

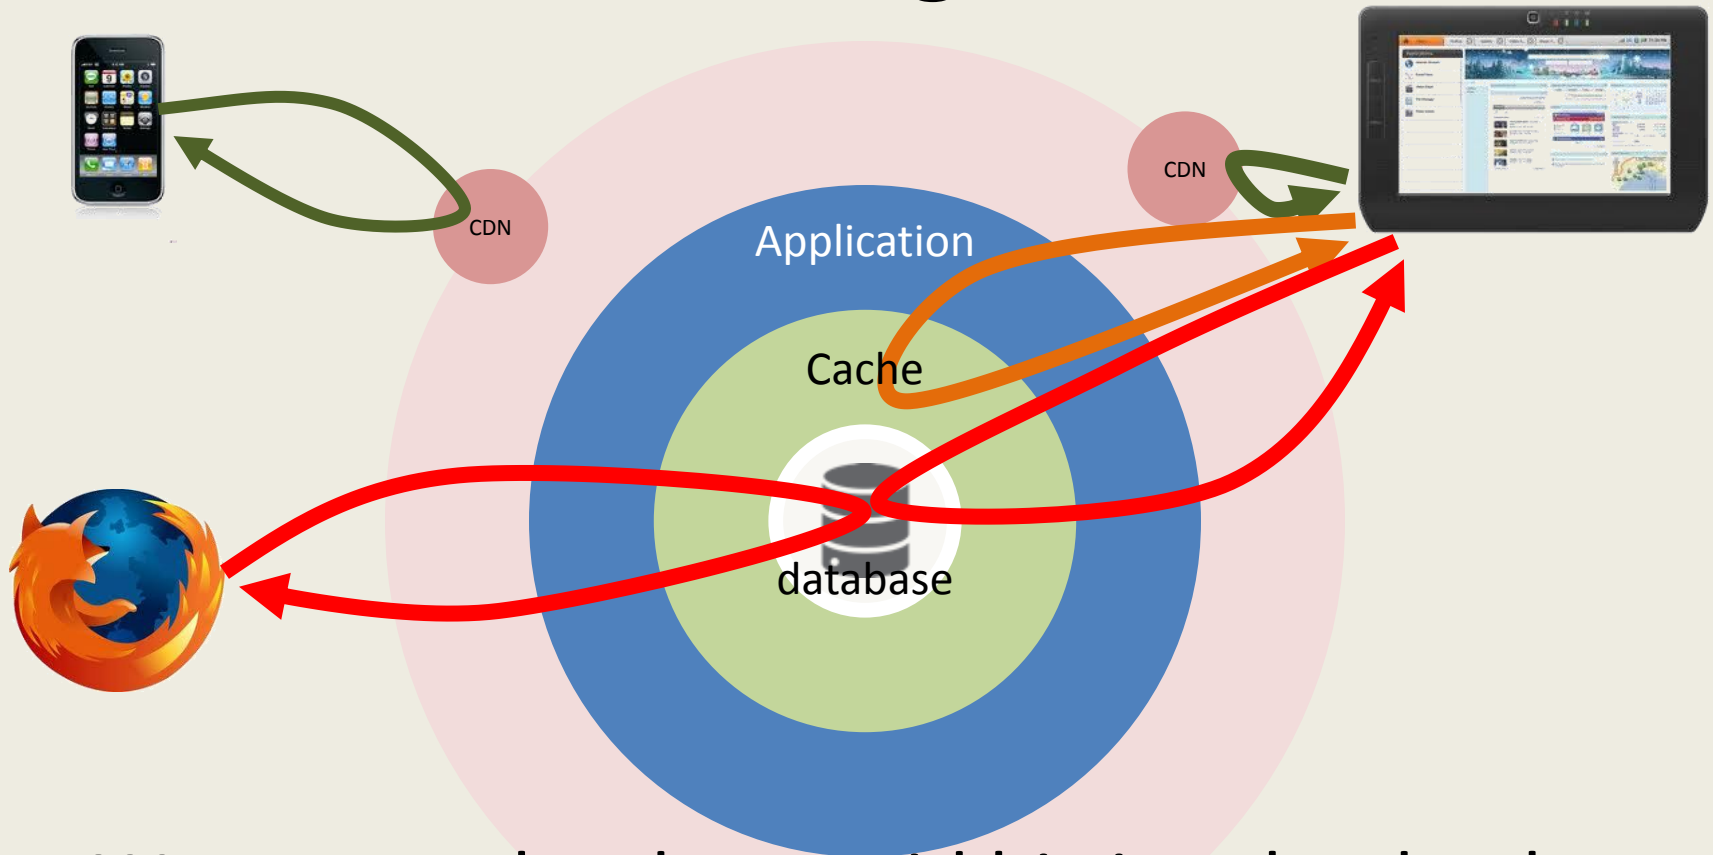
# Database Scalability with Parallel Databases

Boston MySQL Meetup  
Monday, December 10, 2012

# What should I do while you yabber along for 1 hour?

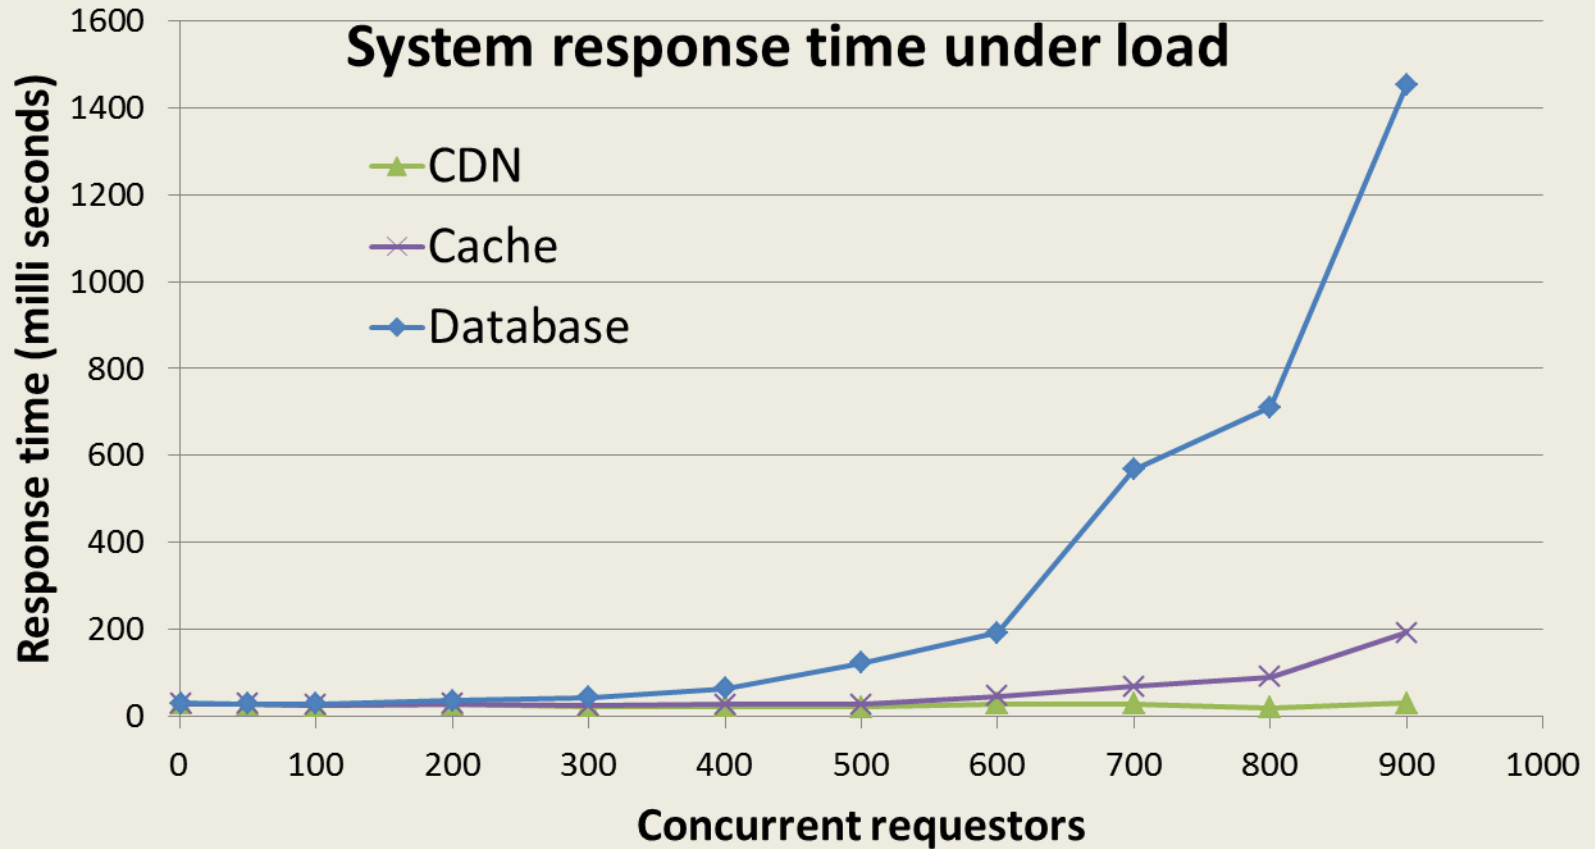
1. Eat Pizza ...
2. Ask questions, and keep me honest
3. Tweet about us, include #parelastic or #mysql and you will be part of the demo!

# The background



- We try very hard to avoid hitting the database
  - But there are times when you cannot avoid it!

# Why?




# Some alternatives

- Scale Up
- Scale Out



# Viability of scale-up (1)

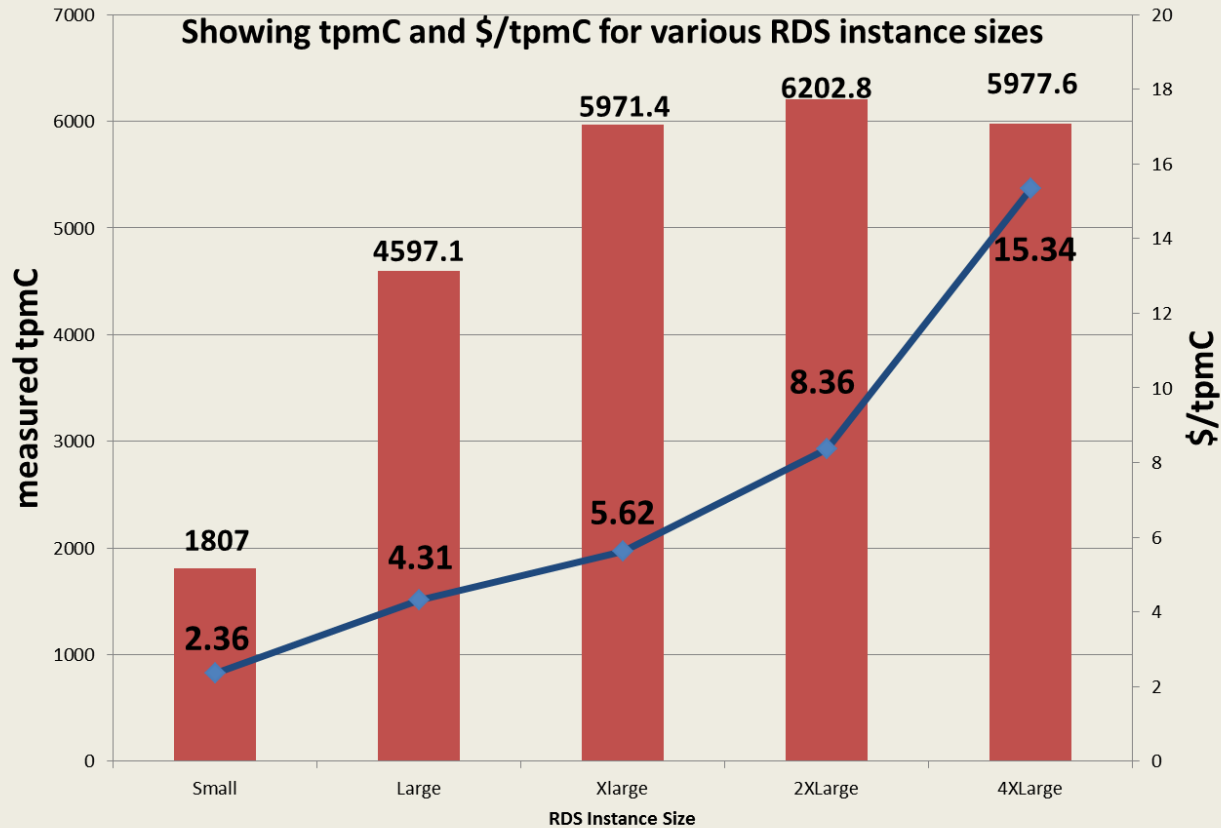
Measured tpmC on Amazon RDS



Source: Amazon RDS TPC-C Benchmark. Md. Borhan Uddin, Bo He, Radu Sion, Cloud Computing Center, SUNY Stony Brook.

Viewed online <http://digitalpiglet.org/research/sion2010cloud-rds.pdf>

# Viability of scale-up (2)



Source: Amazon RDS TPC-C Benchmark. Md. Borhan Uddin, Bo He, Radu Sion, Cloud Computing Center, SUNY Stony Brook.

Viewed online <http://digitalpiglet.org/research/sion2010cloud-rds.pdf>

# Database Scale Out: It's been around!

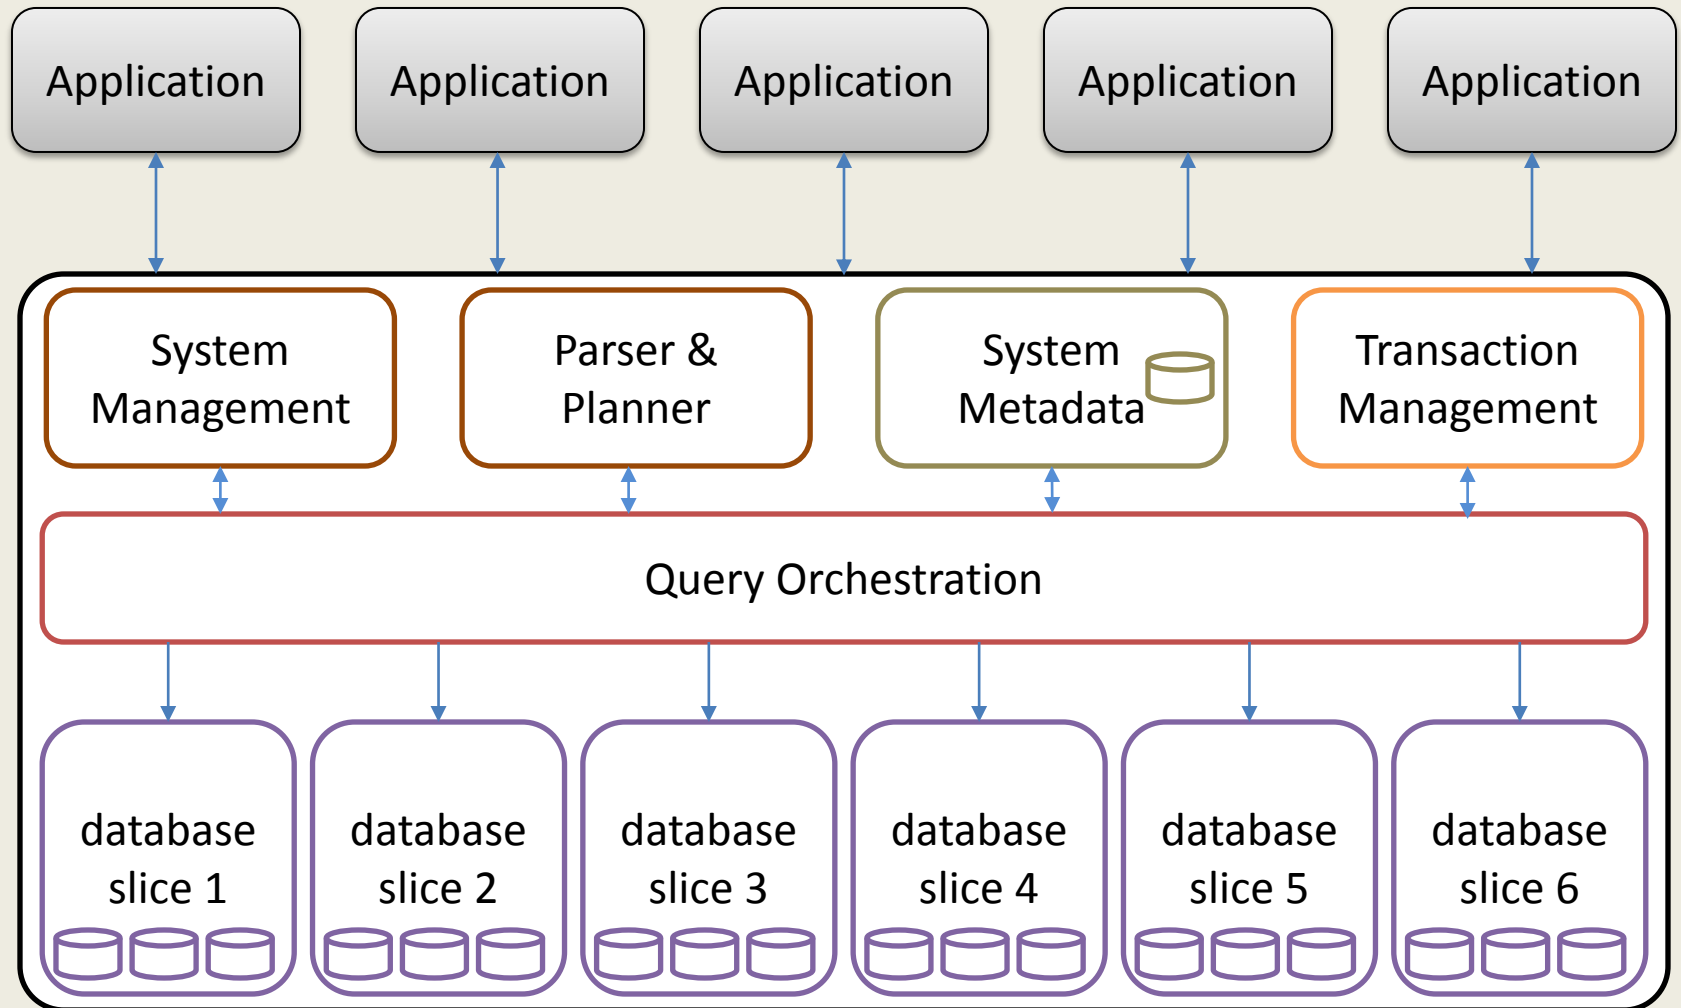
- [1986] The Case for Shared Nothing (Michael Stonebraker)
- [1991] Parallel database systems: the future of high performance database systems (David Dewitt, Jim Gray)
- [2004] MapReduce: simplified data processing on large clusters (Jeffrey Dean, Sanjay Ghemawat)
- [1984] Teradata releases the world's first parallel data warehouses and data marts.
- [1997] Teradata customer creates world's largest production database at 24 terabytes.
- [2001-3] Netezza, Vertica, Greenplum, Datallegro, ...
- [2004] MapReduce
- [2010] ParElastic



