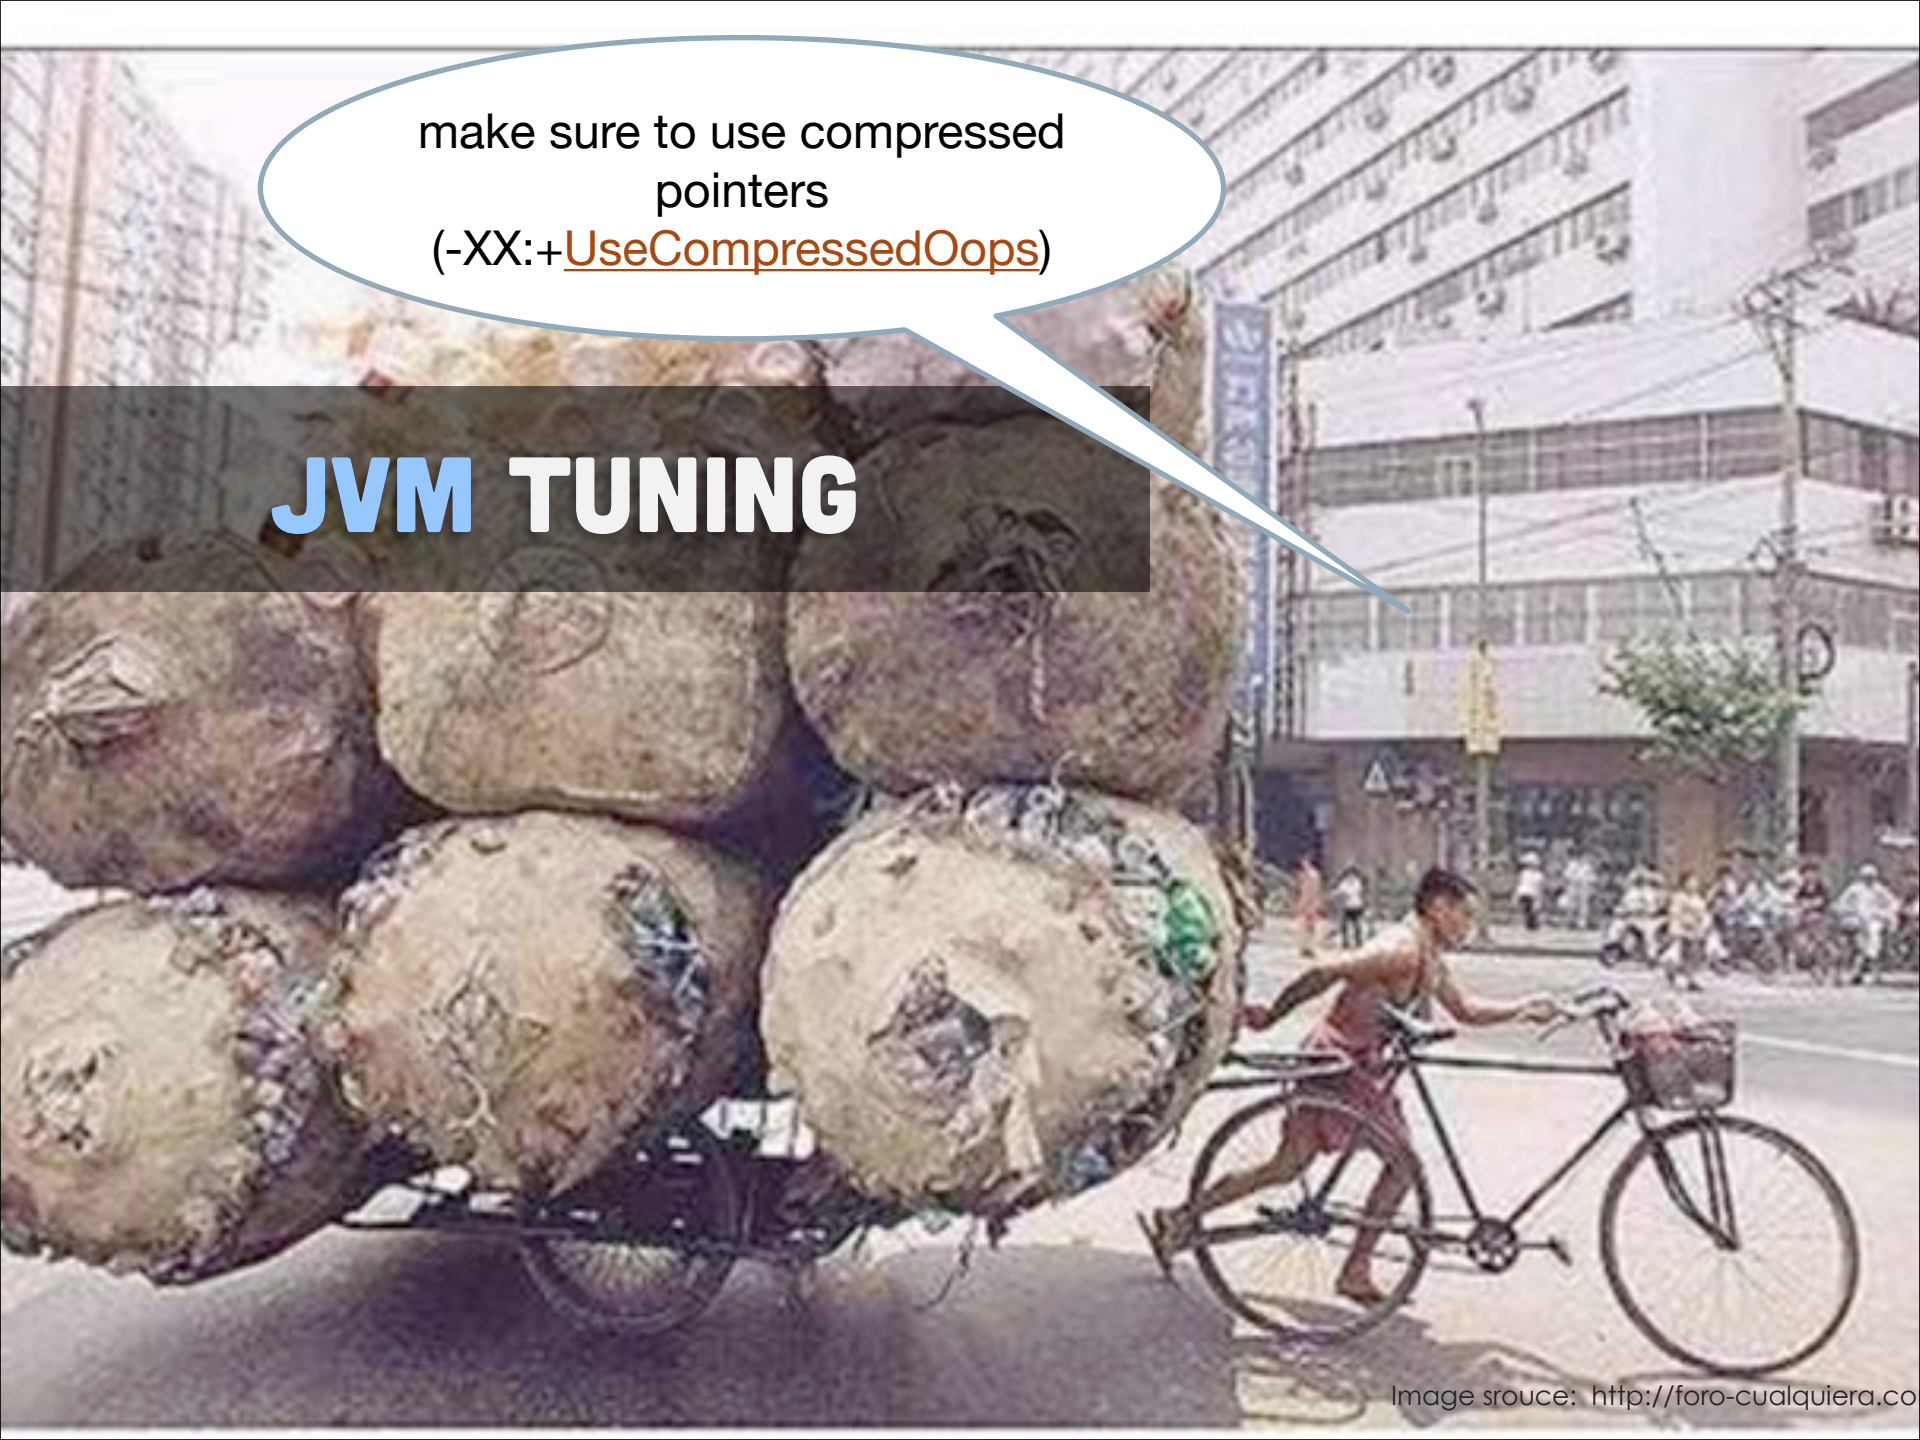


JVM TUNING



make sure to use compressed
pointers
(-XX:+UseCompressedOops)