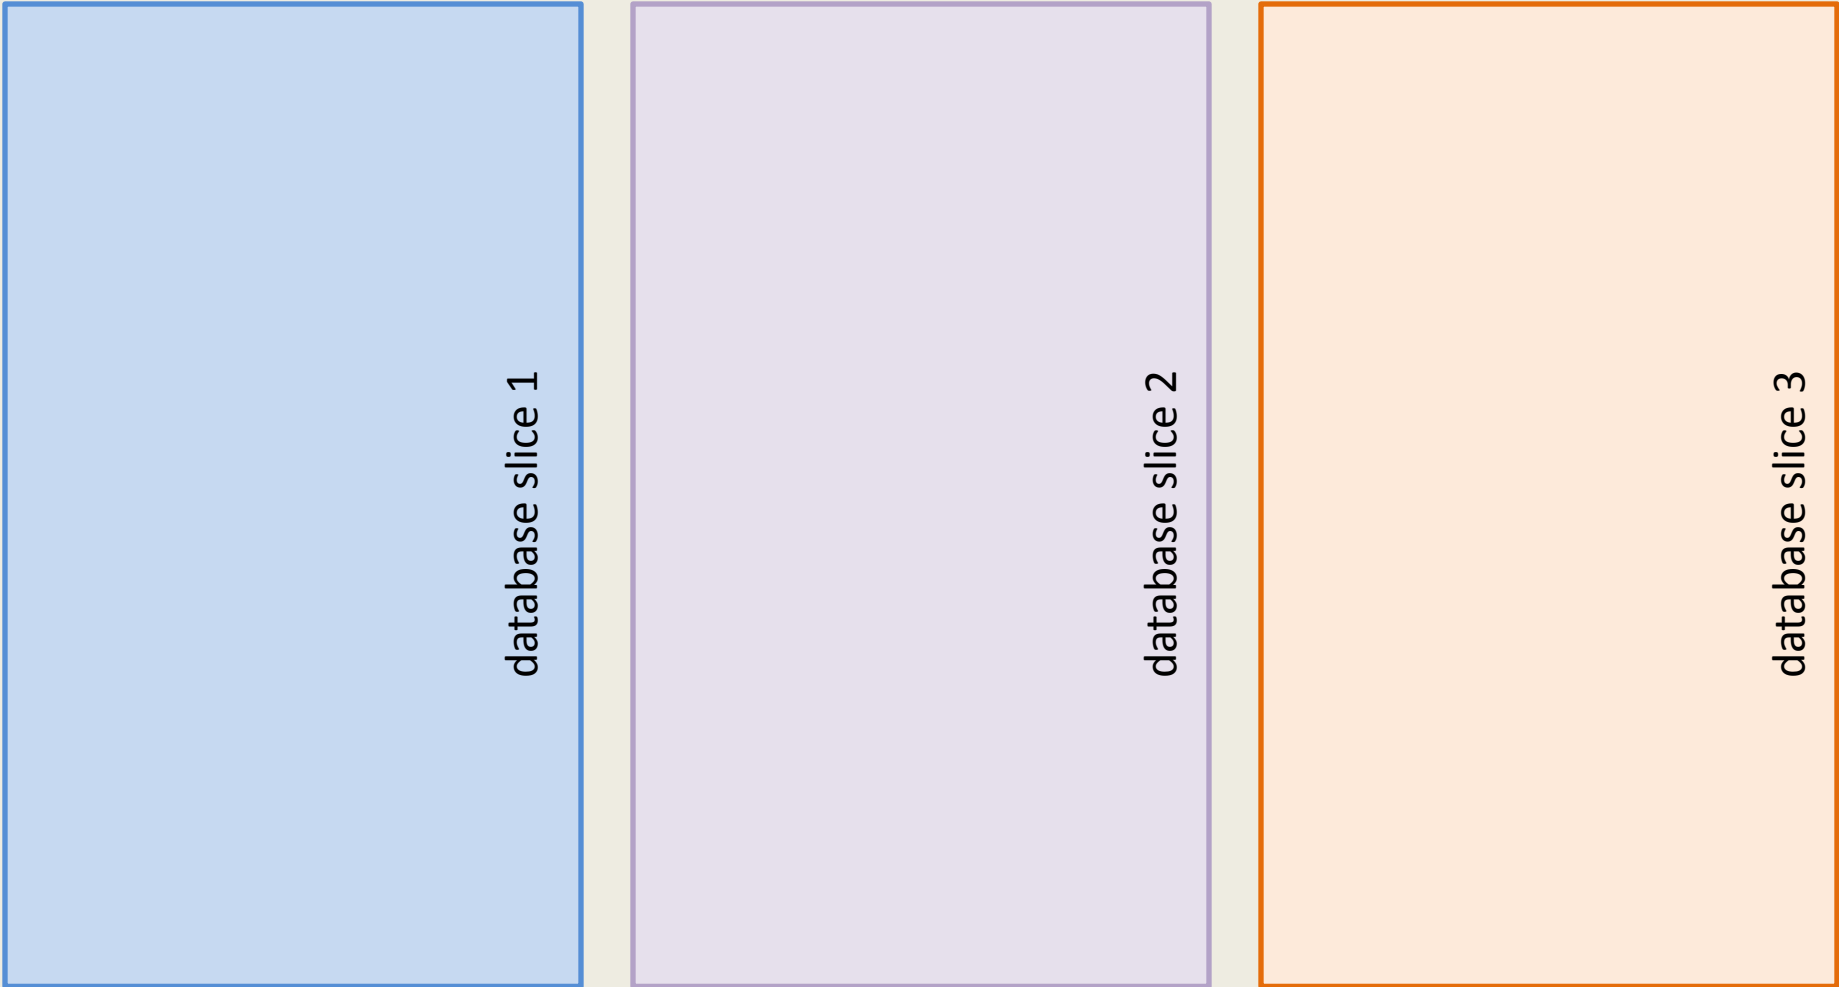
# What is a parallel database?

- A system that seeks to improve performance through the parallelization of various operations such as scanning, restriction, aggregation, sorting, loading, index management, and system recovery.
- Typically based on the “Shared Nothing” architecture
  - Shared disk is also common
  - Shared memory is less common
- Often for Analytics (Netezza, Vertica, ...)
  - But not always (Oracle Exadata)

# A “shared-nothing” parallel database



# How it works: The scenario

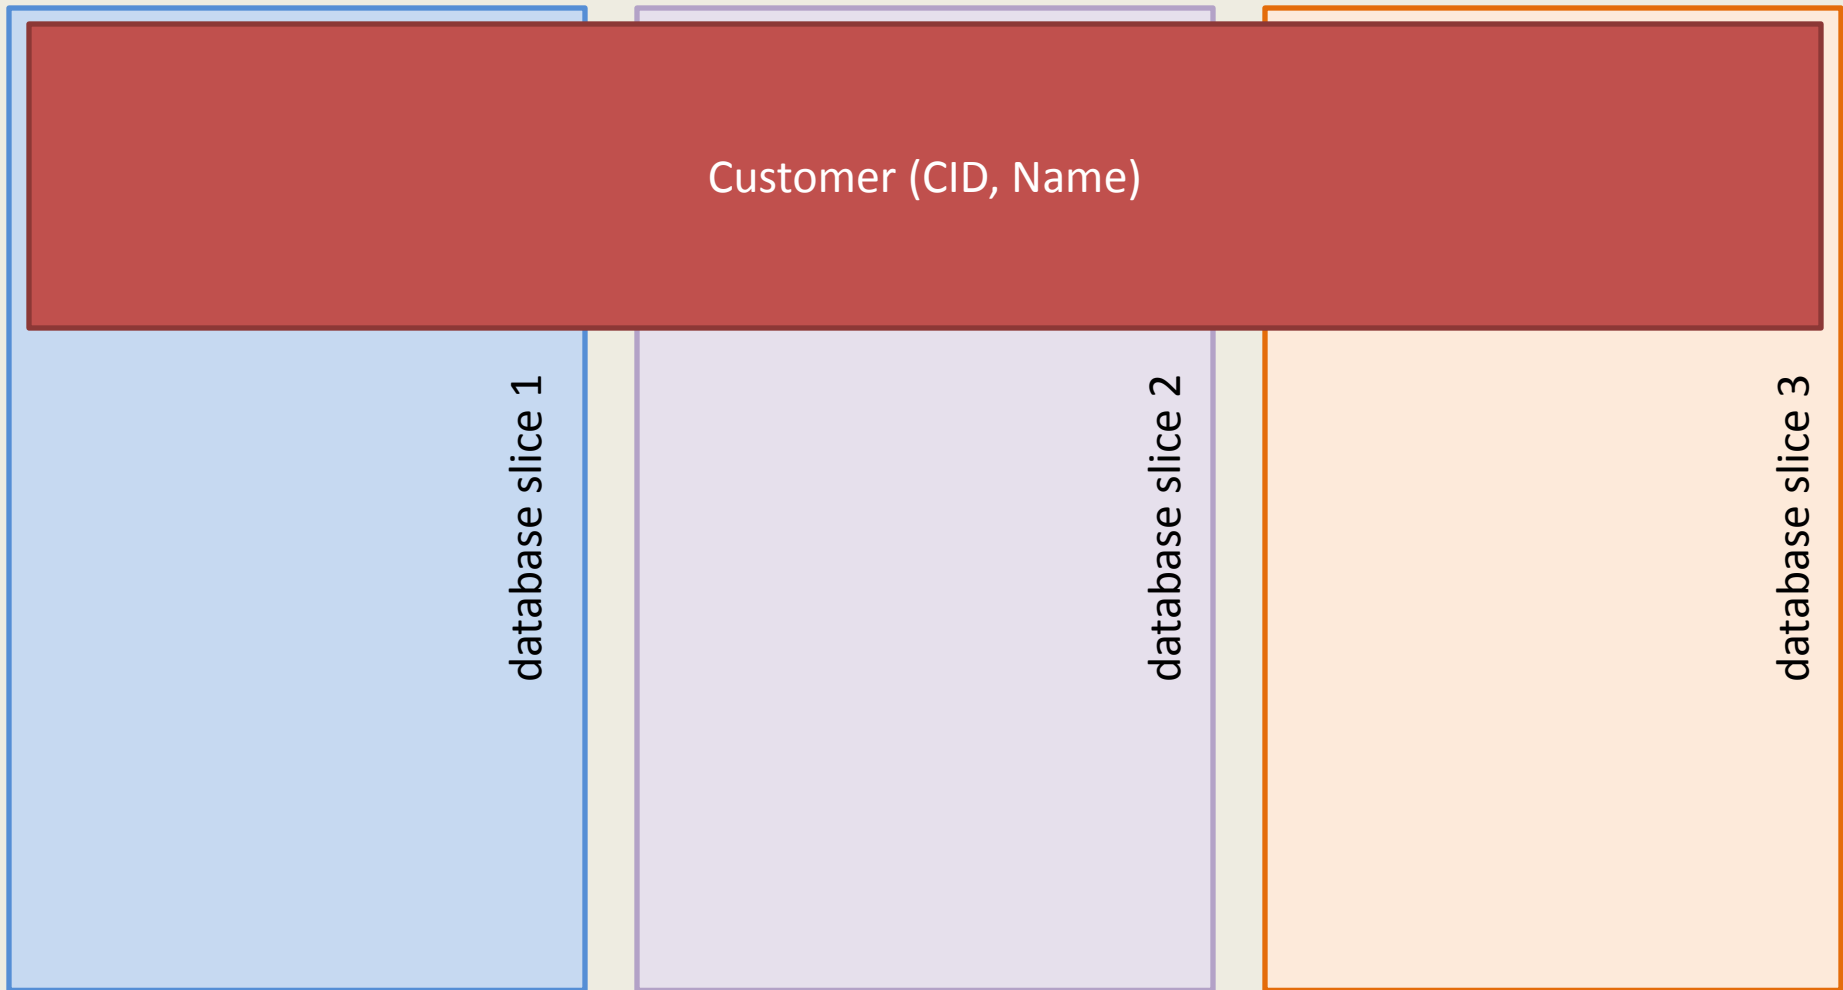


database slice 1

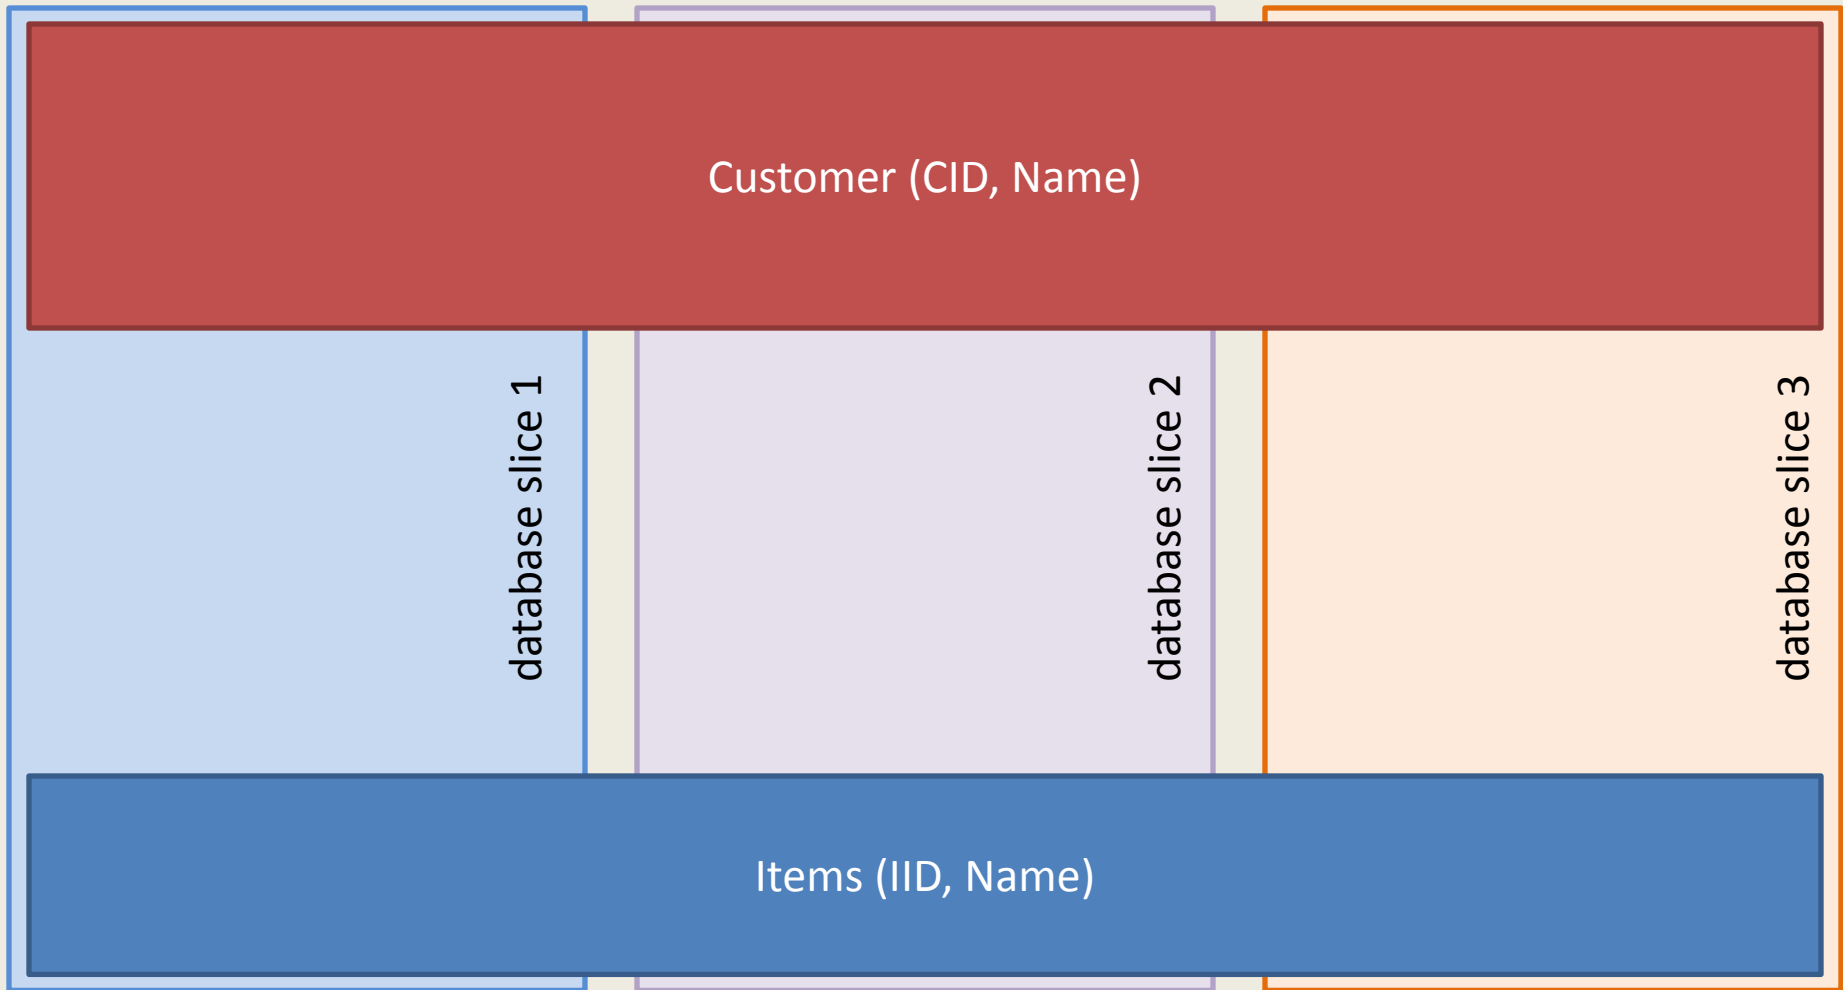
database slice 2

database slice 3

# How it works: The scenario



# How it works: The scenario



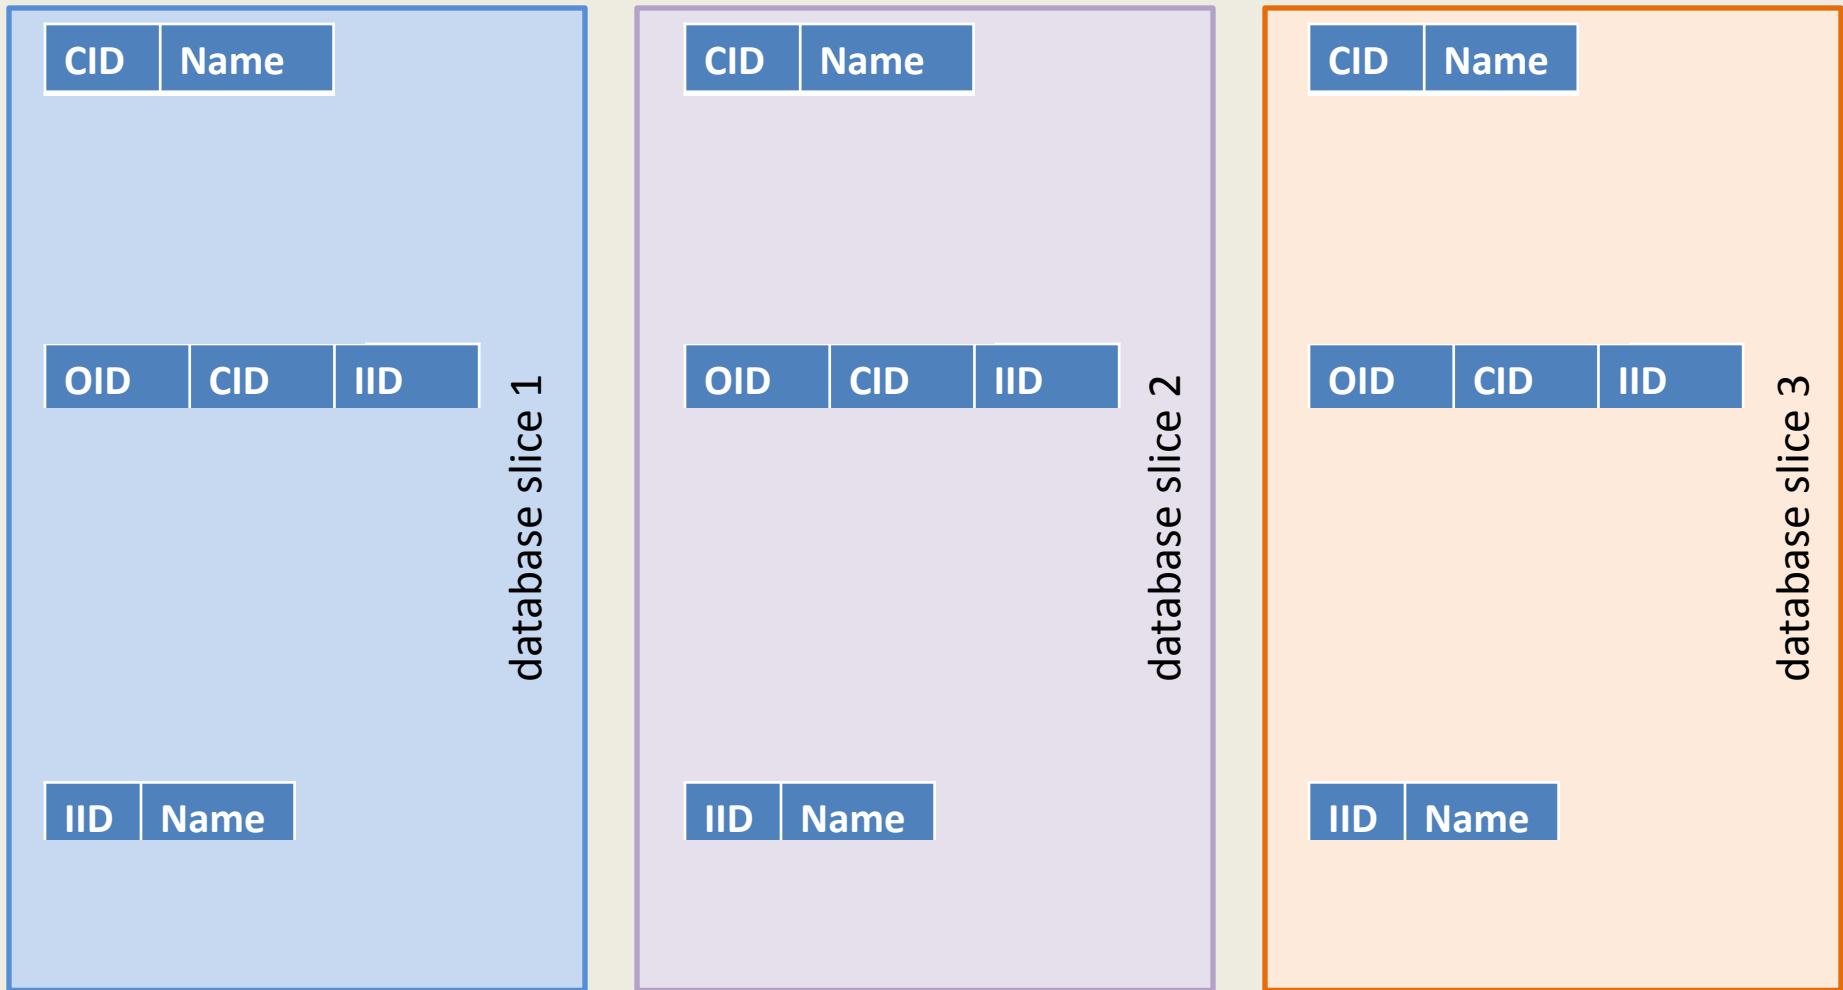
# How it works: The scenario

Customer (CID, Name)

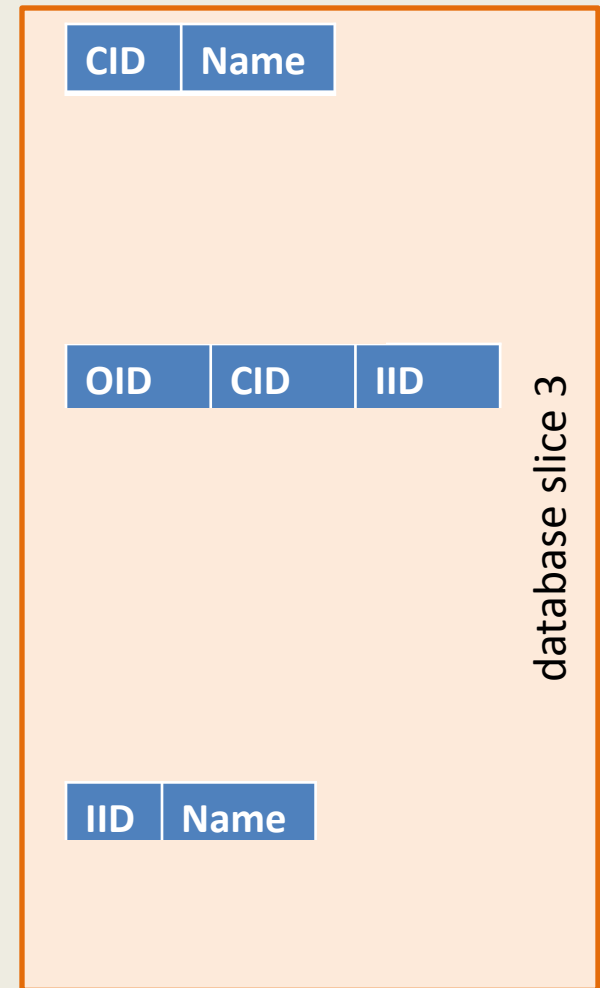
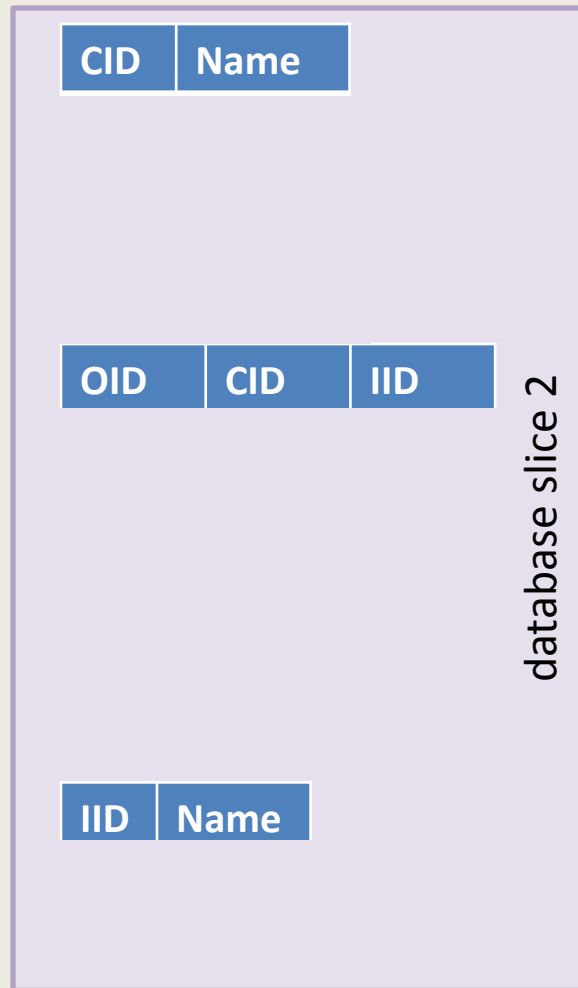
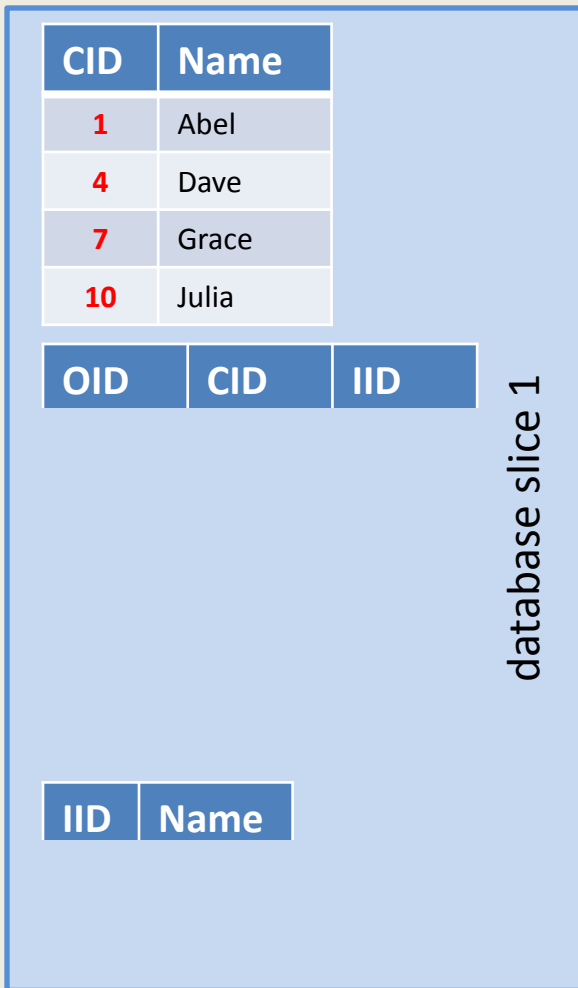
Orders(OID, CID, IID)

Items (IID, Name)

# How it works: The scenario



# How it works: The scenario





# How it works: The scenario

CID	Name
1	Abel
4	Dave
7	Grace
10	Julia

OID	CID	IID
-----	-----	-----

database slice 1

IID	Name
-----	------

CID	Name
2	Baker
5	Ed
8	Henry
11	Kim

OID	CID	IID
-----	-----	-----

database slice 2

IID	Name
-----	------

CID	Name
3	Charles
6	Fred
9	Ivan
12	Louis

OID	CID	IID
-----	-----	-----

database slice 3

IID	Name
-----	------

# How it works: The scenario

CID	Name
1	Abel
4	Dave
7	Grace
10	Julia

OID	CID	IID
-----	-----	-----

database slice 1

IID	Name
101	Cup
104	Plate

CID	Name
2	Baker
5	Ed
8	Henry
11	Kim

OID	CID	IID
-----	-----	-----

database slice 2

IID	Name
102	Bowl
105	Spoon

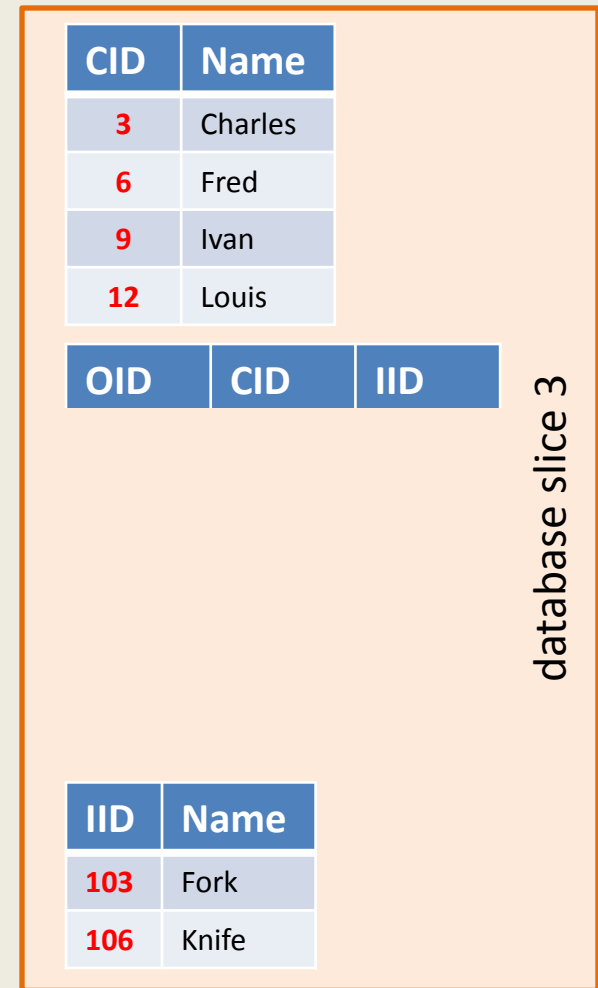
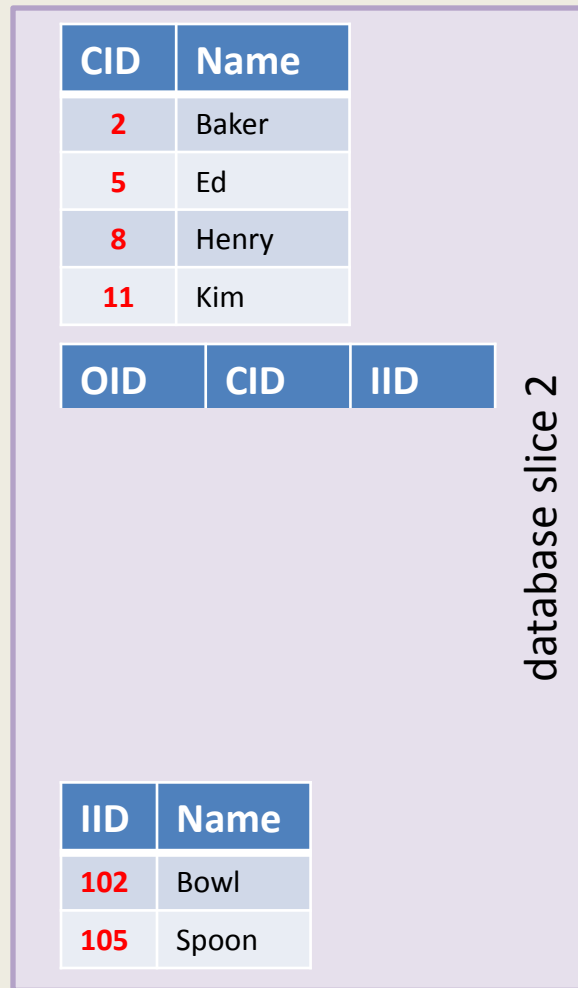
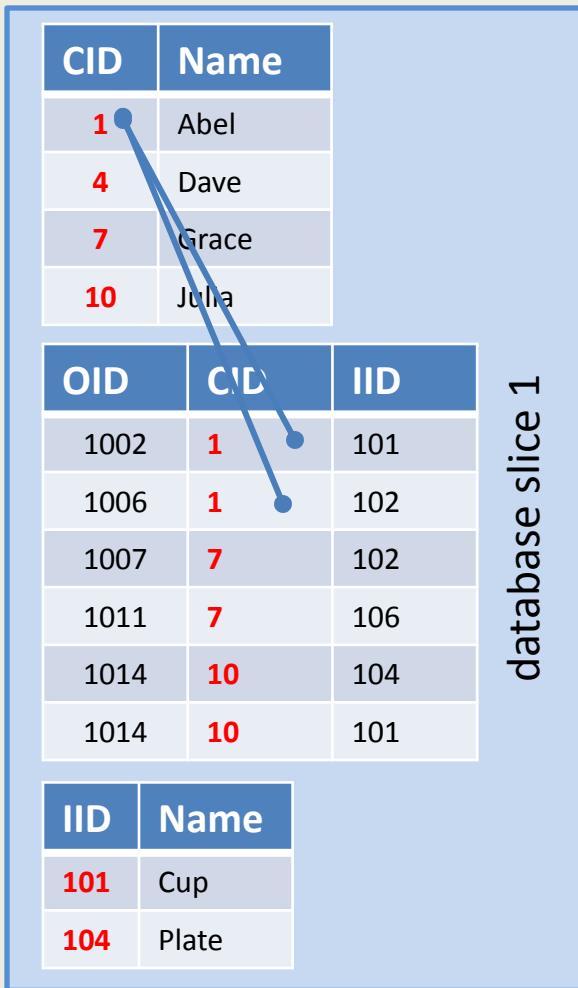
CID	Name
3	Charles
6	Fred
9	Ivan
12	Louis