JVM TUNING

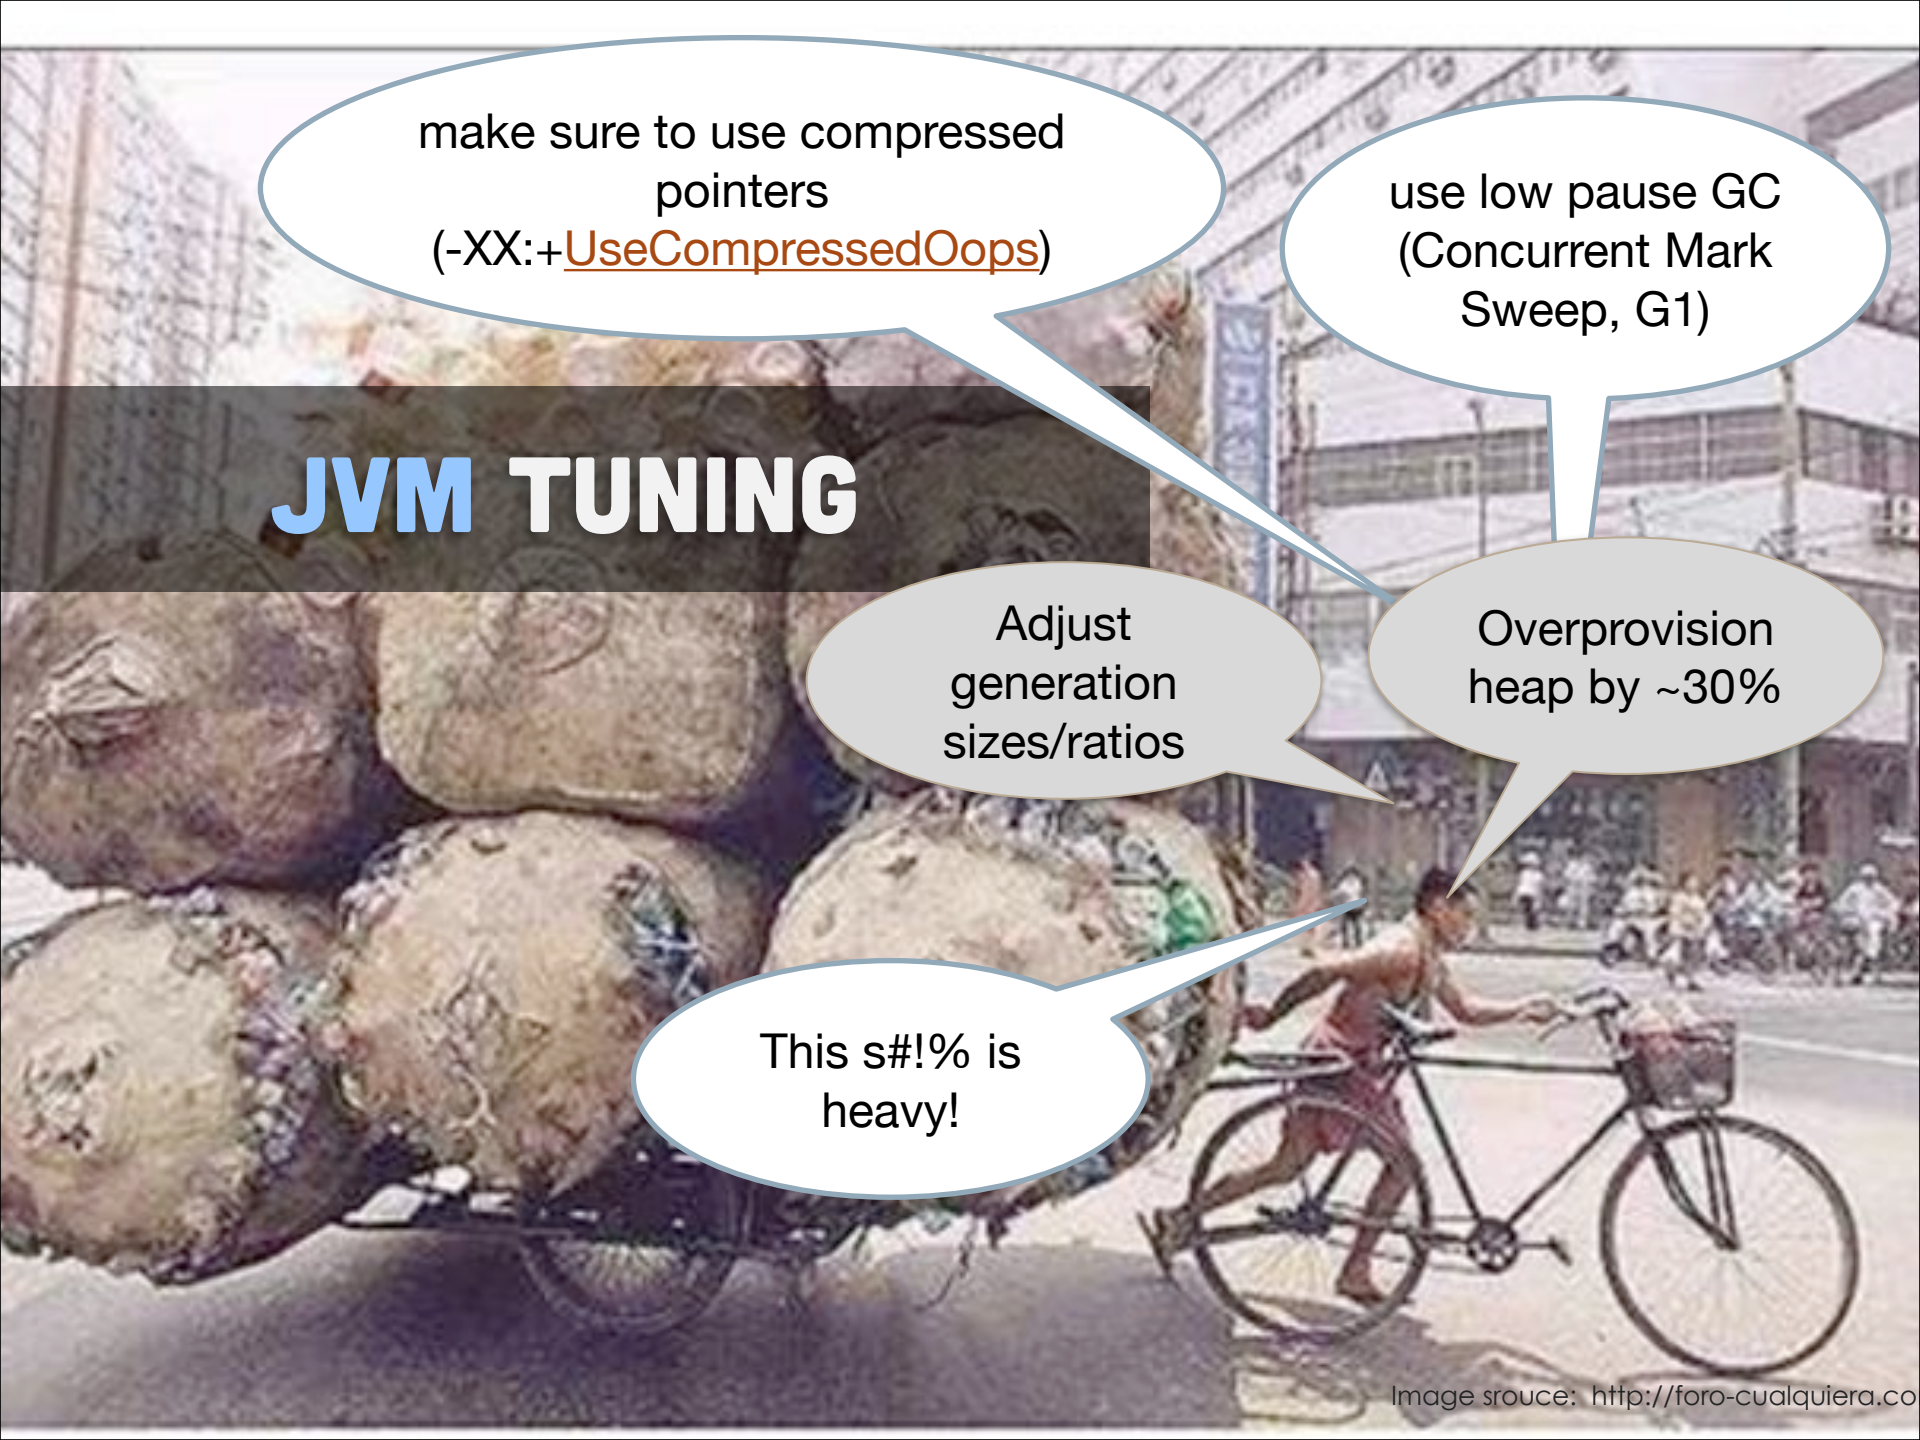


make sure to use compressed
pointers
(-XX:+UseCompressedOops)

use low pause GC
(Concurrent Mark
Sweep, G1)

JVM TUNING

This s#!% is
heavy!



make sure to use compressed pointers
(-XX:+UseCompressedOops)

use low pause GC
(Concurrent Mark Sweep, G1)


JVM TUNING

Adjust generation sizes/ratios


Overprovision heap by ~30%

This s#!% is heavy!

JVM TUNING



JVM TUNING



JVM TUNING



Print garbage
collection



JVM TUNING

Print garbage
collection

If GC pauses still
prohibitive then
consider partitioning



IN SUMMARY



KNOW YOUR BUSINESS



KNOW YOUR DATA



KNOW WHEN TO OPTIMIZE



THE END



My company: <https://popforms.com>
My website: <http://katemats.com>

And much of the credit for this talk goes to Leon Stein for developing the technology. **Thank you, Leon.**

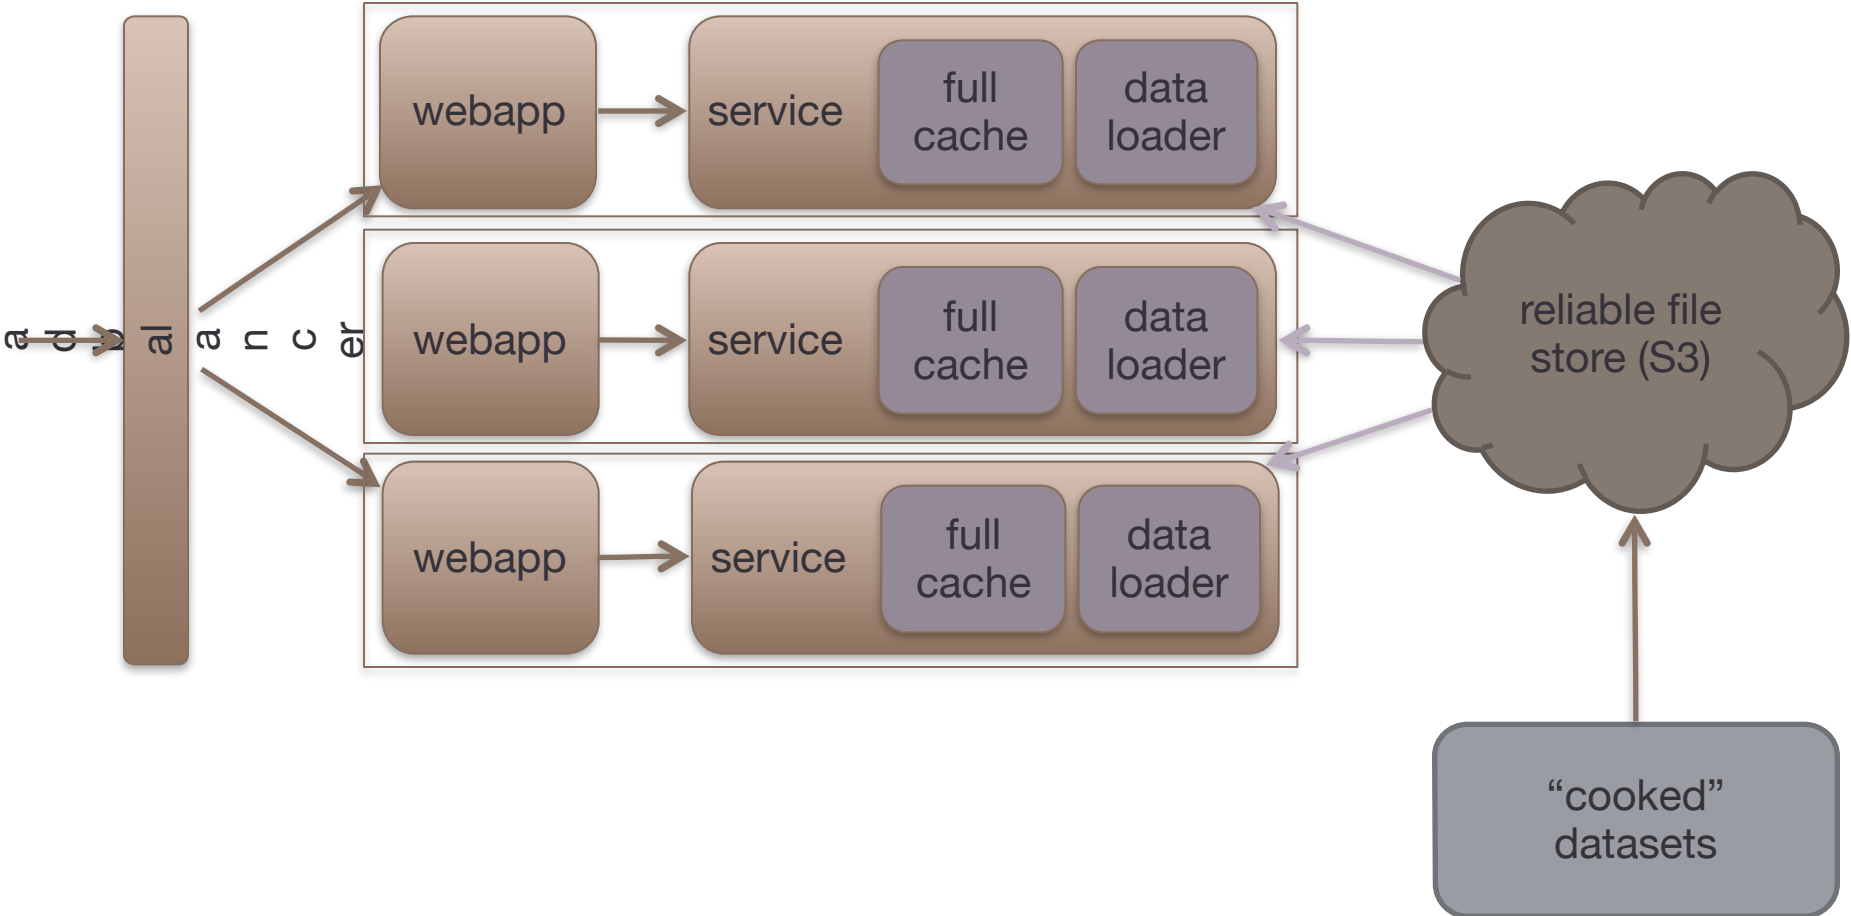
HOW DO YOU LOAD THE DATA?