OID	CID	IID
-----	-----	-----

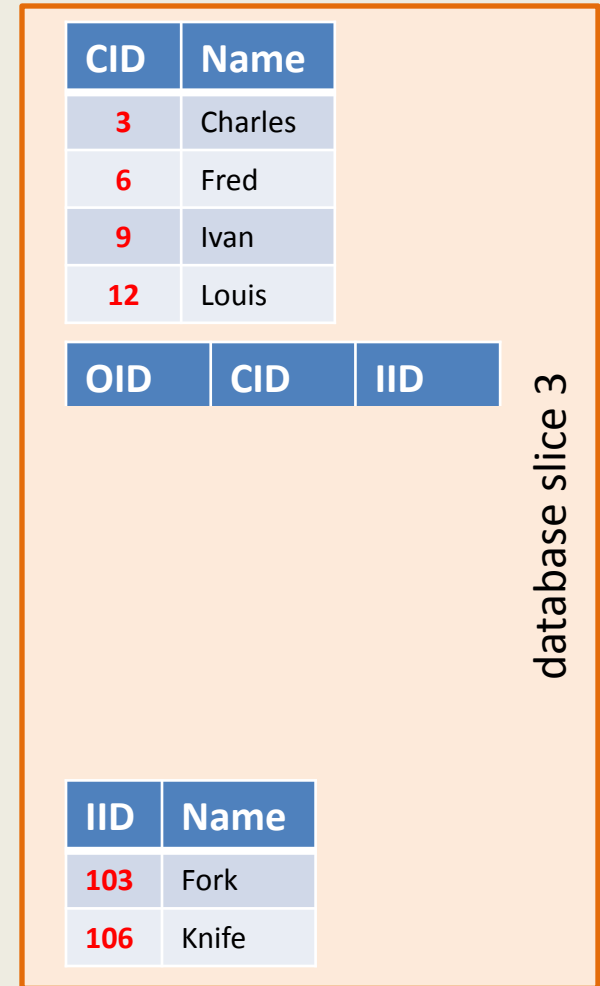
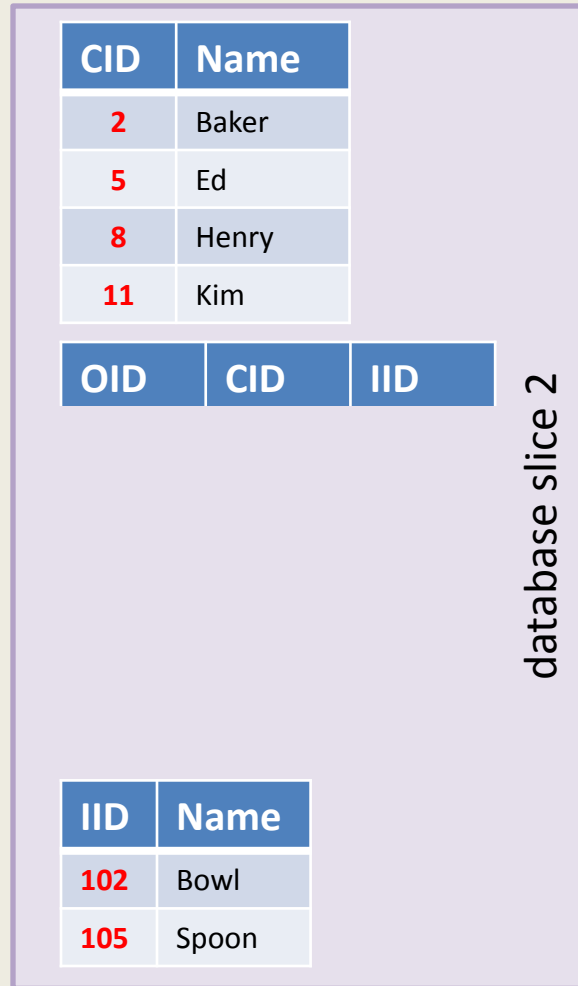

database slice 3

IID	Name
103	Fork
106	Knife

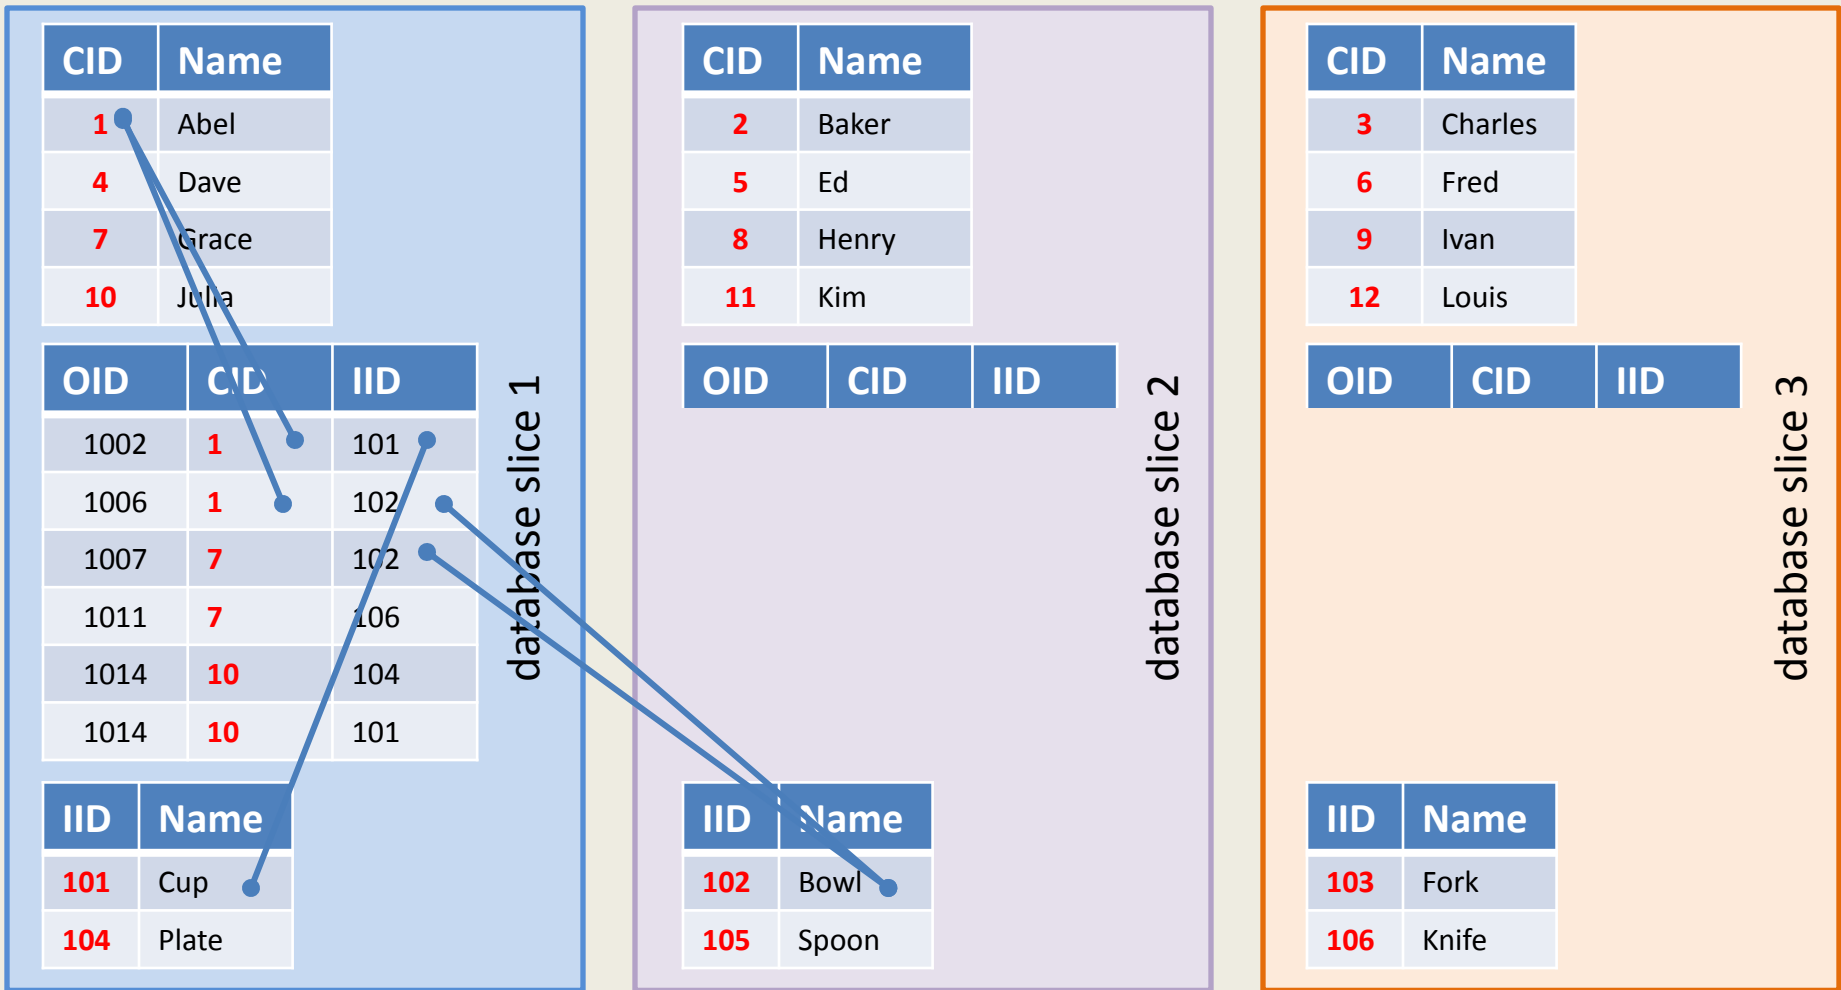
# How it works: The scenario



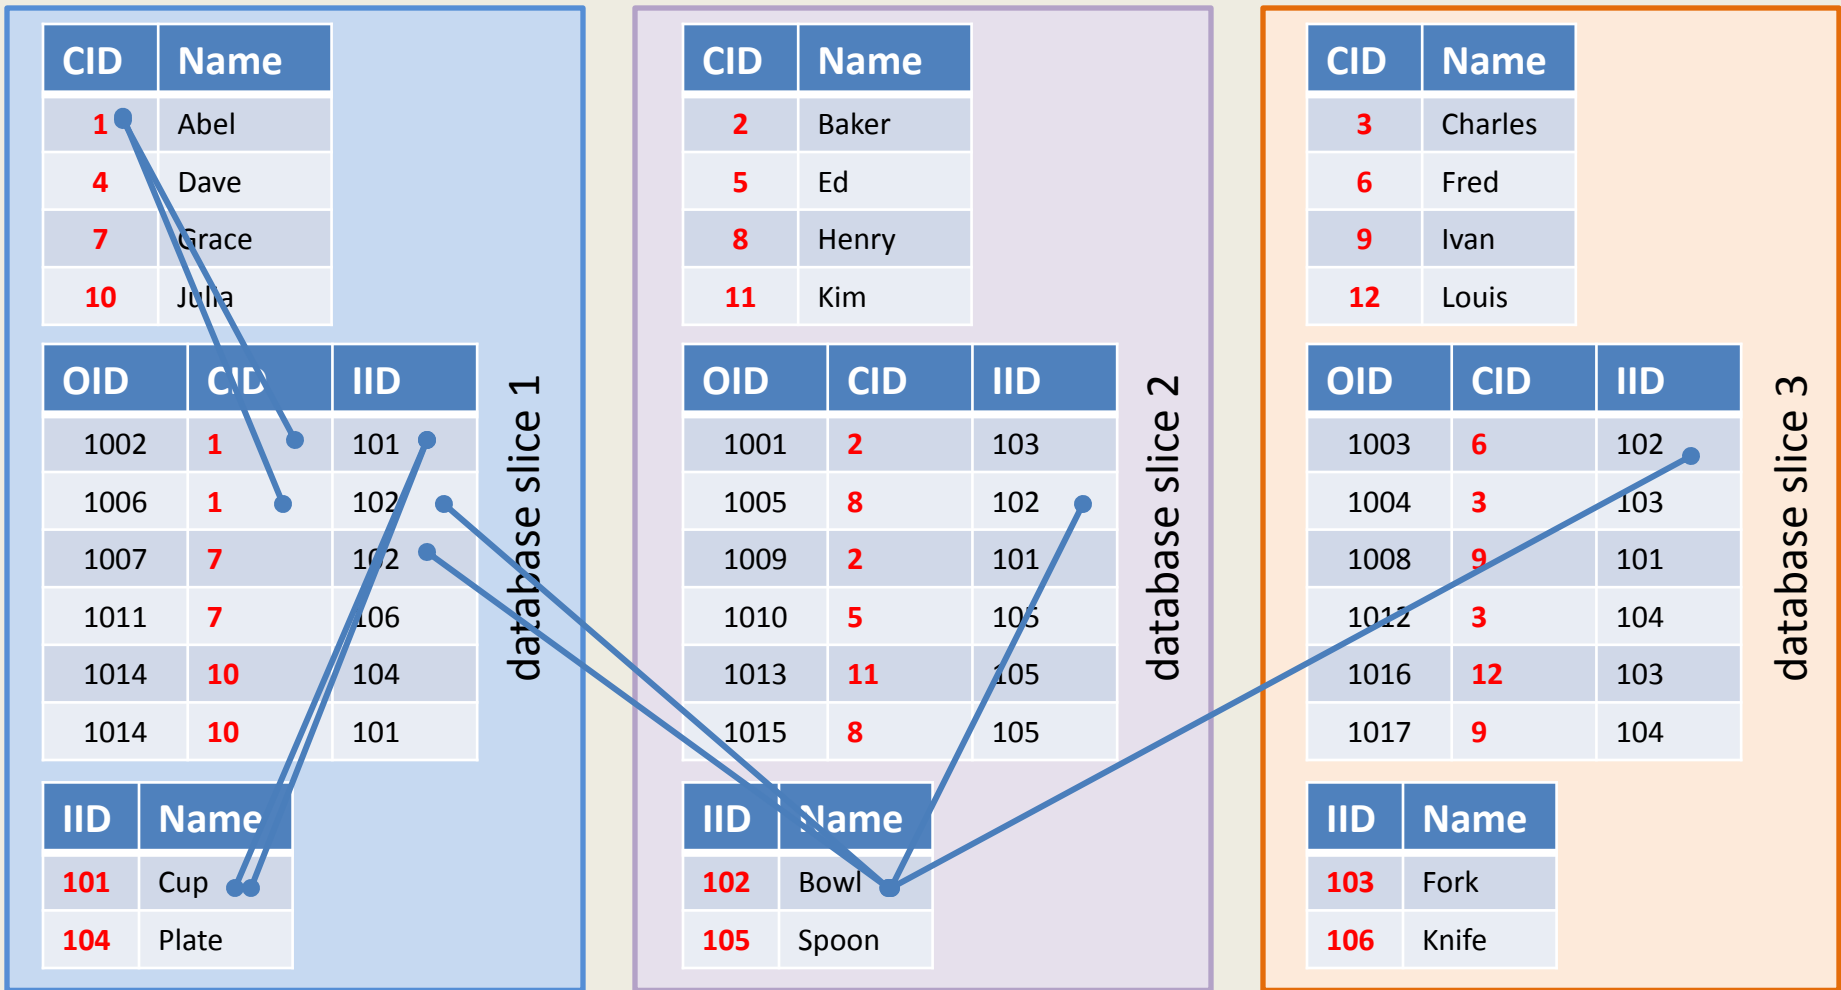
# How it works: The scenario



# How it works: The scenario



# How it works: The scenario



# Who bought a bowl?

database slice 1

CID	Name
1	Abel
4	Dave
7	Grace
10	Julia

OID	CID	IID
1002	1	101
1006	1	102
1007	7	102
1011	7	106
1014	10	104
1014	10	101

IID	Name
101	Cup
104	Plate

database slice 2

CID	Name
2	Baker
5	Ed
8	Henry
11	Kim

OID	CID	IID
1001	2	103
1005	8	102
1009	2	101
1010	5	105
1013	11	105
1015	8	105

IID	Name
102	Bowl
105	Spoon

database slice 3

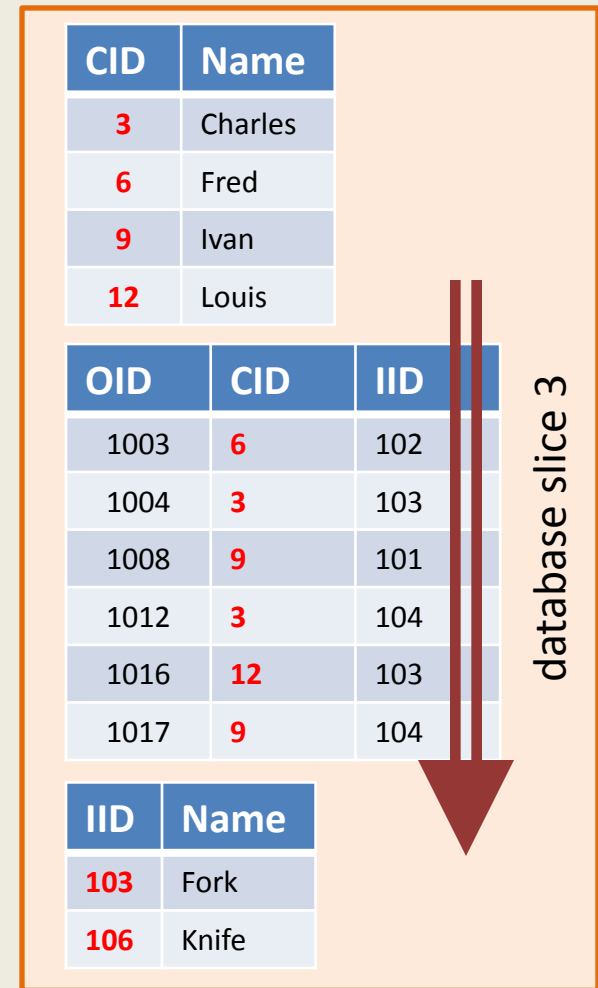


CID	Name
3	Charles
6	Fred
9	Ivan
12	Louis

OID	CID	IID
1003	6	102
1004	3	103
1008	9	101
1012	3	104
1016	12	103
1017	9	104

IID	Name
103	Fork
106	Knife

```
SELECT C.NAME FROM CUSTOMERS C, ORDERS
O, ITEMS I WHERE C.CID = O.CID AND
O.IID = I.IID AND I.NAME = 'Bowl' ;
```

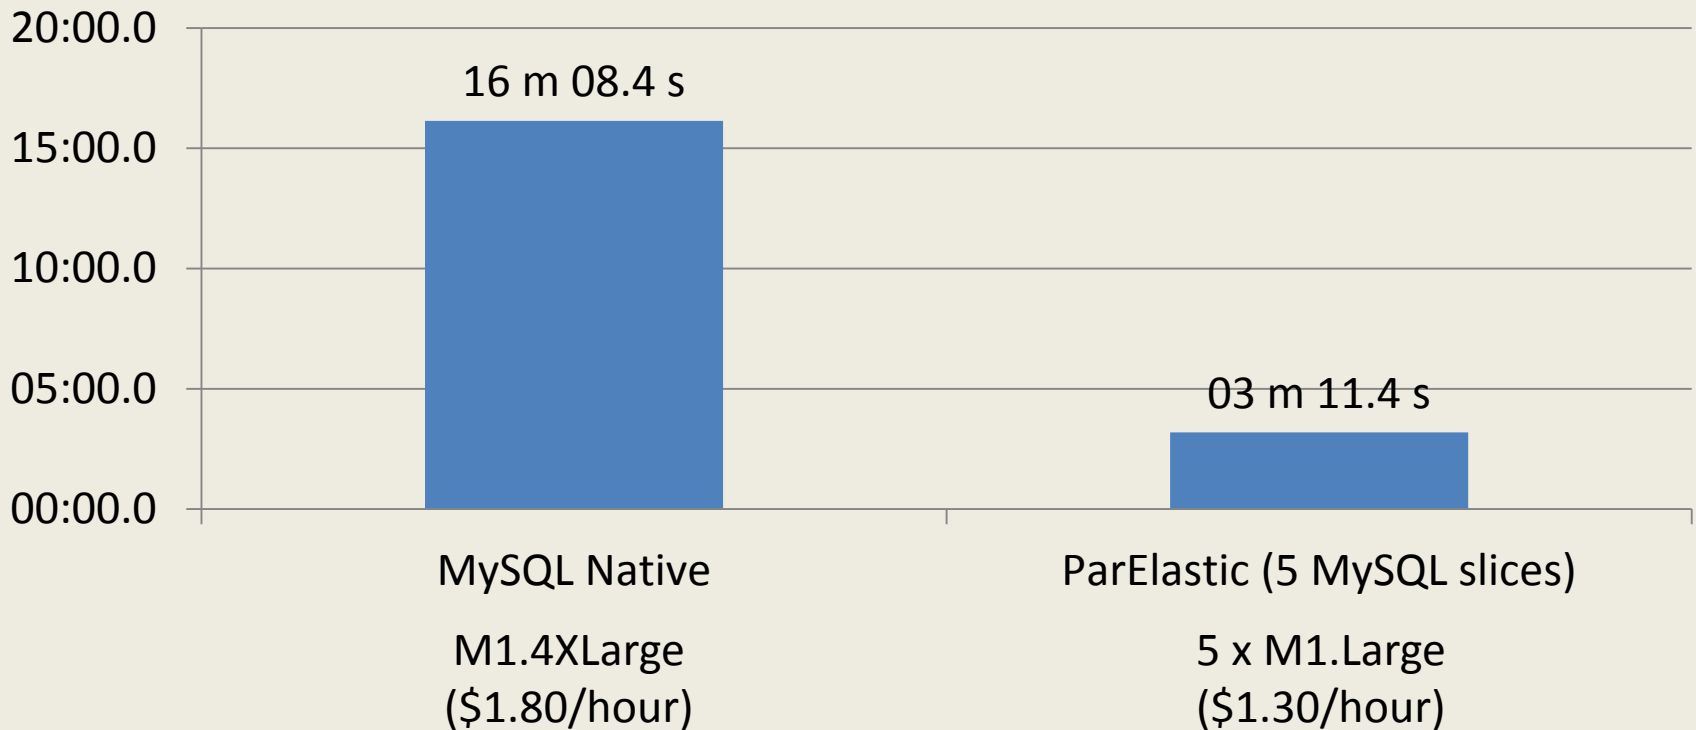
# Scanning in parallel





# Scanning in parallel

## Time to scan a 2b row table



Note: Amazon EC2 pricing for spot instances, December 2012. m1.4xlarge \$1.80/hour, m1.large \$0.26/hour

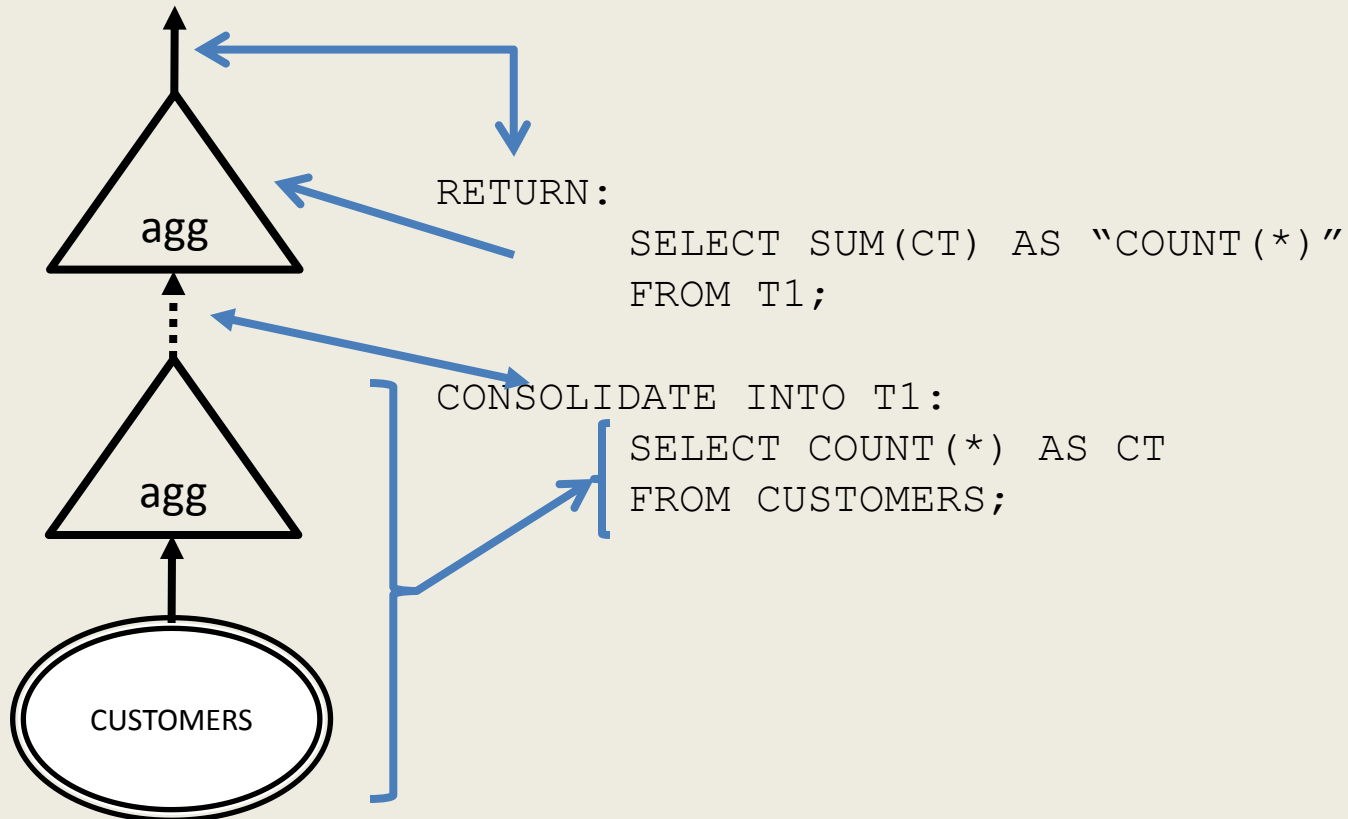
# Counting rows in a table

- Table is distributed on a number of database sites
- I want to count the number of rows in the table

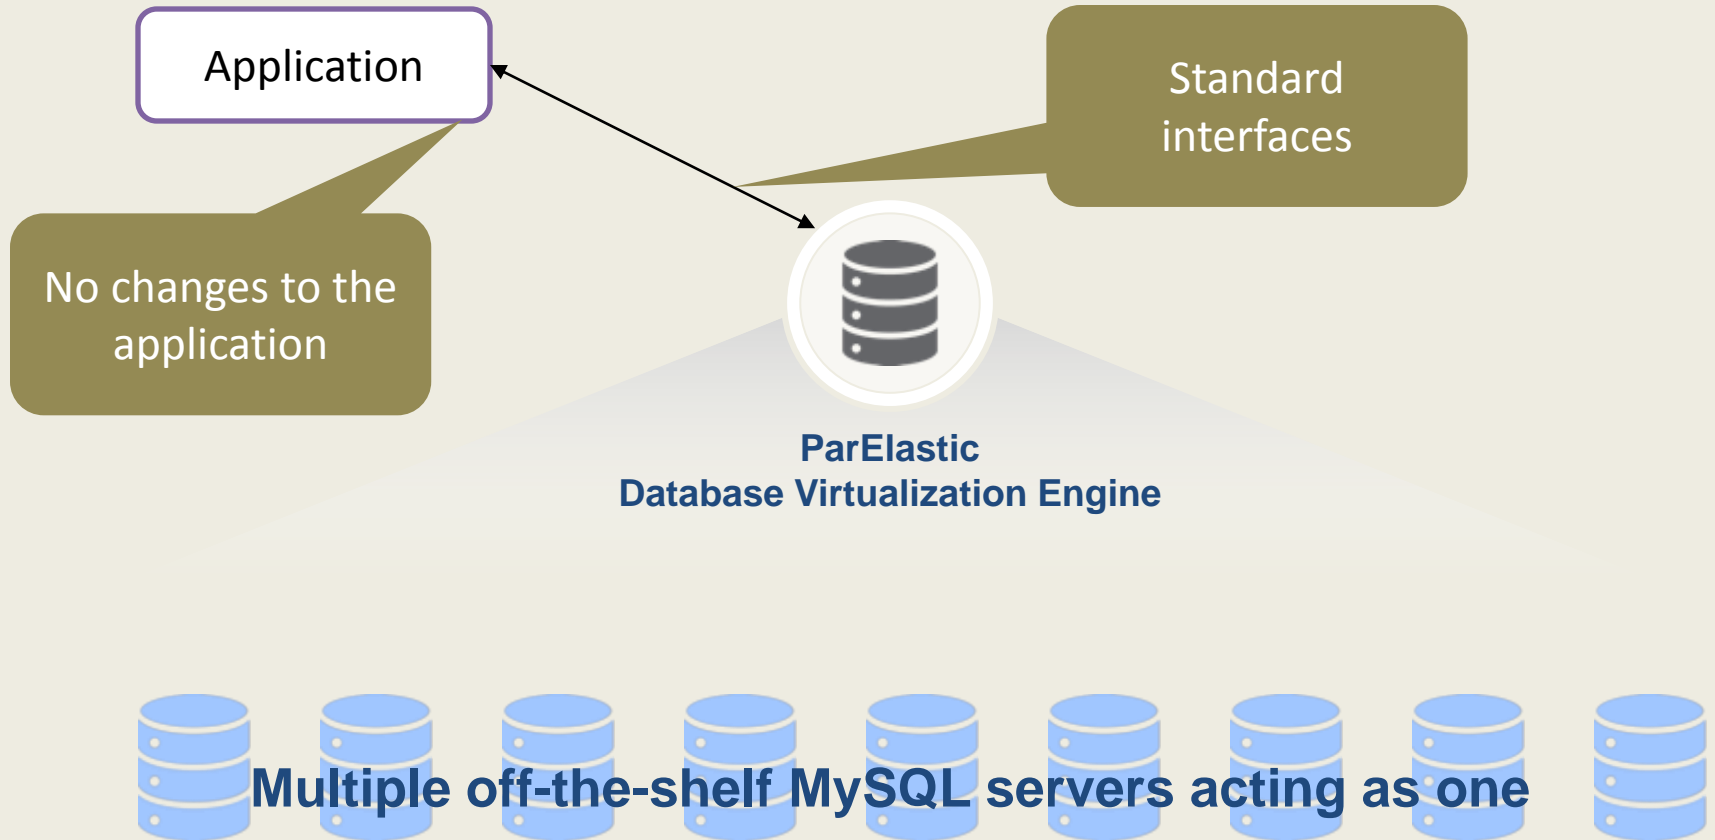
```
SELECT COUNT (*)  
FROM CUSTOMERS ;
```

# Counting rows in a table

```
SELECT COUNT(*) FROM CUSTOMERS;
```

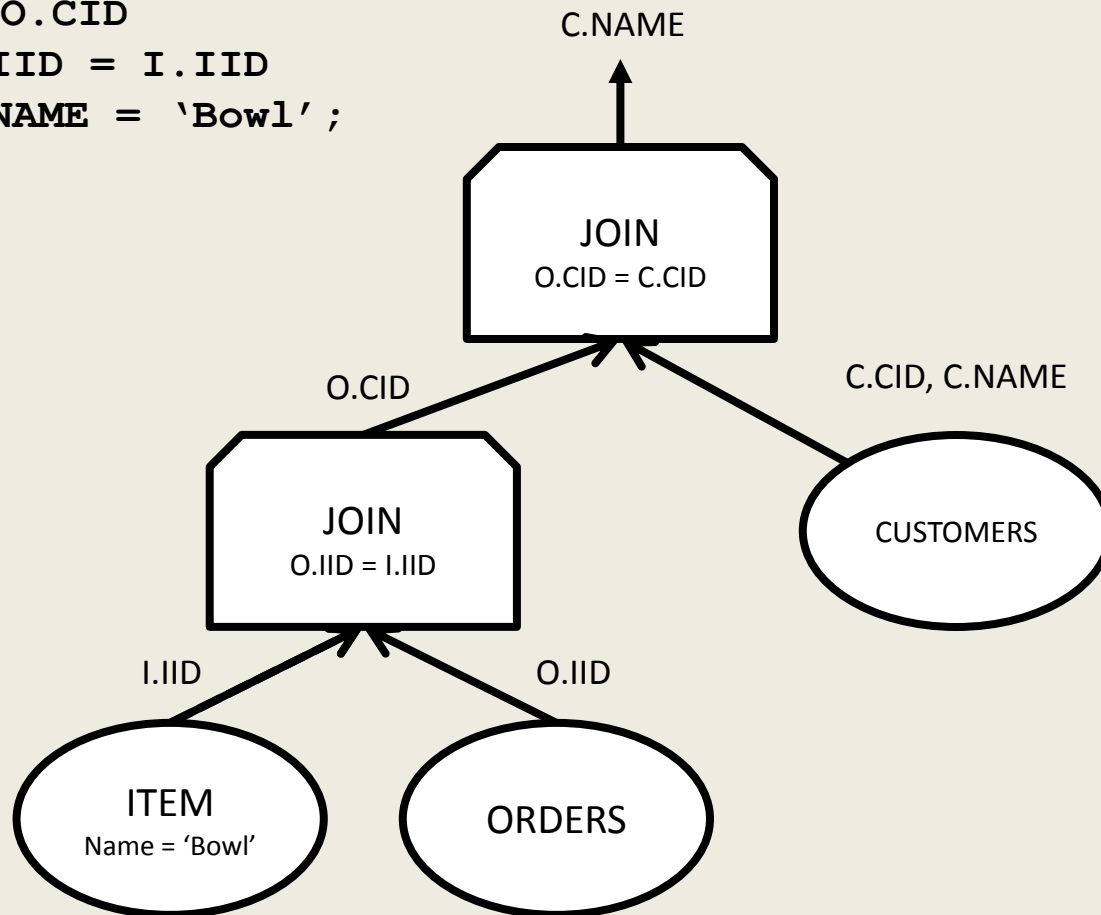


# ParElastic: A parallel database



# A simple join


```
SELECT C.NAME  
FROM CUSTOMERS C, ORDERS O, ITEMS I  
WHERE C.CID = O.CID  
      AND O.IID = I.IID  
      AND I.NAME = 'Bowl' ;
```



# The parallel (distributed) join

```
SELECT C.NAME
FROM CUSTOMERS C, ORDERS O, ITEMS I
WHERE C.CID = O.CID
      AND O.IID = I.IID
      AND I.NAME = 'Bowl' ;
```


How is this  
stream  
distributed?