PEPSI

FAIL

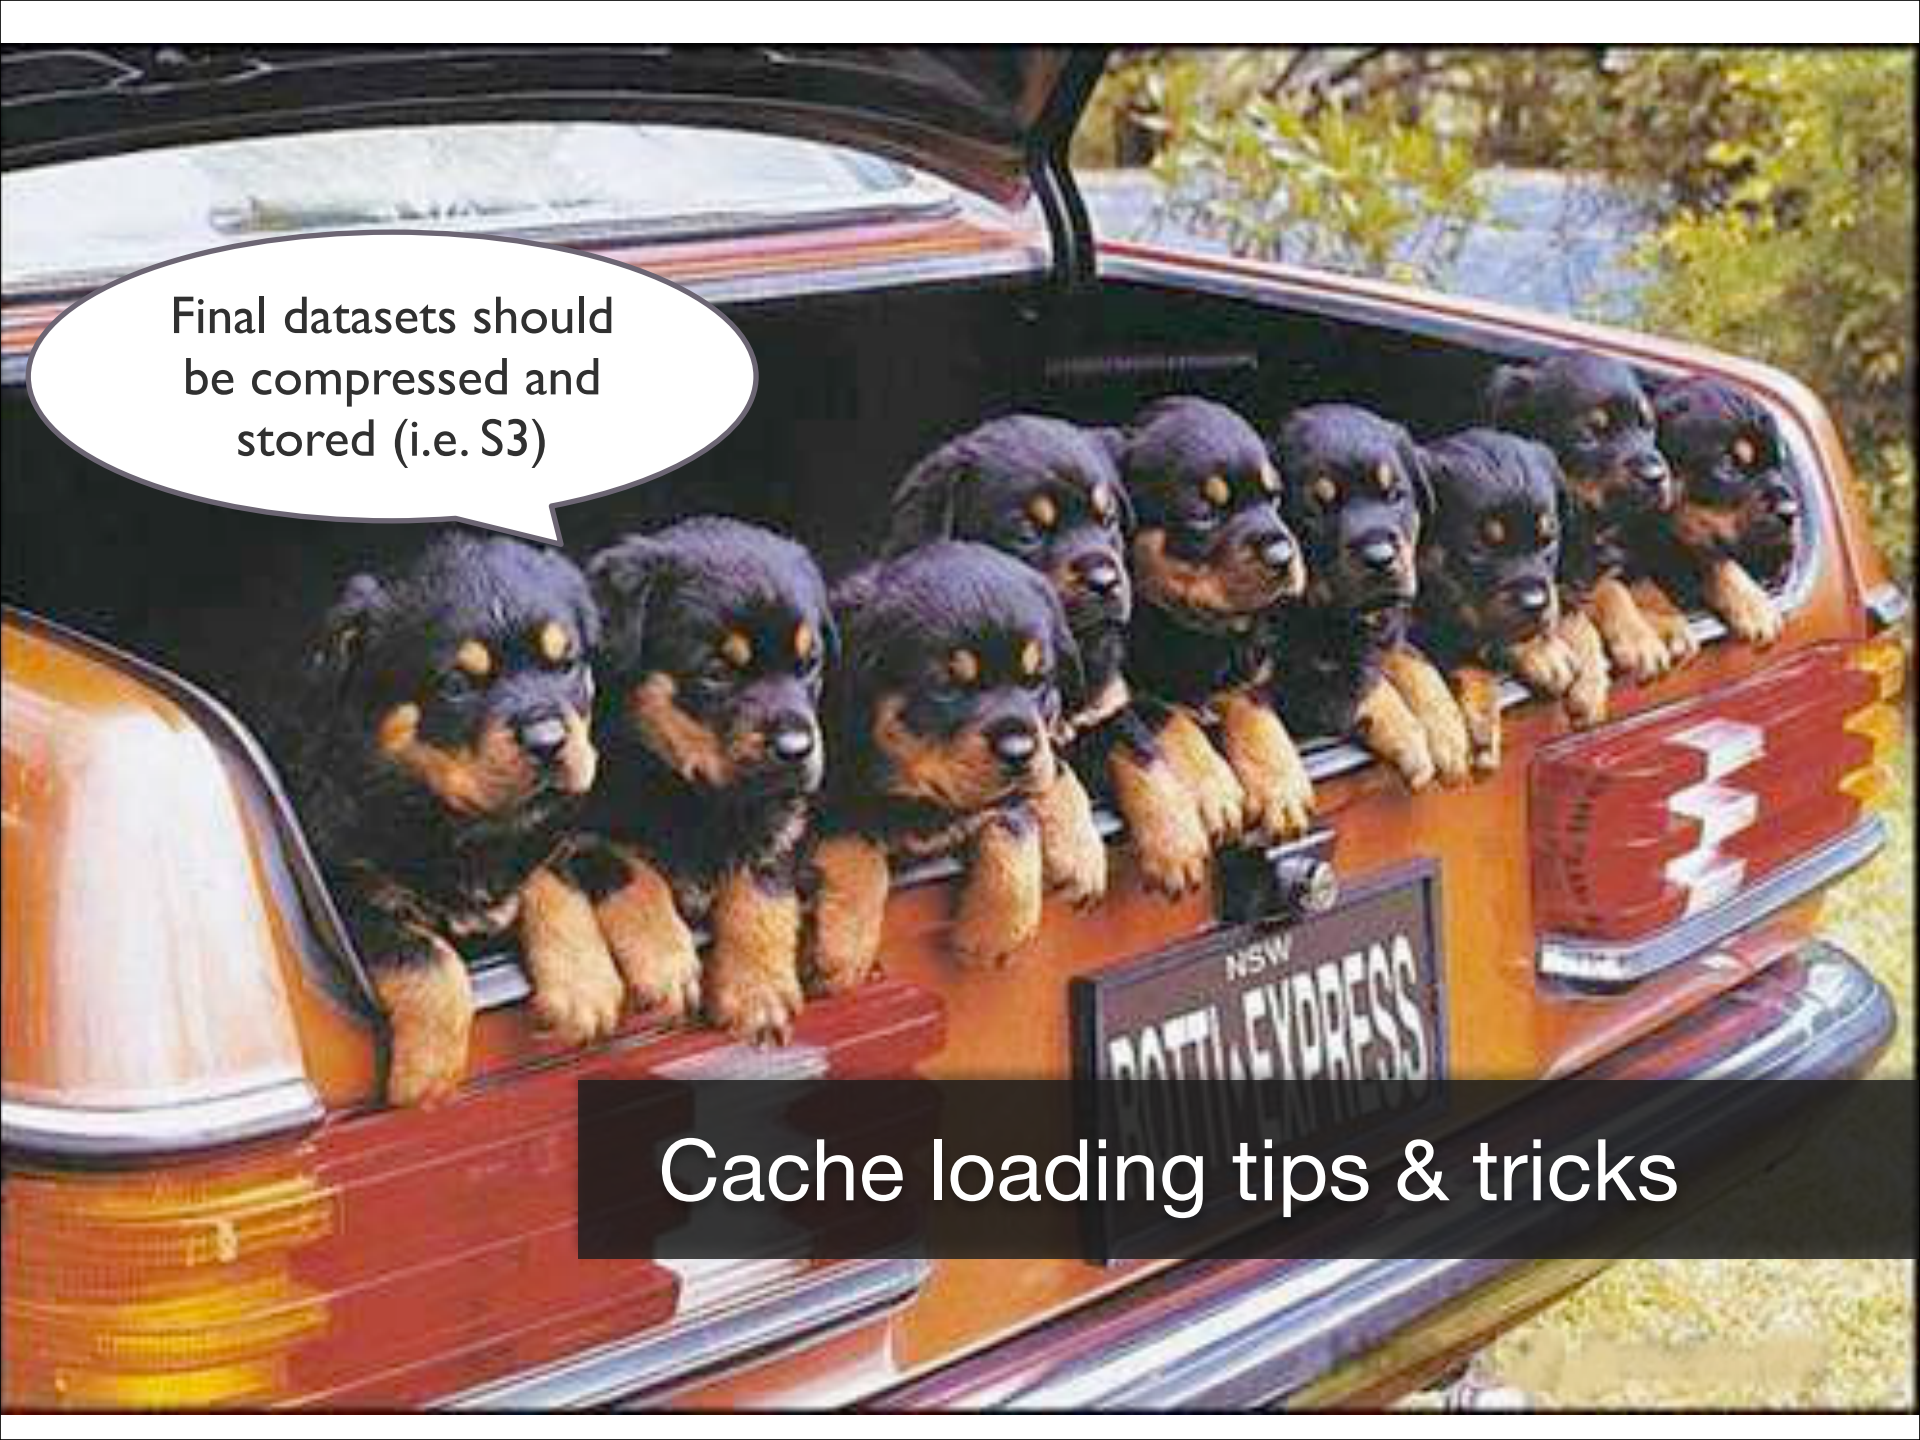


CACHE LOADING





Cache loading tips & tricks



Final datasets should be compressed and stored (i.e. S3)

Cache loading tips & tricks

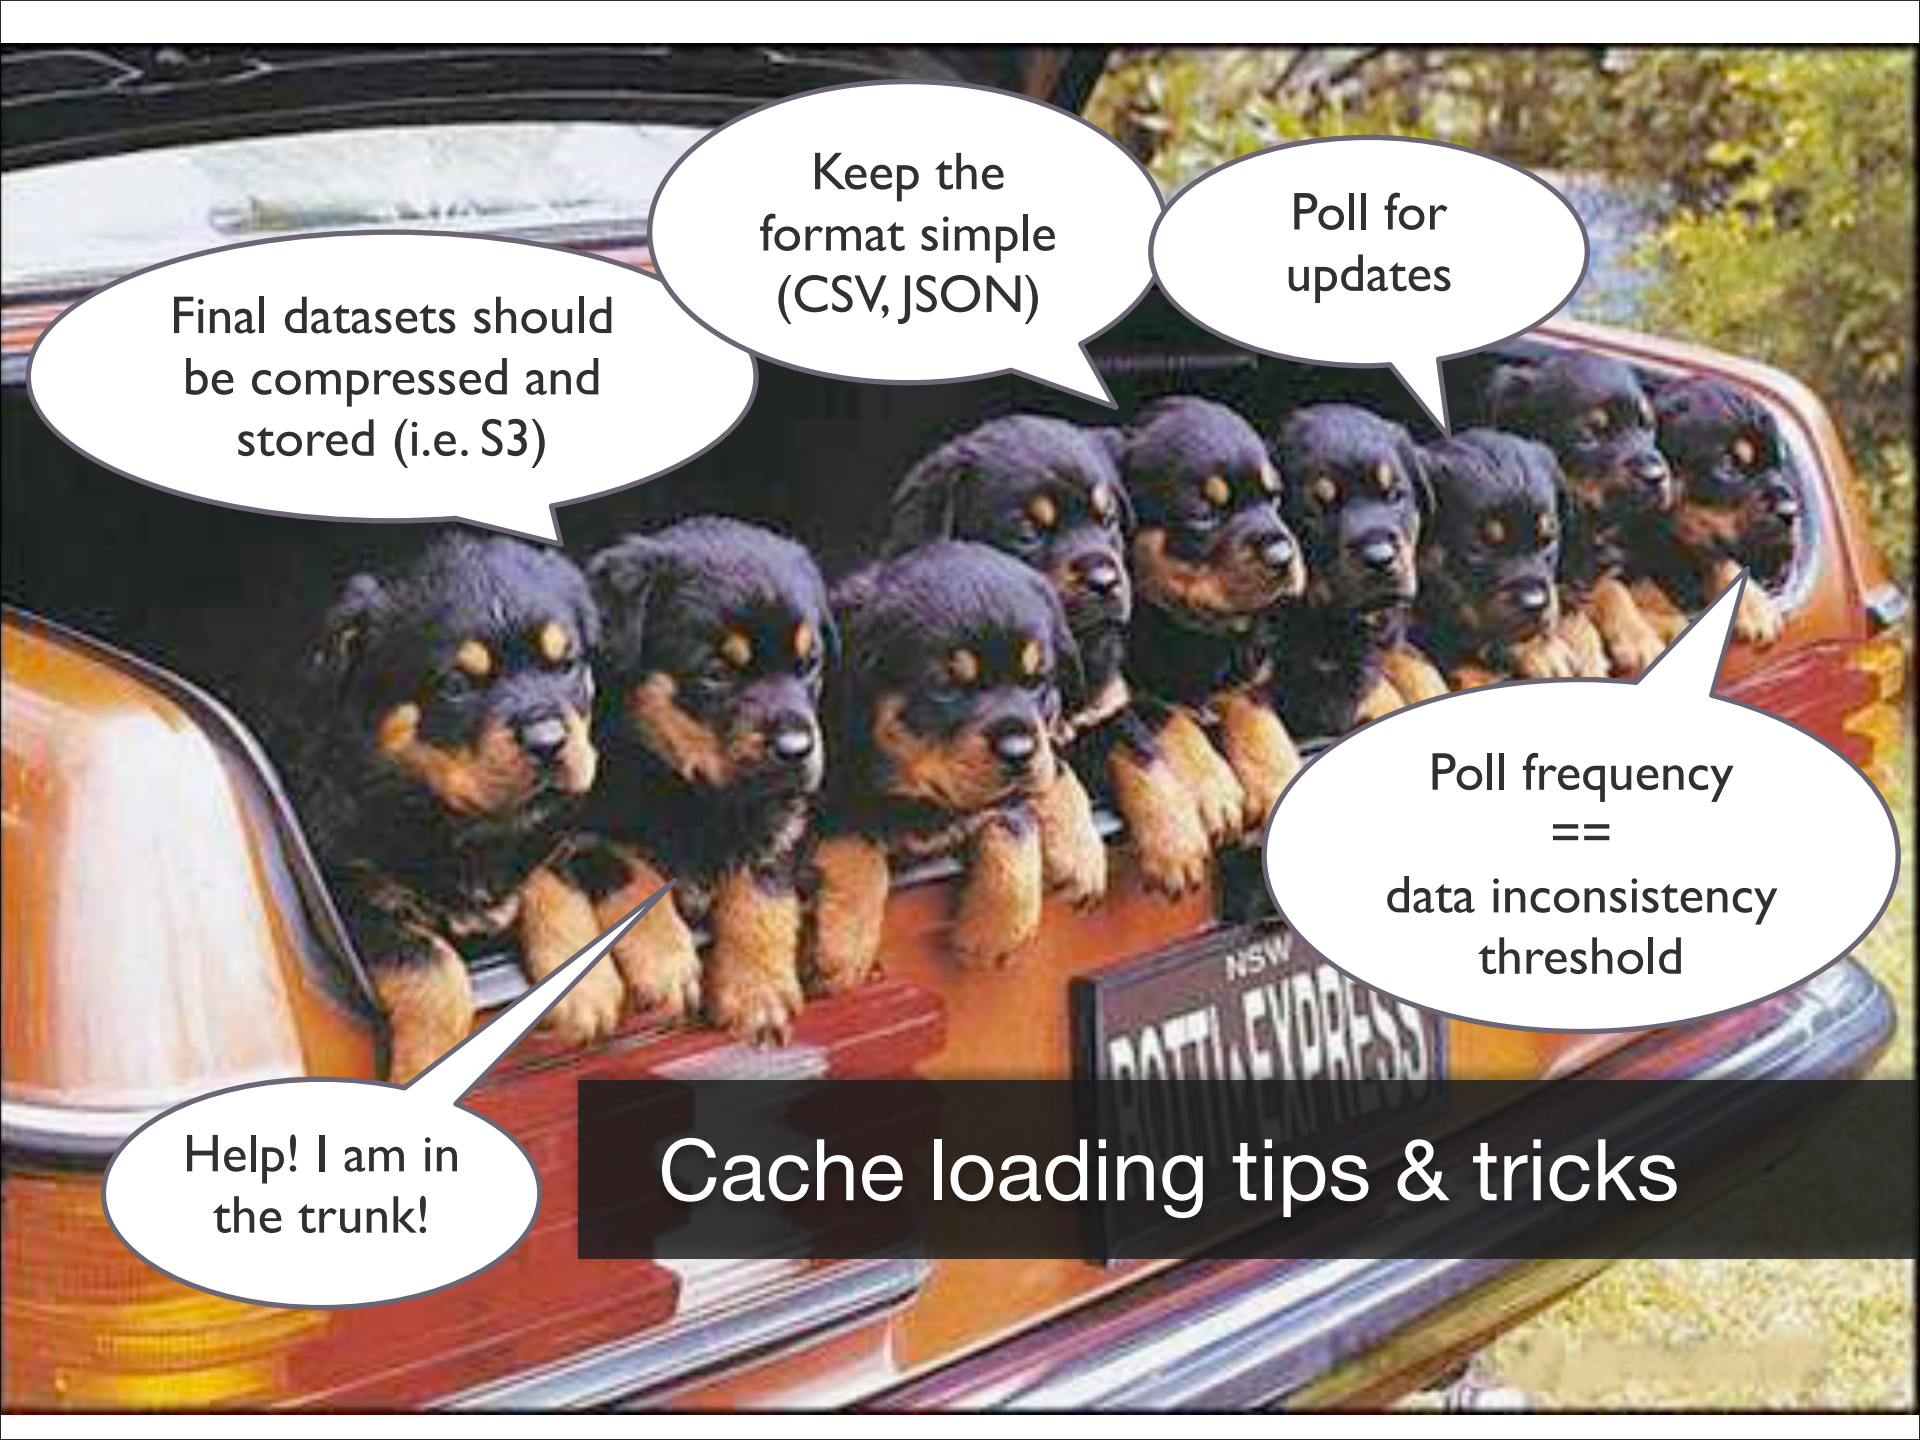


Final datasets should be compressed and stored (i.e. S3)

Keep the format simple (CSV, JSON)

Help! I am in the trunk!

Cache loading tips & tricks



Final datasets should be compressed and stored (i.e. S3)

Keep the format simple (CSV, JSON)

Poll for updates

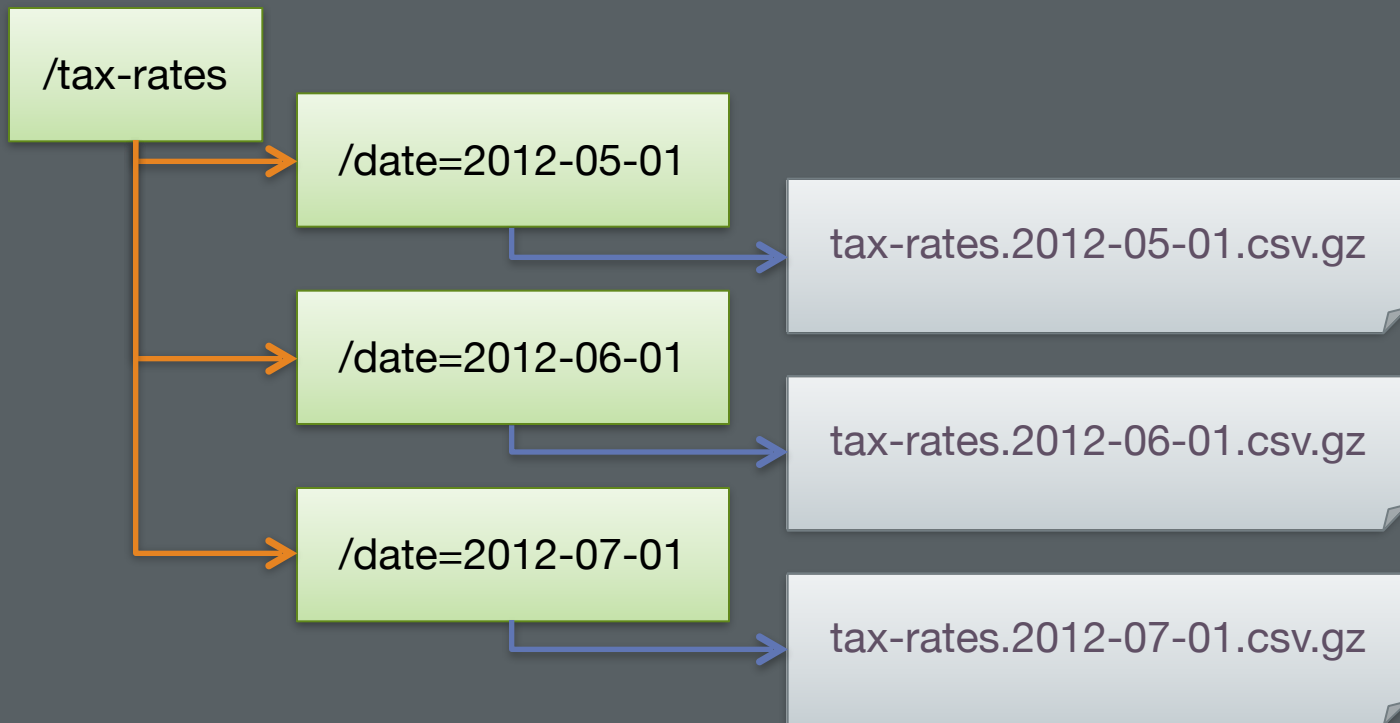
Poll frequency
==
data inconsistency
threshold

Help! I am in the trunk!