# The parallel (distributed) join


```
SELECT C.NAME
FROM CUSTOMERS C, ORDERS O, ITEMS I
WHERE C.CID = O.CID
      AND O.IID = I.IID
      AND I.NAME = 'Bowl' ;
```

Distributed  
according to  
table ITEMS (IID)



# The parallel (distributed) join

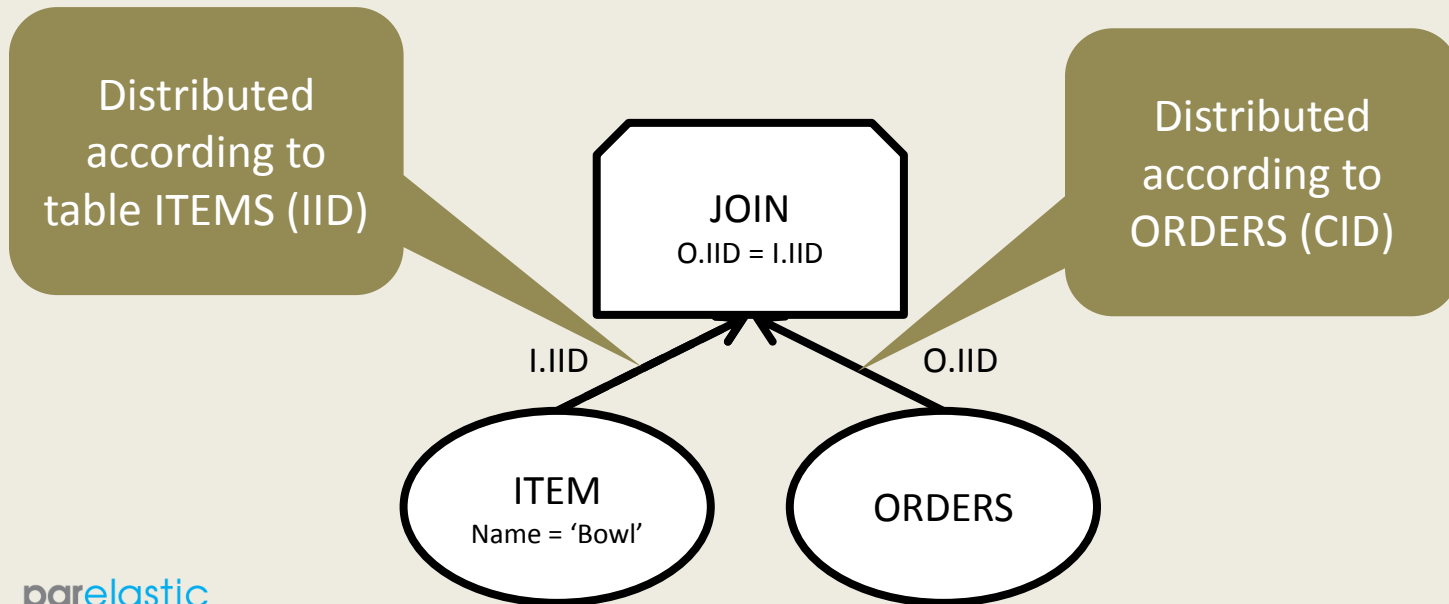
```
SELECT C.NAME
FROM CUSTOMERS C, ORDERS O, ITEMS I
WHERE C.CID = O.CID
      AND O.IID = I.IID
      AND I.NAME = 'Bowl' ;
```






# The parallel (distributed) join

```
SELECT C.NAME
FROM CUSTOMERS C, ORDERS O, ITEMS I
WHERE C.CID = O.CID
      AND O.IID = I.IID
      AND I.NAME = 'Bowl' ;
```



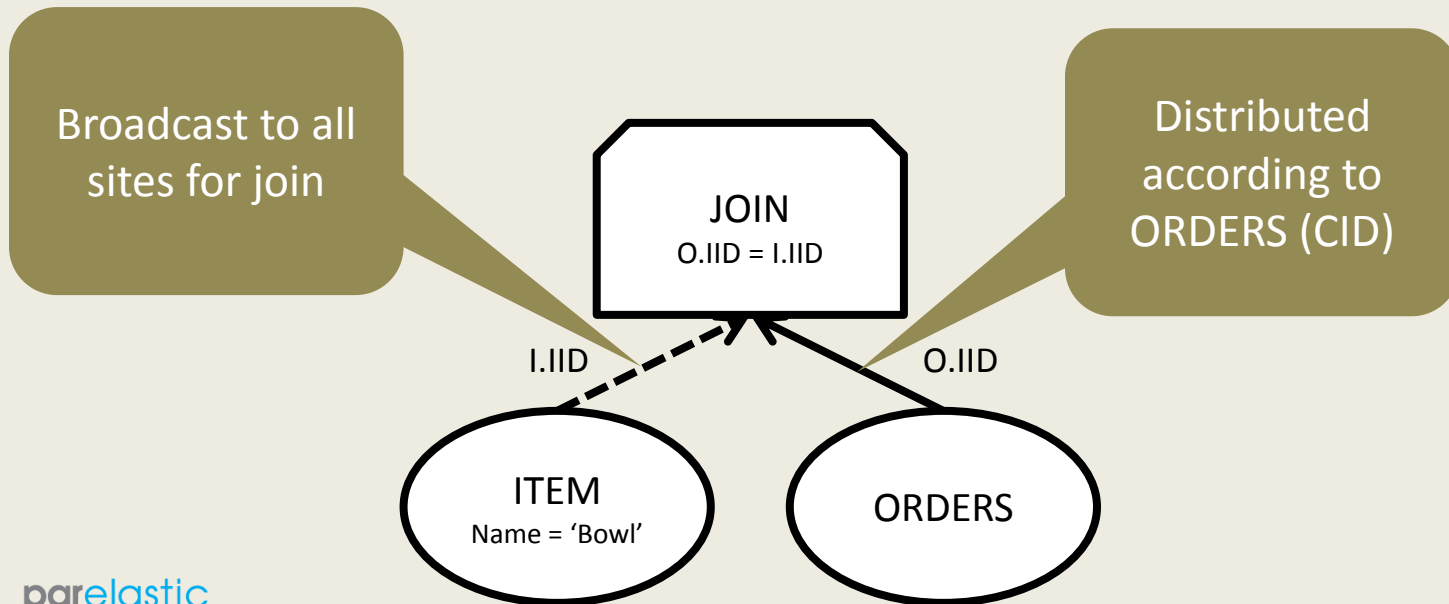
# The parallel (distributed) join

```
SELECT C.NAME
FROM CUSTOMERS C, ORDERS O, ITEMS I
WHERE C.CID = O.CID
      AND O.IID = I.IID
      AND I.NAME = 'Bowl' ;
```



# The parallel (distributed) join

```
SELECT C.NAME  
FROM CUSTOMERS C, ORDERS O, ITEMS I  
WHERE C.CID = O.CID  
      AND O.IID = I.IID  
      AND I.NAME = 'Bowl' ;
```

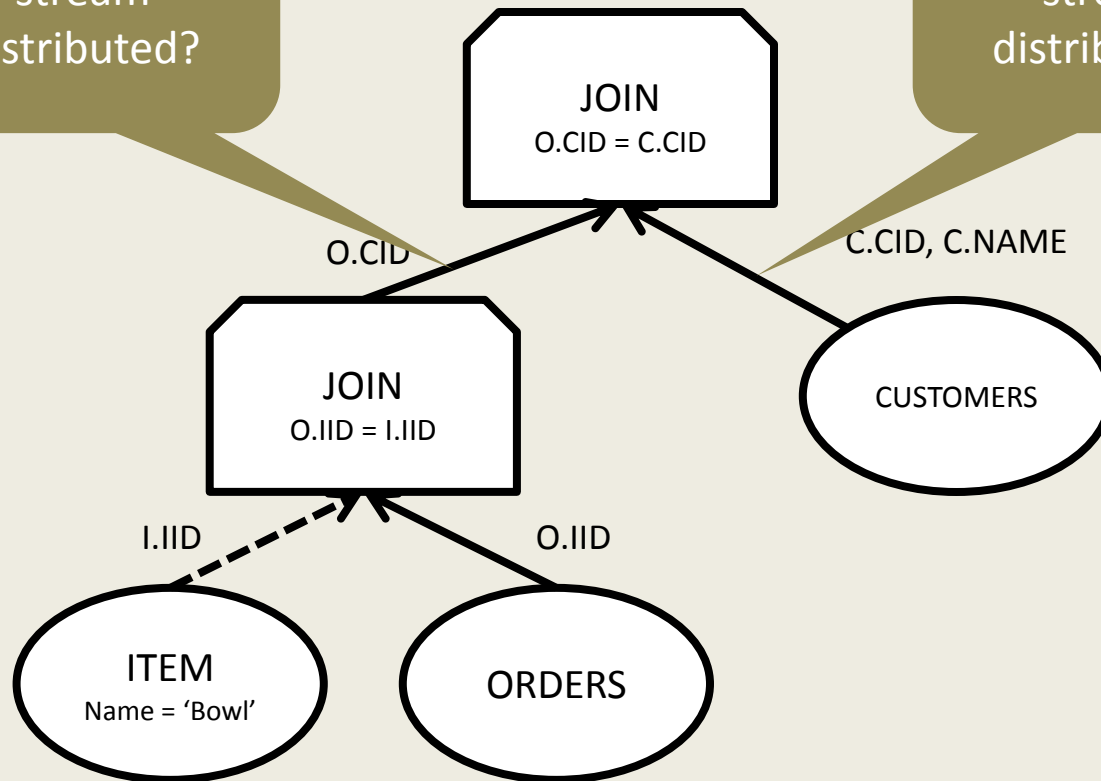


# The parallel (distributed) join

```
SELECT C.NAME  
FROM CUSTOMERS C, ORDERS O, ITEMS I  
WHERE C.CID = O.CID  
AND O.IID = I.IID  
AND I.NAME = 'Bowl';
```


How is this stream distributed?

How is this stream distributed?




# The parallel (distributed) join

```
SELECT C.NAME  
FROM CUSTOMERS C, ORDERS O, ITEMS I  
WHERE C.CID = O.CID  
      AND O.IID = I.IID  
      AND I.NAME = 'Bowl' ;
```



# The parallel (distributed) join

```
SELECT C.NAME  
FROM CUSTOMERS C, ORDERS O, ITEMS I  
WHERE C.CID = O.CID  
      AND O.IID = I.IID  
      AND I.NAME = 'Bowl';
```

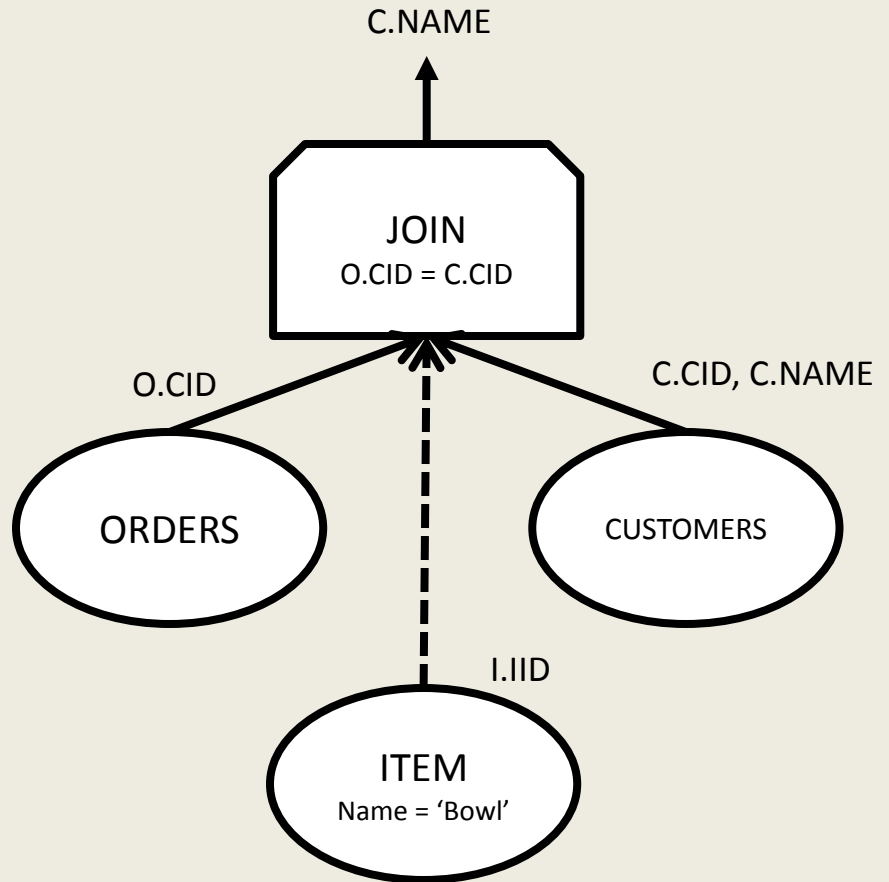


# The parallel (distributed) join

```
SELECT C.NAME
FROM CUSTOMERS C, ORDERS O, ITEMS I
WHERE C.CID = O.CID
      AND O.IID = I.IID
      AND I.NAME = 'Bowl';
```

```
BROADCAST TO T1:
  SELECT IID FROM ITEMS
  WHERE NAME = 'Bowl';
```

```
RETURN:
  SELECT C.NAME
  FROM CUSTOMERS C,
  ORDERS O, T1
  WHERE C.CID = O.CID
        AND O.IID = T1.IID;
```

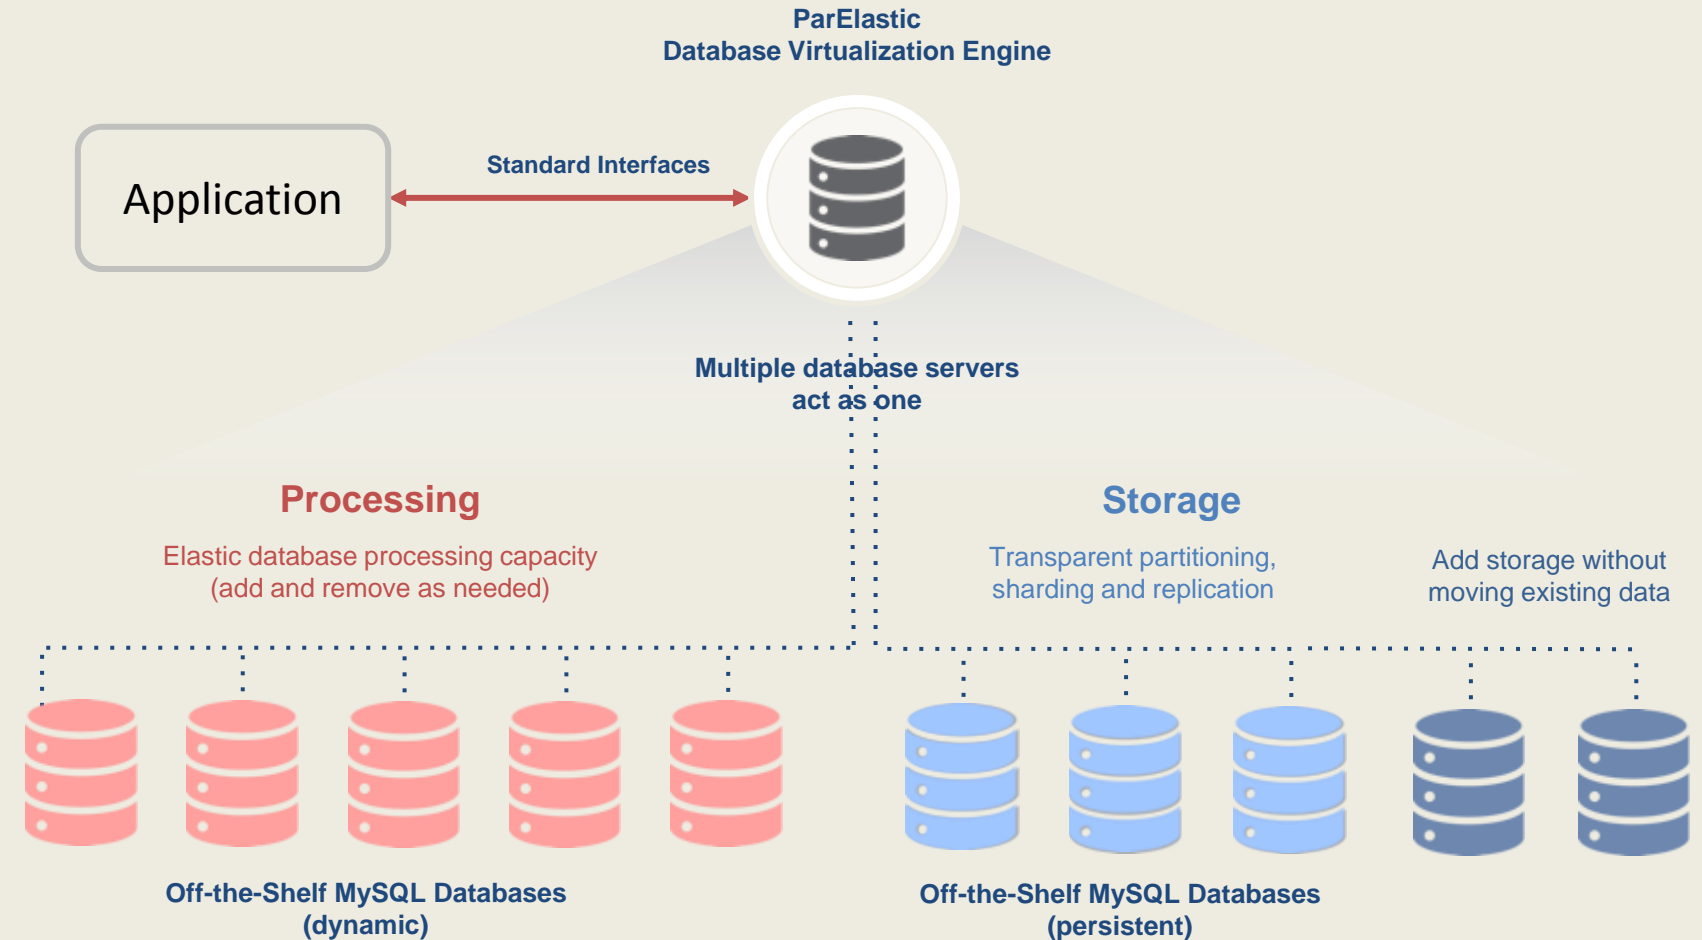


# Parallel Databases vs. Sharding

- Parallel Database
  - Database architecture
  - Application is data location agnostic
  - Application perceives a single database
  - Requires **no** application rewrites
- Application **is not** constrained by parallel database architecture
- A parallel database handles any schema
- Sharding
  - Application architecture
  - Application is data location aware
  - Application perceives a collection of databases
  - Requires **application** rewrites
- Application **is** constrained to the limitations of the sharding architecture
- Not all schemas are shard'able