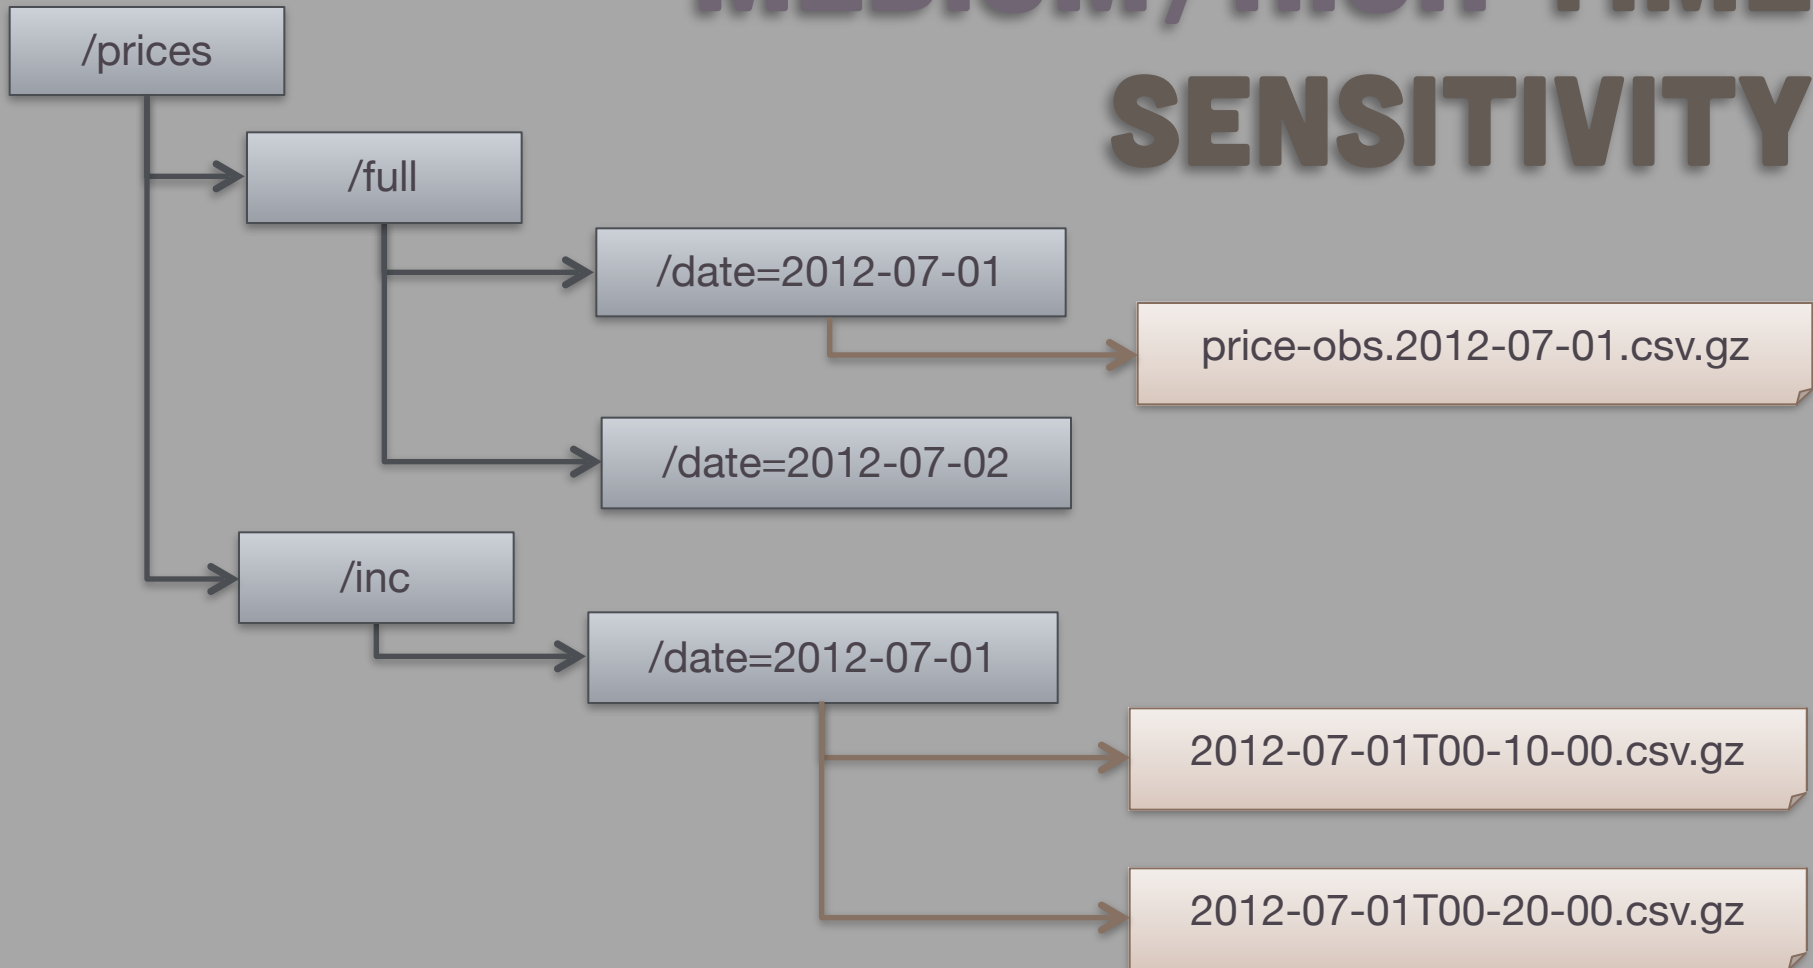
Cache loading tips & tricks

CACHE LOADING TIME SENSITIVITY

CACHE LOADING: LOW TIME SENSITIVITY DATA



CACHE LOADING: MEDIUM / HIGH TIME SENSITIVITY



CACHE LOADING STRATEGY SWAP



Image src:http://static.fjcdn.com/pictures/funny_22d73a_372351.jpg

Cache is immutable,
so no locking is
required

CACHE LOADING STRATEGY SWAP



Image src:http://static.fjcdn.com/pictures/funny_22d73a_372351.jpg

CACHE LOADING STRATEGY SWAP

Cache is immutable,
so no locking is
required

meow

Works well for
infrequently
updated data sets

And for datasets that
need to be refreshed
each update



**HAVE A PIRATE
FOR BREAKFAST**



CACHE LOADING STRATEGY: CRUD

U.S.A.

HAVE A PIRATE
FOR BREAKFAST

Deletions can
be tricky

YARRRRRR!



CACHE LOADING STRATEGY: **CRUD**



Avoid full synchronization

Deletions can be tricky

YARRRRRR!

CACHE LOADING STRATEGY: **CRUD**



Avoid full synchronization

Deletions can be tricky

Consider loading cache in small batches. Use one container per partition

YARRRRRR!

CACHE LOADING STRATEGY: **CRUD**

CONCURRENT LOCKING WITH *Trove Map*

```
public class LongCache<V> {  
    private TLongObjectMap<V> map = new TLongObjectHashMap<V>();  
    private ReentrantReadWriteLock lock = new ReentrantReadWriteLock();  
    private Lock r = lock.readLock(), w = lock.writeLock();  
    public V get(long k) {  
        r.lock();  
        try { return map.get(k); } finally { r.unlock(); }  
    }  
    public V update(long k, V v) {  
        w.lock();  
        try { return map.put(k, v); } finally { w.unlock(); }  
    }  
    public V remove(long k) {  
        w.lock();  
        try { return map.remove(k); } finally { w.unlock(); }  
    }  
}
```




CACHE LOADING OPTIMIZATIONS



I am “cooking”
the data sets.
Ha!



Keep local
copies



Periodically
generate serialized
data/state



CACHE LOADING OPTIMIZATIONS



I am “cooking”
the data sets.
Ha!

Validate with
CRC or hash



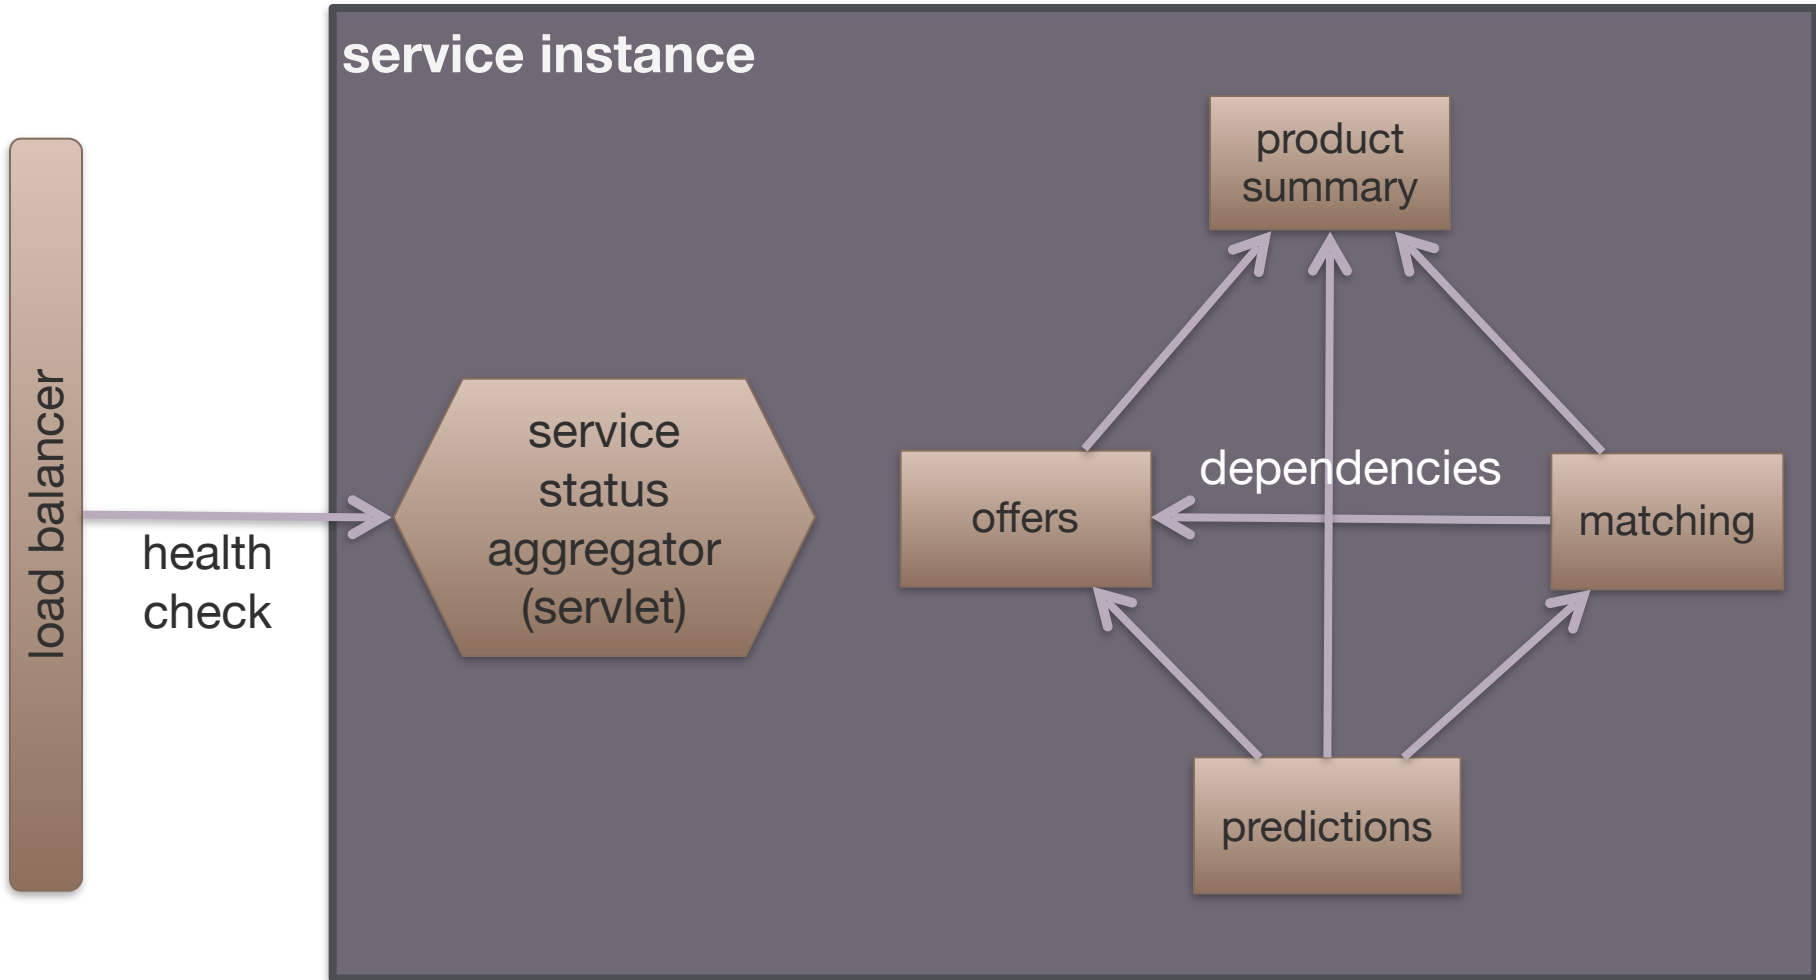
Keep local
copies

Periodically
generate serialized
data/state

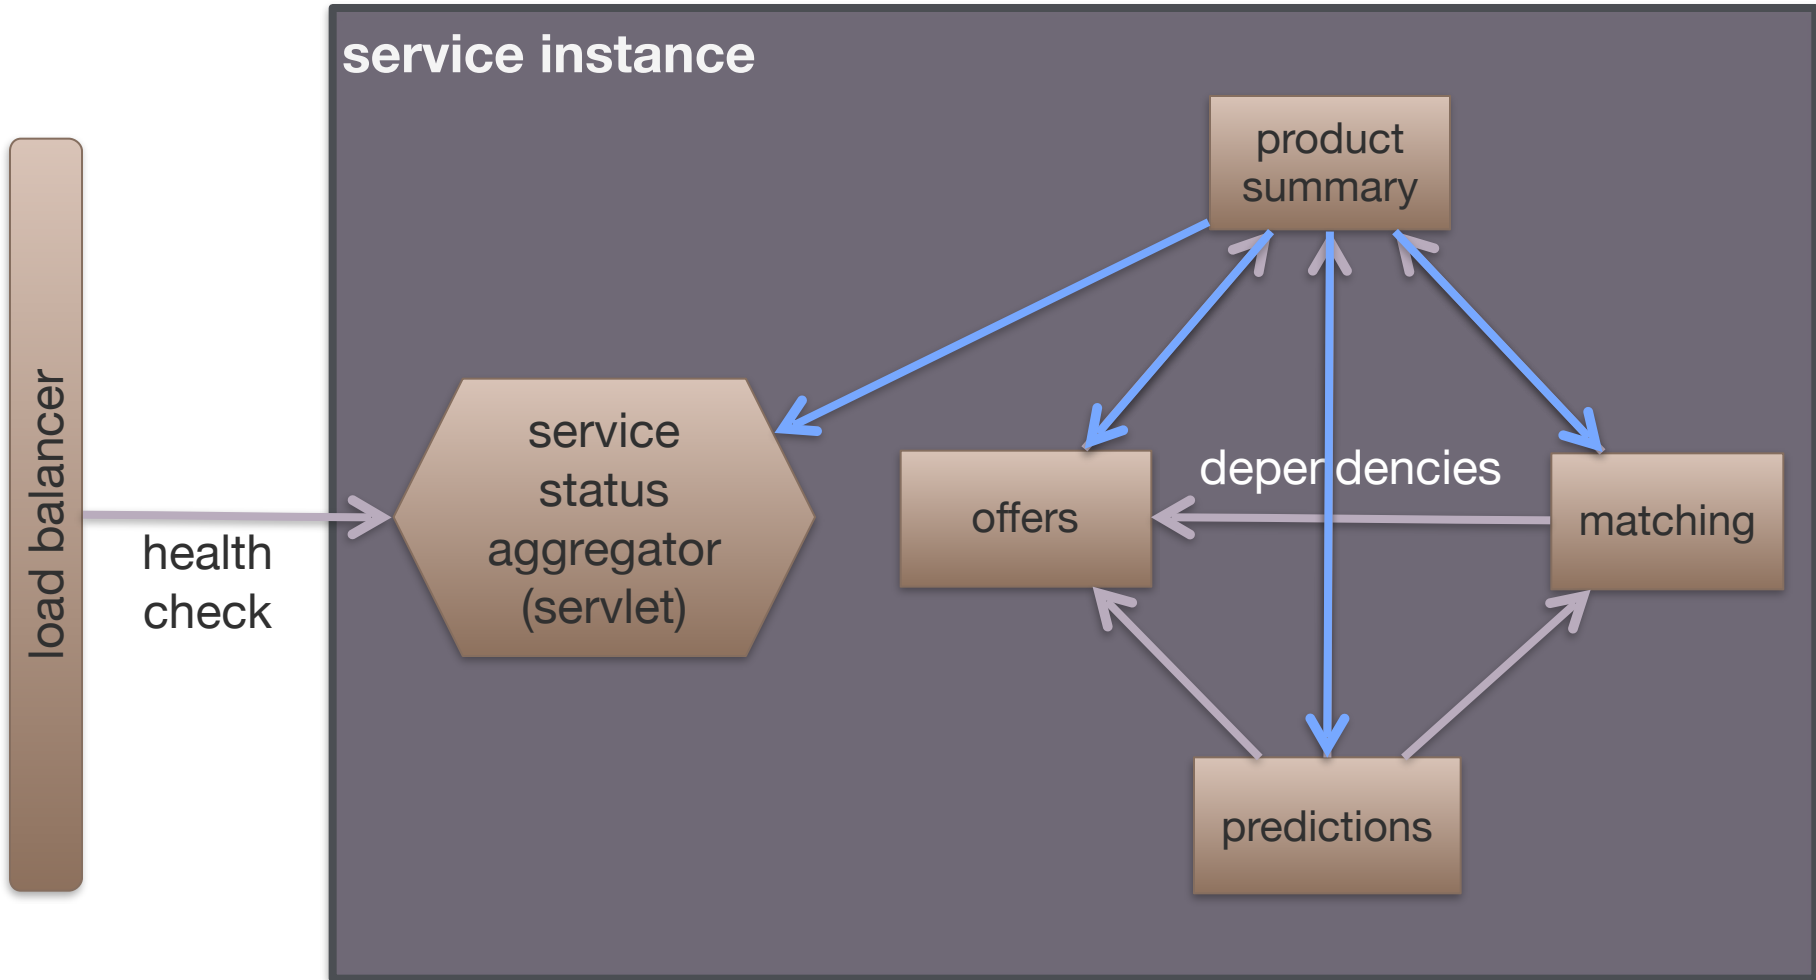


CACHE LOADING OPTIMIZATIONS

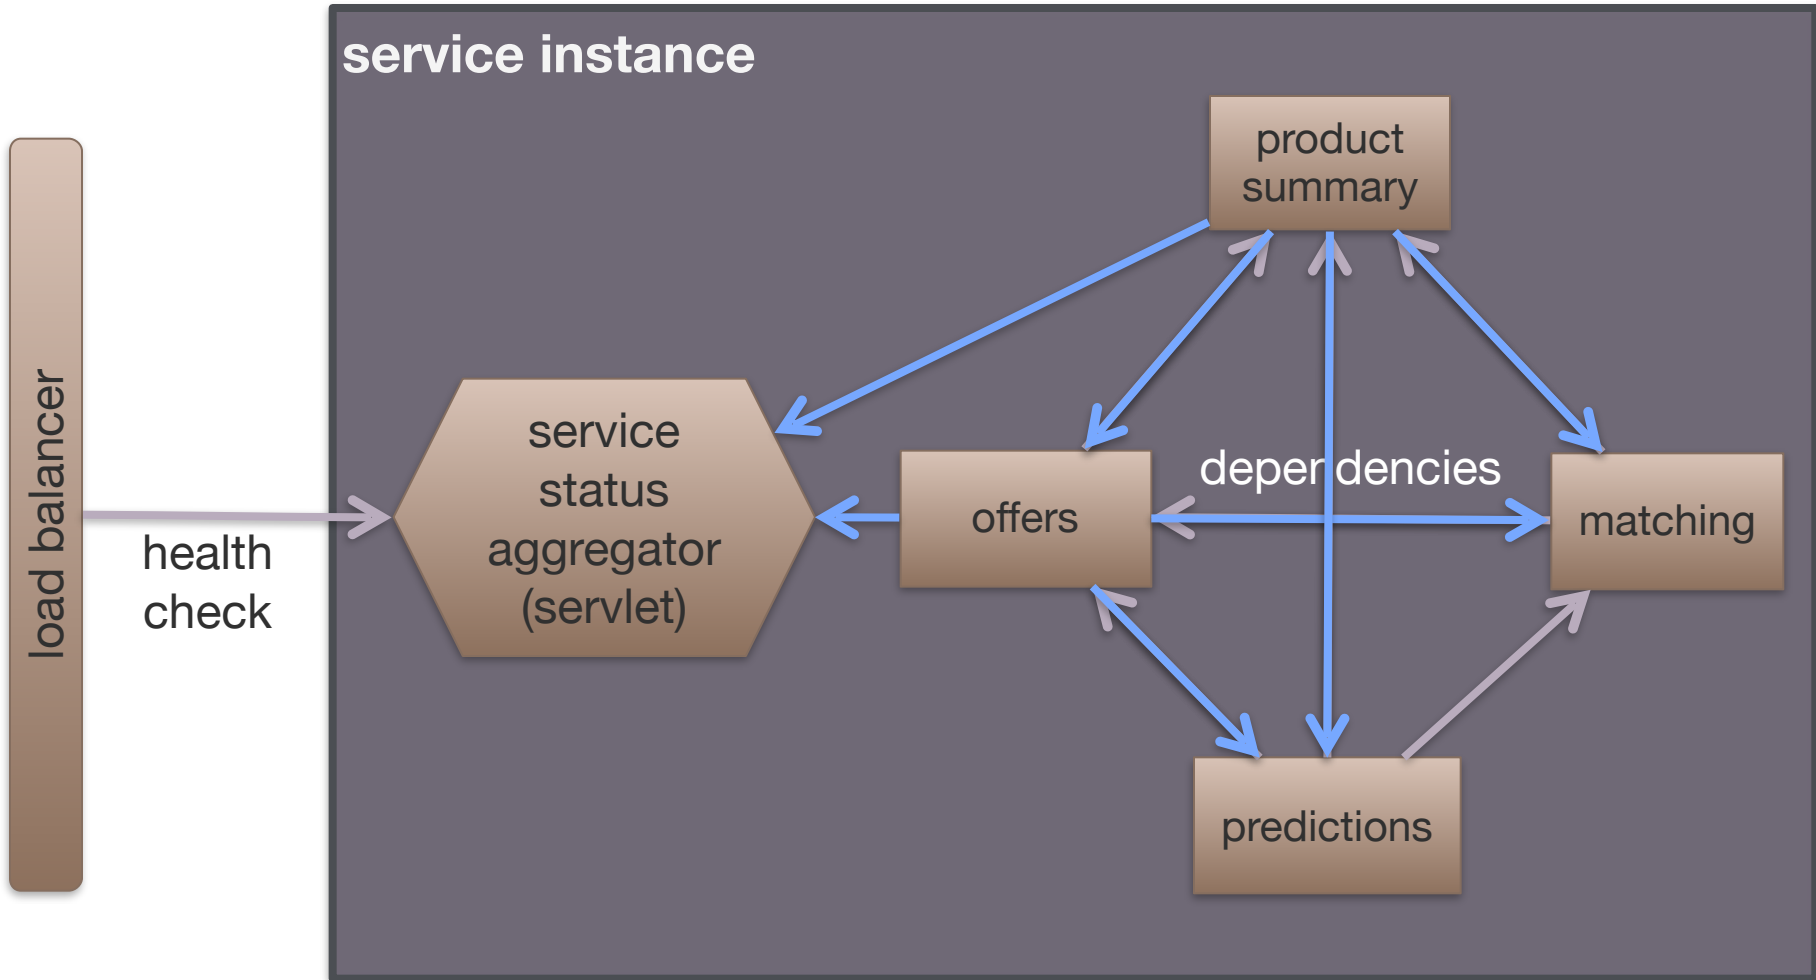
DEPENDENT CACHES



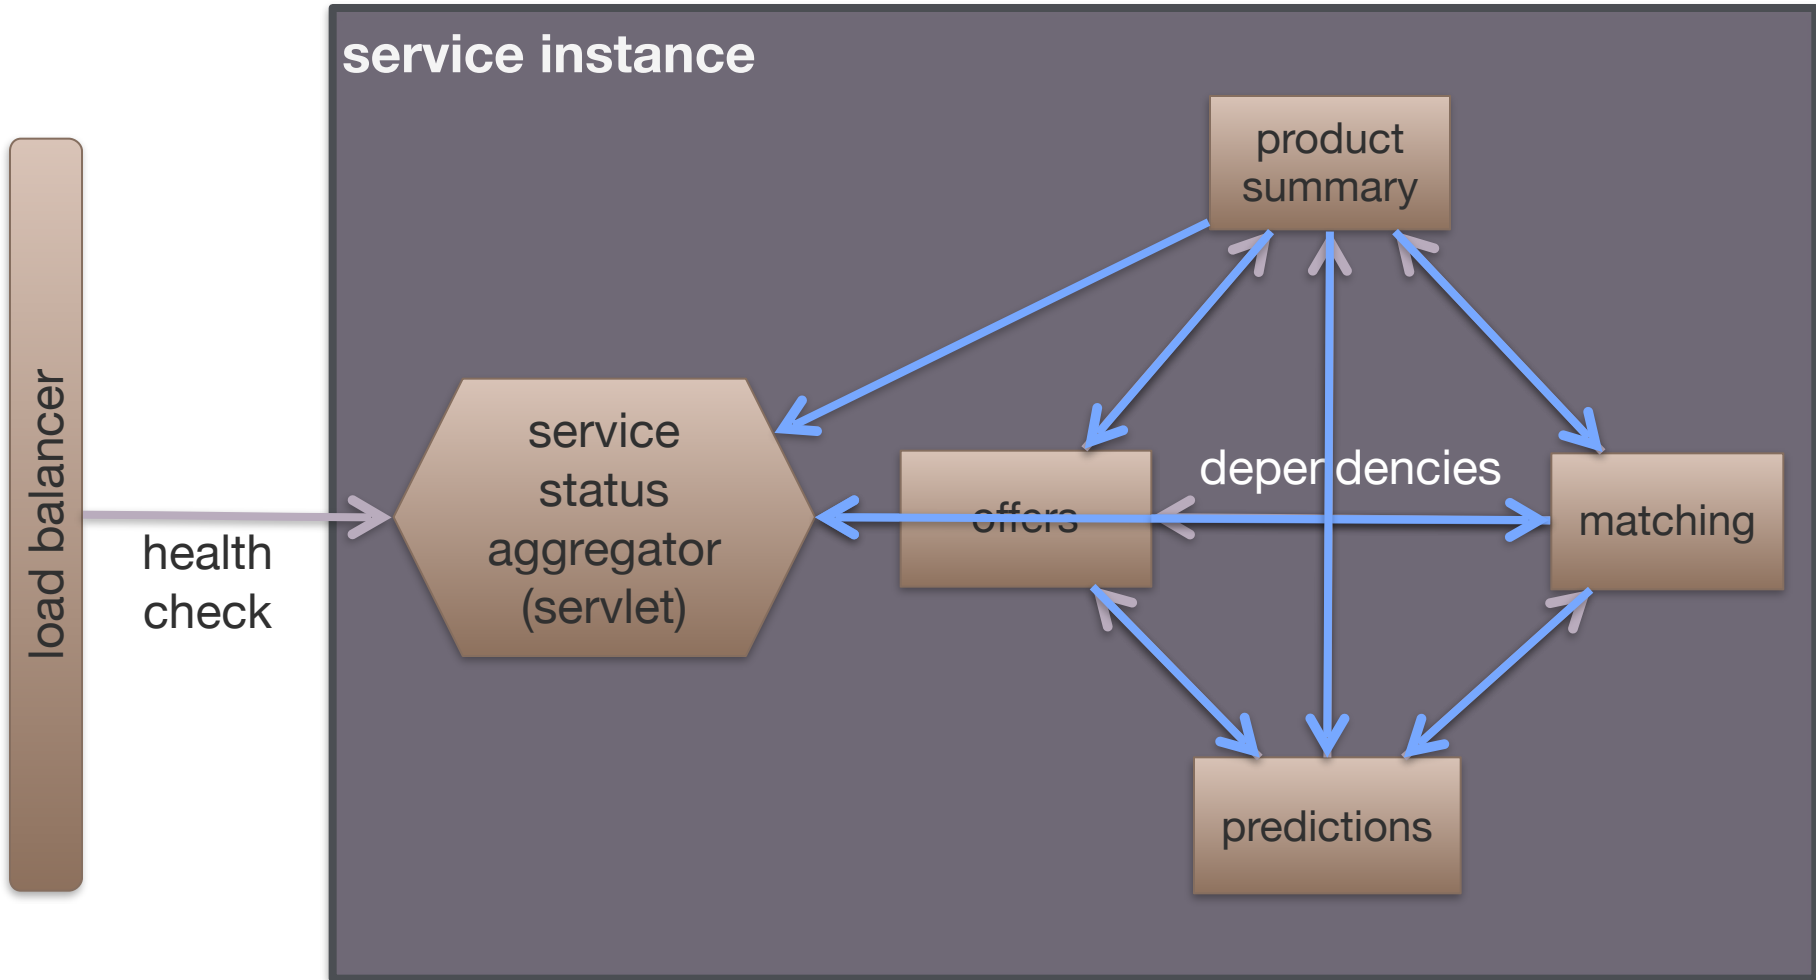
DEPENDENT CACHES



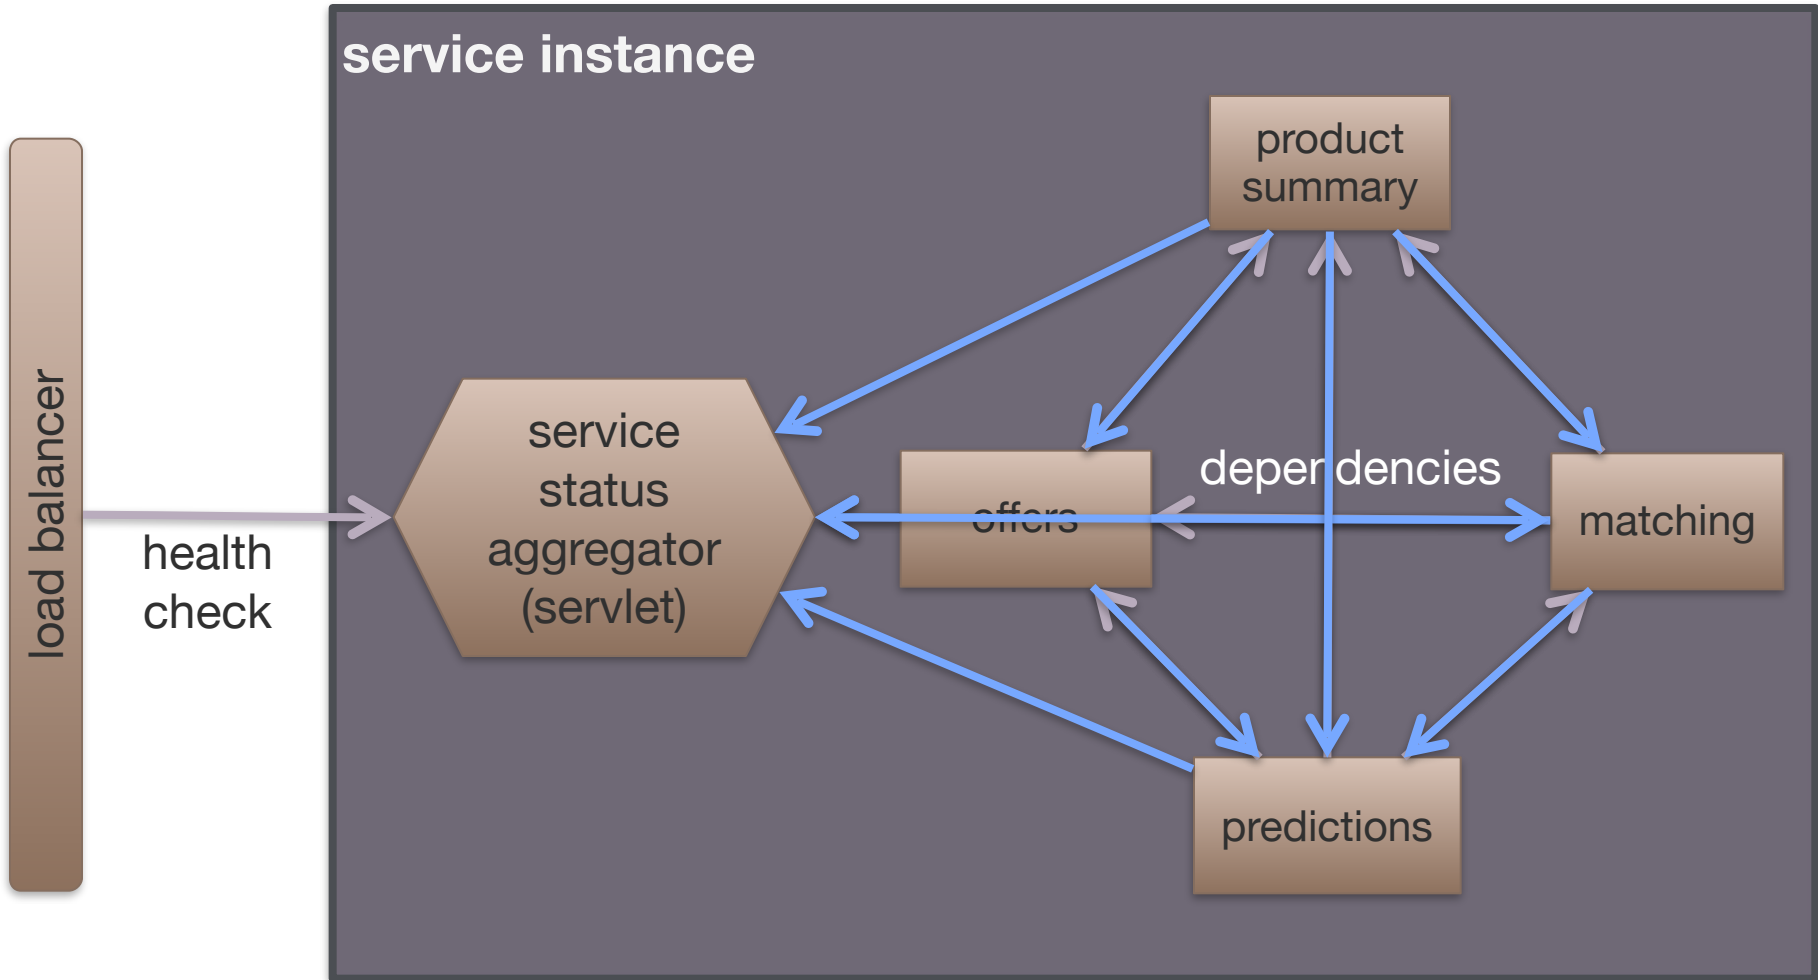
DEPENDENT CACHES



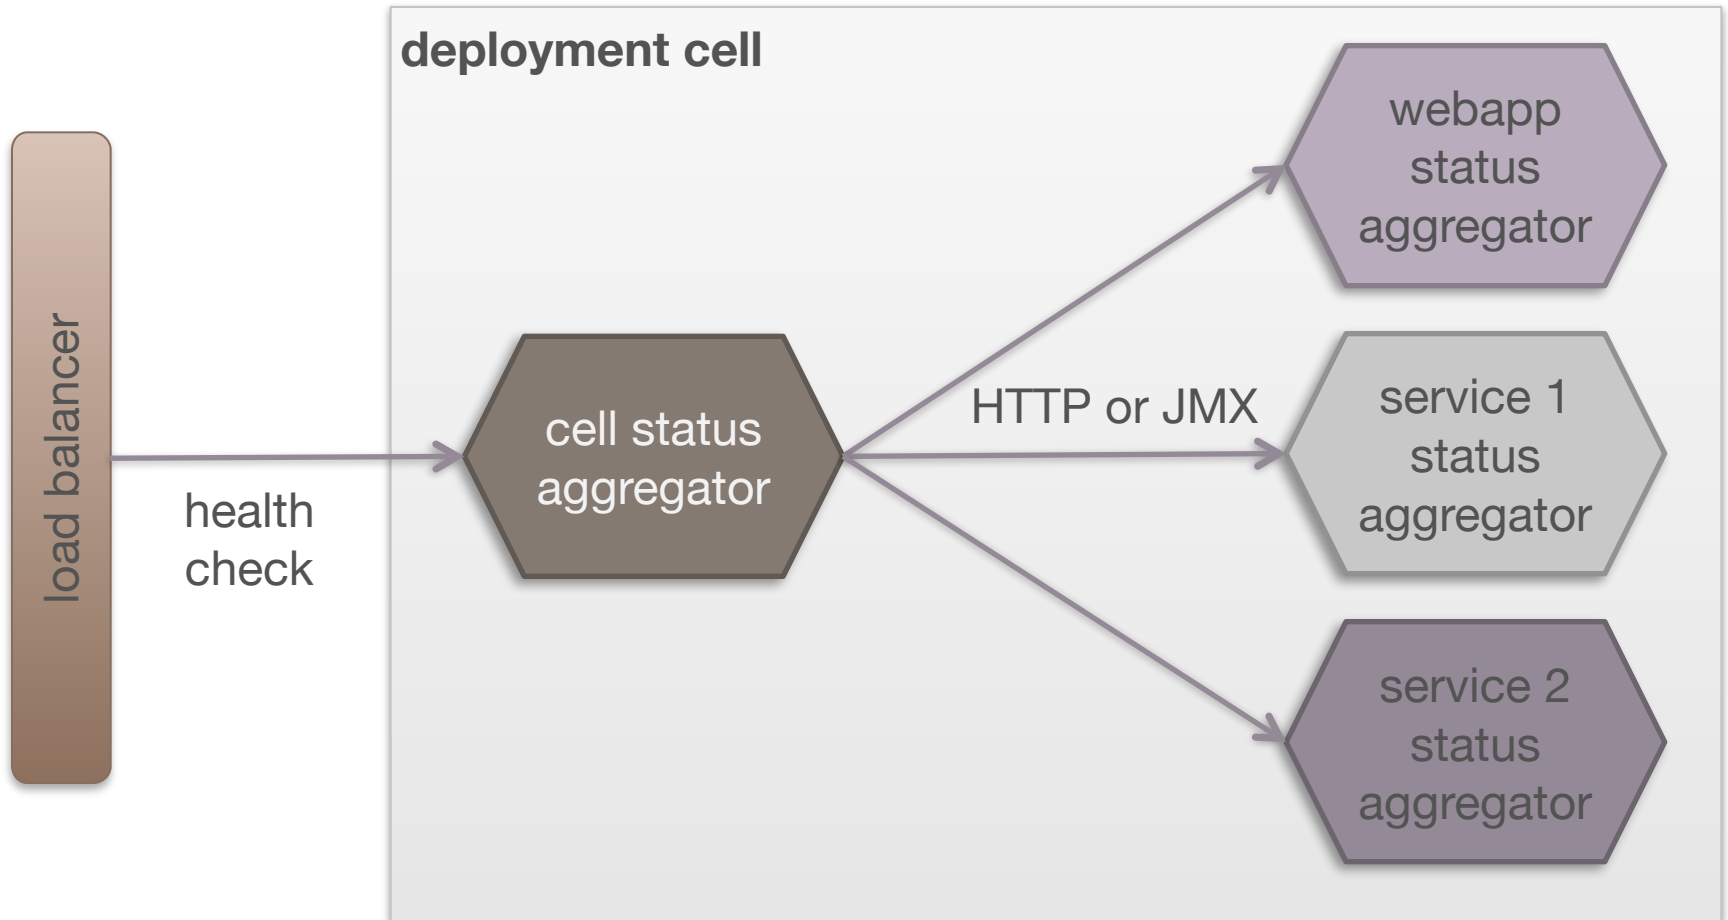
DEPENDENT CACHES



DEPENDENT CACHES



DEPLOYMENT CELL STATUS



HIERARCHICAL STATUS AGGREGATION

Environment Status: prod

Status:	READY
Status Time:	2012-07-12 19:43:18 +0000
Startup Time:	2012-06-18 18:49:03 +0000
Current Time:	2012-07-12 19:43:18 +0000

Detail

	Bean Name	Version	State	Timestamp	Message
1	frontend@prod-us-east-01.decideit.org				
2	frontend@prod-us-east-02.decideit.org		READY	2012-07-12 19:43:18 +0000	NORMAL
3	frontend@prod-us-east-03.decideit.org		READY	2012-07-12 19:43:18 +0000	NORMAL
4	frontend@prod-us-east-04.decideit.org		READY	2012-07-12 19:43:18 +0000	NORMAL
5	instance-salmon@prod-us-east-01.decideit.org		READY	2012-07-12 19:43:18 +0000	
6	instance-salmon@prod-us-east-01.decideit.org		READY	2012-07-12 19:43:18 +0000	

Group Status Status: frontend

Status:	READY
Status Time:	2012-07-12 19:42:38 +0000
Startup Time:	2012-06-18 18:49:03 +0000
Current Time:	2012-07-12 19:42:39 +0000

Detail

	Bean Name	Version	State	Timestamp	Message
1	instance-cod@prod-us-east-01.decideit.org:11	4.0.10	READY	2012-07-12 19:42:38 +0000	NORMAL