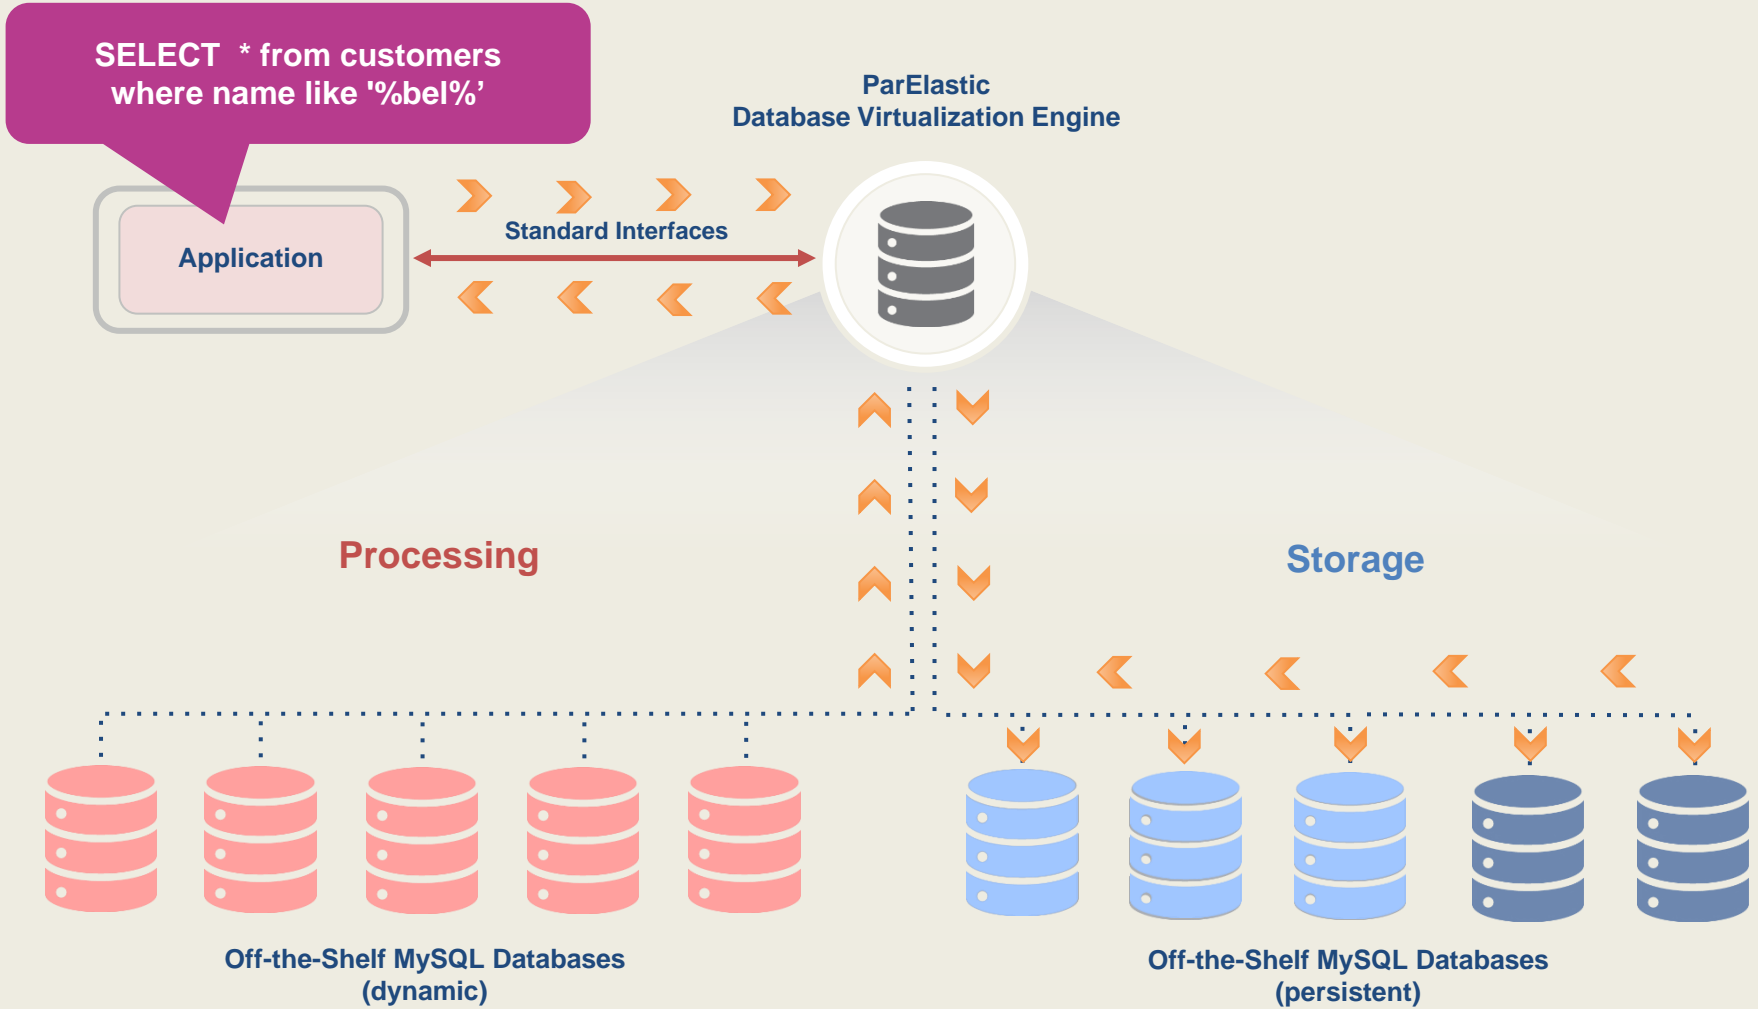
# ParElastic Architecture



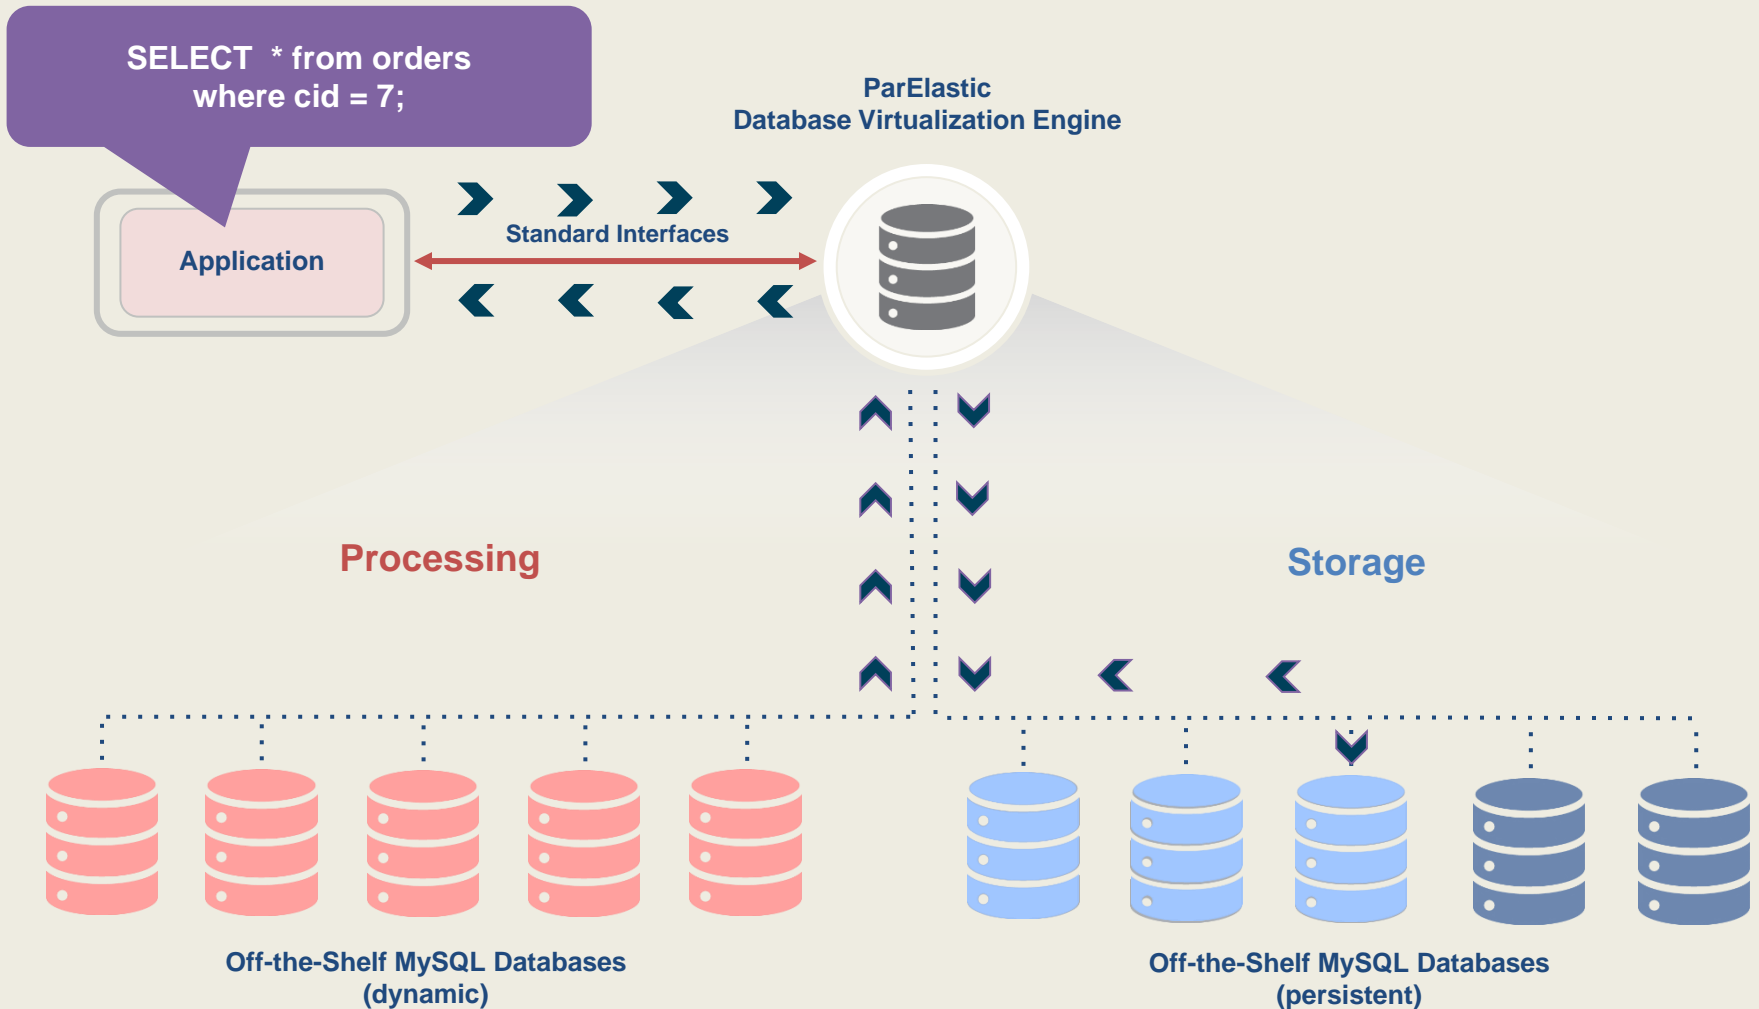
# ParElastic data distribution

- Data distribution based on role in the schema
- Table data can be distributed
  - Randomly
  - Deterministically based on some attributes
  - Broadcast, XA for transaction consistency
- Other distribution methods for multi-tenancy
  - Later in this presentation

# Simple Select



# Query by Distribution Key



# Local Join

select customers.name, orders.iid from customers,  
orders where customers.cid = orders.cid;

ParElastic  
Database Virtualization Engine

Application

Standard Interfaces

Processing

Storage

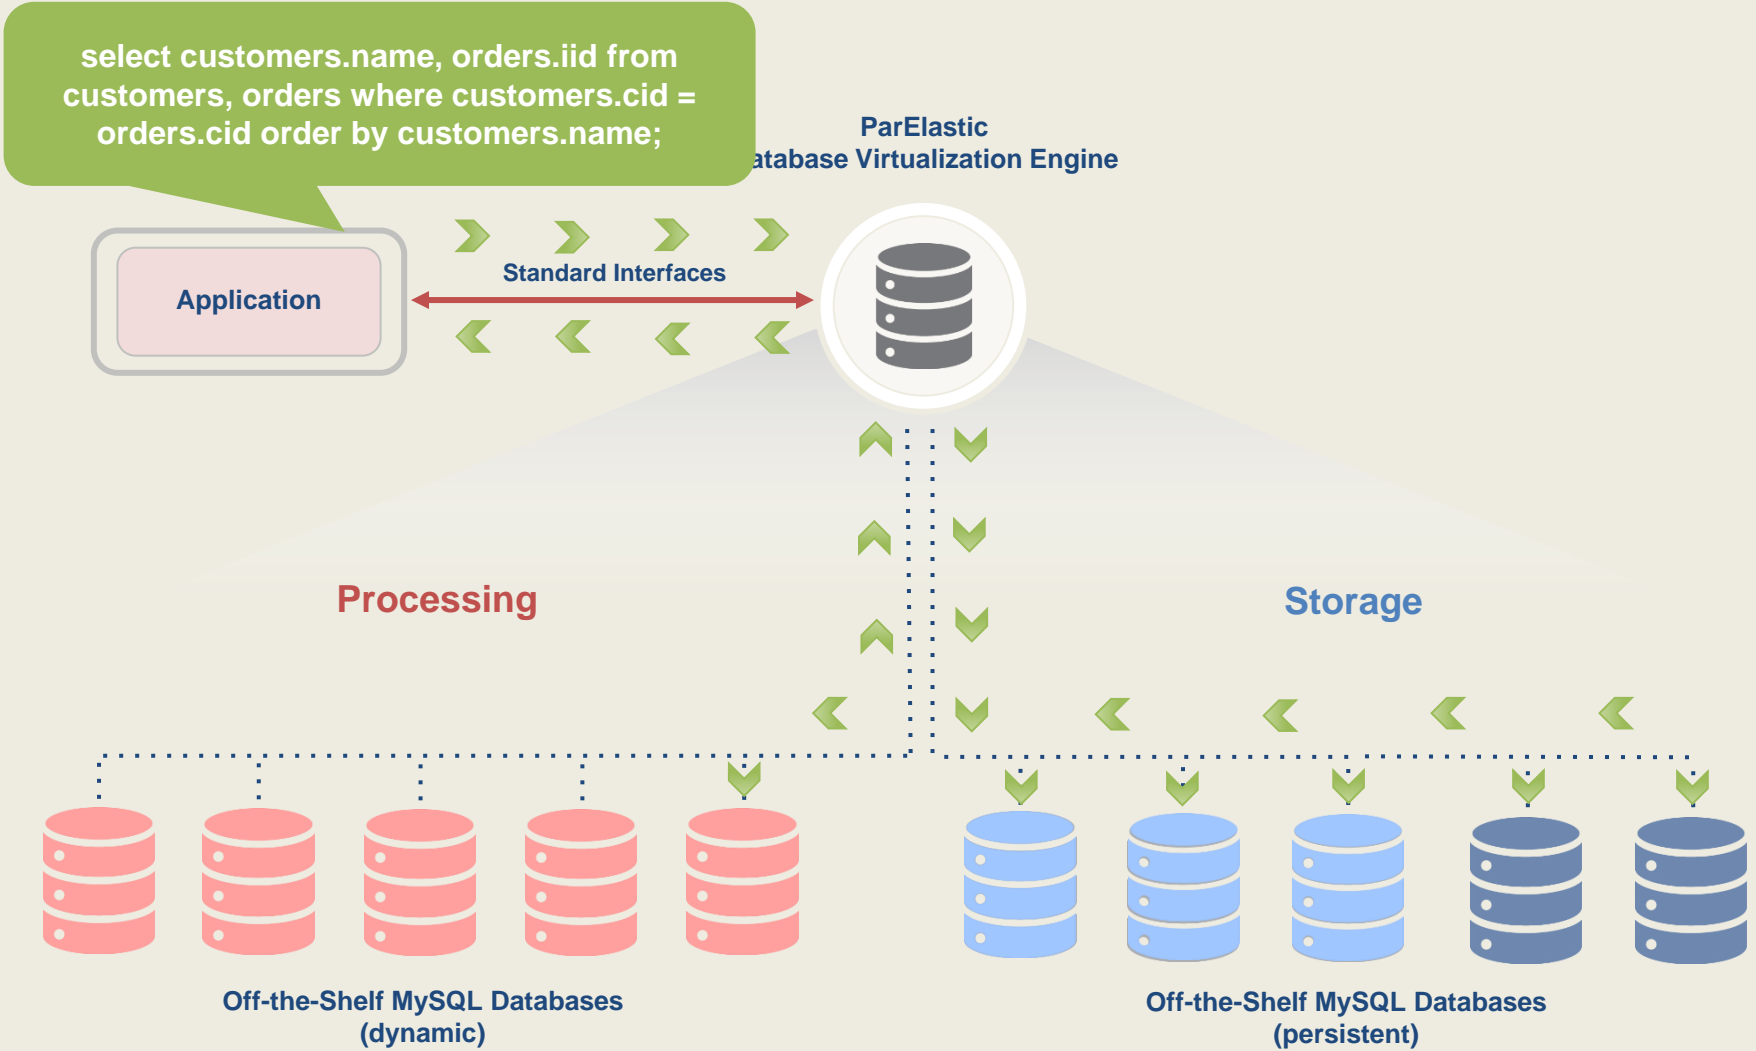


Off-the-Shelf MySQL Databases  
(dynamic)



Off-the-Shelf MySQL Databases  
(persistent)