2	instance-hulk@localvm	5.0.7.1	READY	2012-07-12 19:42:38 +0000	
3	instance-salmon@prod-us-east-01.decideit.org:5	4.0.12.1	READY	2012-07-12 19:42:38 +0000	NORMAL
4	instance-tron@prod-us-east-01.decideit.org:2	2.0.15	READY	2012-07-12 19:42:38 +0000	
5	instance-yna@prod-us-east-01.decideit.org:3	3.0.12	READY	2012-07-12 19:42:38 +0000	

Servlet Context Status: /product-ws

Status:	READY
Status Time:	2012-07-12 19:40:08 +0000
Startup Time:	2012-06-01 20:58:03 +0000
Current Time:	2012-07-12 19:40:08 +0000

Detail

	Bean Name	Version	State	Timestamp	Message
1	details-service.namedExpressionService		READY	2012-06-01 20:59:01 +0000	
2	geo-service.ipGeoLocationService		READY	2012-07-10 15:45:17 +0000	
3	product-service.channelProductMappingService		READY	2012-06-01 20:59:00 +0000	
4	product-service.decideitBankService		READY	2012-07-12 19:33:45 +0000	
5	product-service.featurePriorityService		READY	2012-07-12 19:33:43 +0000	
6	product-service.modelHistoryService		READY	2012-07-12 19:29:22 +0000	
7	product-service.priorPredictionTableService		READY	2012-07-12 19:25:02 +0000	NORMAL
8	product-service.productEventService		READY	2012-07-12 19:38:05 +0000	
9	product-service.productFeatureRatingService		READY	2012-07-12 19:33:43 +0000	NORMAL
10	product-service.productInfoService		READY	2012-07-12 19:38:15 +0000	
11	product-service.productMappingService		READY	2012-07-12 19:33:43 +0000	NORMAL
12	product-service.productOfferService		READY	2012-07-12 19:35:56 +0000	NORMAL
13	product-service.productPromotionService		READY	2012-06-01 21:02:02 +0000	
14	product-service.productRelationService		READY	2012-07-12 19:33:43 +0000	NORMAL
15	product-service.productReviewSummaryService		READY	2012-07-12 19:33:43 +0000	NORMAL
16	product-service.productSpecificationSchemaService		READY	2012-07-12 19:31:30 +0000	
17	product-service.productSpecificationService		READY	2012-07-12 19:31:30 +0000	
18	product-service.productSummaryService		READY	2012-07-12 19:27:14 +0000	NORMAL
19	product-service.productTitleService		READY	2012-06-01 21:01:03 +0000	
20	product-service.promotionService		READY	2012-06-01 20:59:00 +0000	
21	product-service.numeroService		READY	2012-07-12 19:27:14 +0000	
22	product-service.shippingAndTaxService		READY	2012-07-12 13:03:09 +0000	
23	seller-service.channelReferenceService		READY	2012-07-12 19:38:27 +0000	
24	seller-service.productCategoryReferenceService		READY	2012-07-12 19:38:27 +0000	
25	seller-service.sellerProductURLService		READY	2012-06-07 19:30:04 +0000	
26	seller-service.sellerReferenceService		READY	2012-07-12 19:38:27 +0000	

Instance Status Status: instance-salmon

Status:	READY
Status Time:	2012-07-12 19:44:56 +0000
Startup Time:	2012-06-01 20:58:03 +0000
Current Time:	2012-07-12 19:44:56 +0000

Detail

	Bean Name	Version	State	Timestamp	Message
1	product-ws@localvm		READY	2012-07-12 19:44:56 +0000	NORMAL