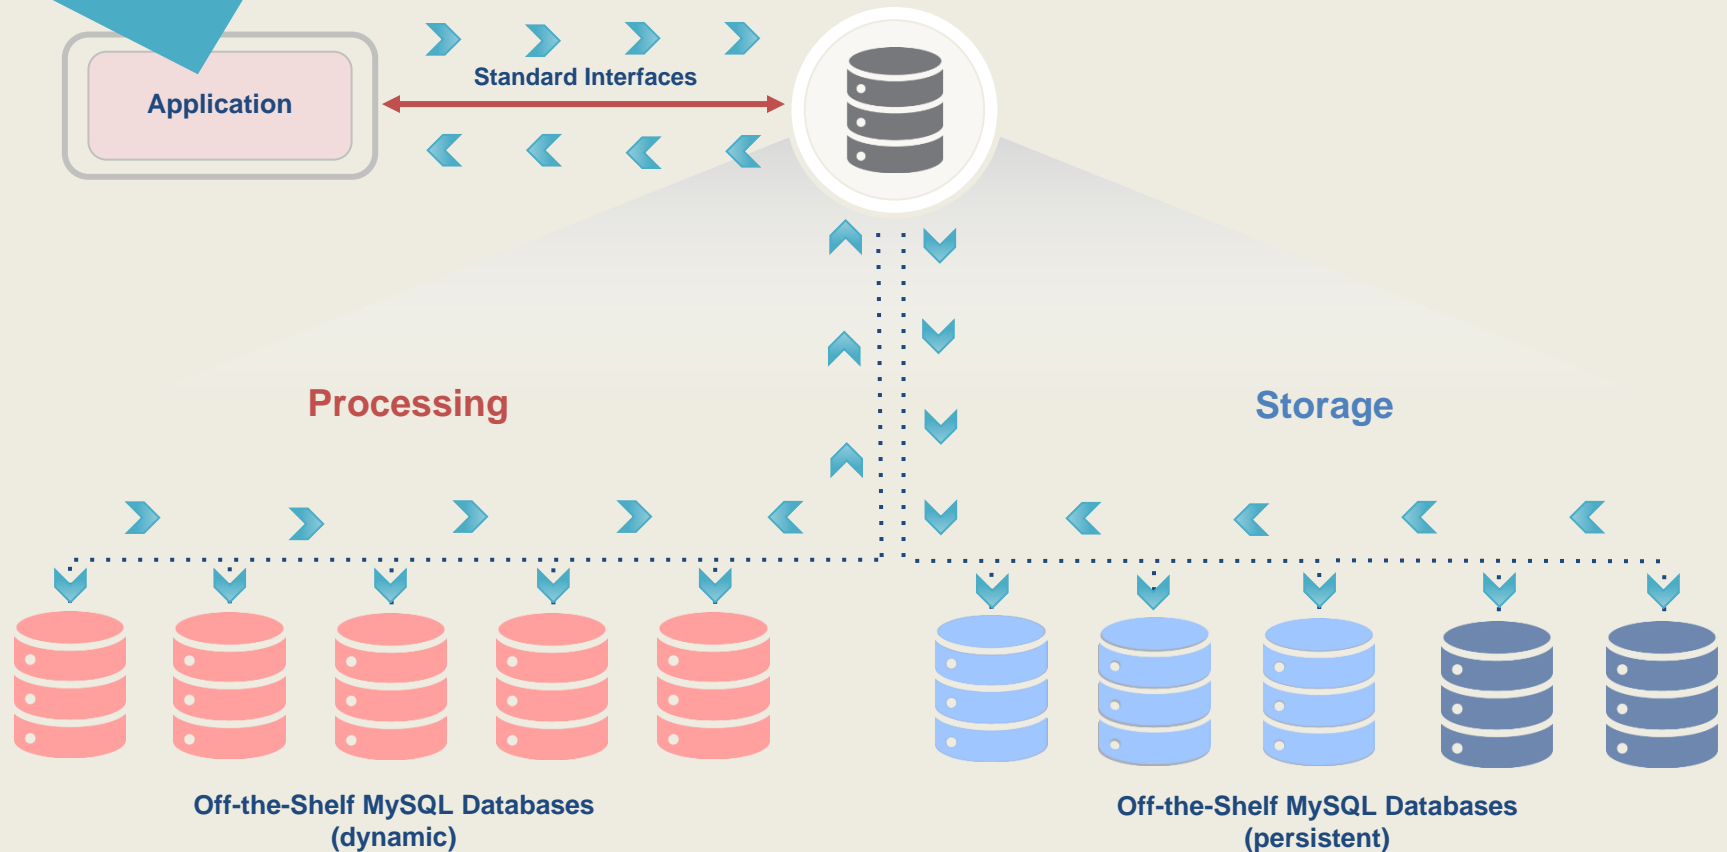
# Ordered Result



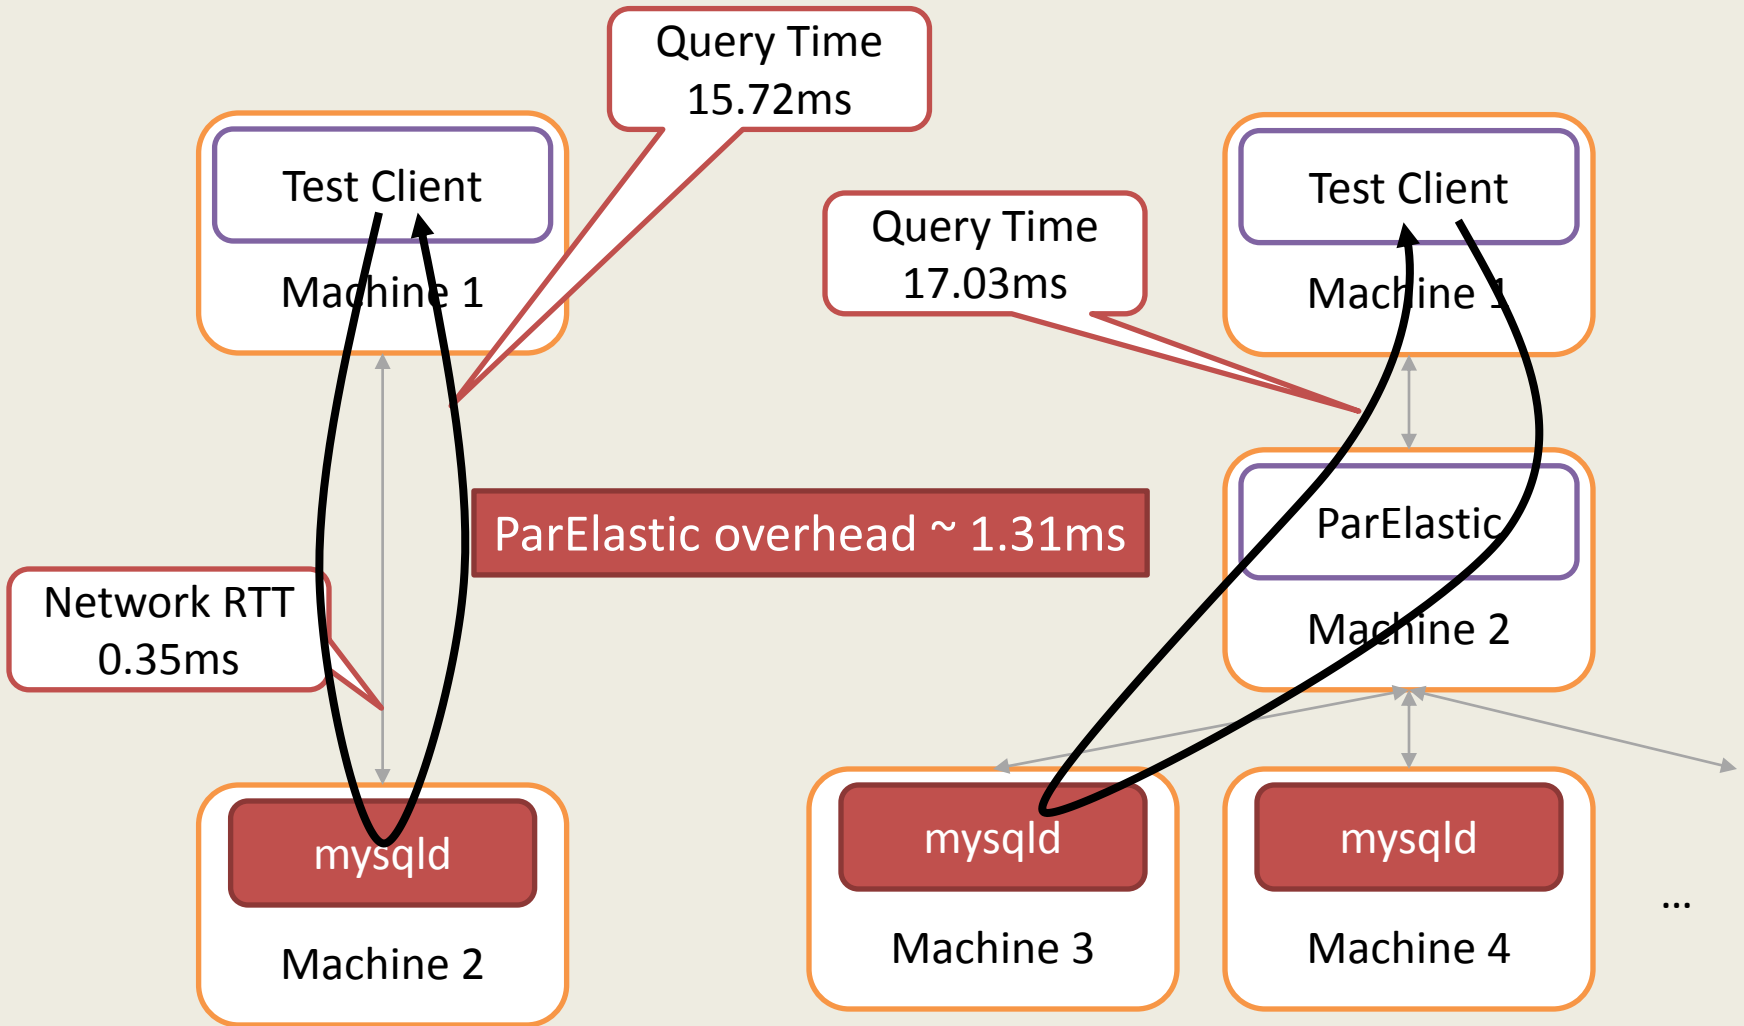
# A Cross Node Query

```
select items.name, count(*) from customers,  
orders, items where customers.cid = orders.cid and  
orders.iid = items.iid and customers.name like  
'%e%' group by items.name;
```

ParElastic  
Database Virtualization Engine




# What's the overhead?

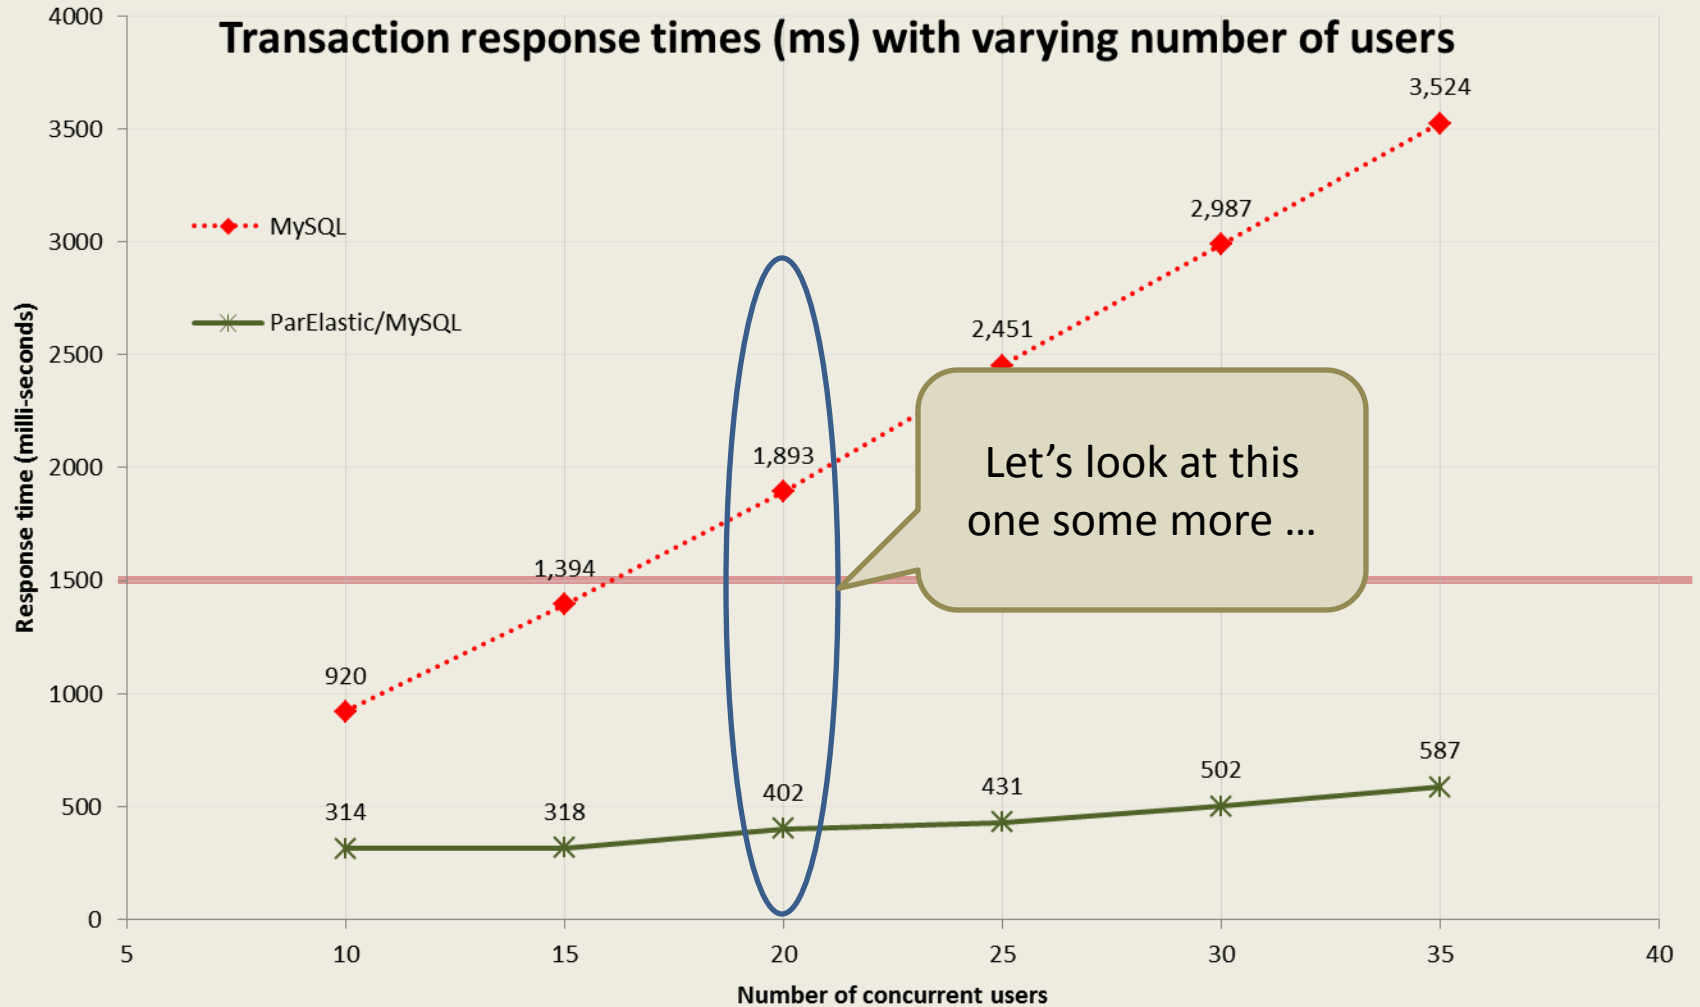




# The power of parallelism



# The power of parallelism



# SaaS Applications and Multi-tenancy

# SaaS application multi-tenancy

## Simple MT

- Each tenant gets their own database & tables
- Application switches to appropriate database

## In application MT

- Application code consolidates tenants into a single database

## ParElastic MT

- Each tenant believes it has its own database and tables
- ParElastic virtualizes and consolidates databases

### PROS

- Simplicity
- Isolation of data

- Improves utilization, reduces # databases

- Simplicity
- Isolation of data
- Rolling upgrades
- Customizations supported

### CONS

- Many databases causes performance issues
- Bad neighbor effect

- Code complexity
- Bad neighbor effect
- Lock-Step Upgrades
- Restricts customization

# ParElastic Adaptive Multi-Tenancy

- Drupal: A SaaS Application
- Tens of thousands of Drupal sites
- Each site has its own database
  - With its own tables and data
  - Schema customizations / modules
  - Thousands of MySQL servers
- Thousands of Application Servers
  - Rolling software upgrades
  - Some with schema changes



# Database Scalability at Acquia

Before ParElastic

ZERO lines of code change in Drupal!




After ParElastic



parelastic  
Flex Your Database™

<http://www.parelastic.com/>

# Impact of large database count



# ParElastic data distribution

- Data distribution based on role in the schema
- Table data can be distributed
  - Randomly
  - Deterministically based on some attributes
  - Broadcast, XA for transaction consistency
- Multi-tenant systems use all of the above plus
  - Tenant based distribution semantics



# And now, a demonstration

# The demonstration

- It's hard to demonstrate plumbing
  - Much easier to demonstrate a faucet!
- We demonstrate a Drupal site powered by ParElastic
- And to simulate traffic
  - We ingest data with specific #hashtags from Twitter
  - So go ahead and tweet with #mysql or #parelastic and you will generate traffic for the demo

# So what's ParElastic?

- A local start-up
  - Based in Waltham, R&D office in Toronto
  - We're hiring!
- ParElastic is Database Virtualization Software
  - Make groups of MySQL databases act as one
  - The cardinality of the group can change
    - Based on data size
    - Based on workload

# I want it, where can I get it?

- We're in closed beta
  - Working with some very exciting early adopters
  - Happy to work with you as well!

# Contact information

- ParElastic
  - Web: <http://www.parelastic.com/>
  - Email: [info@parelastic.com](mailto:info@parelastic.com)
  - Twitter: @parelastic
  
- Me
  - Email: [amrith@parelastic.com](mailto:amrith@parelastic.com)
  - Twitter: @amrithkumar