



Scaling Hibernate Applications with Postgres

Bruce Momjian

Jim Mlodgenski

JbossWord 2009

www.enterprisedb.com

**JBoss
WORLD**
CHICAGO 2009

- To show some the typical scaling techniques used by Postgres
- To demonstrate how Hibernate can leverage those techniques
- To highlight some issues using Hibernate with Postgres and how to overcome them

Bruce Momjian

Sr Database Architect
EnterpriseDB Corporation

Jim Mlodgenski

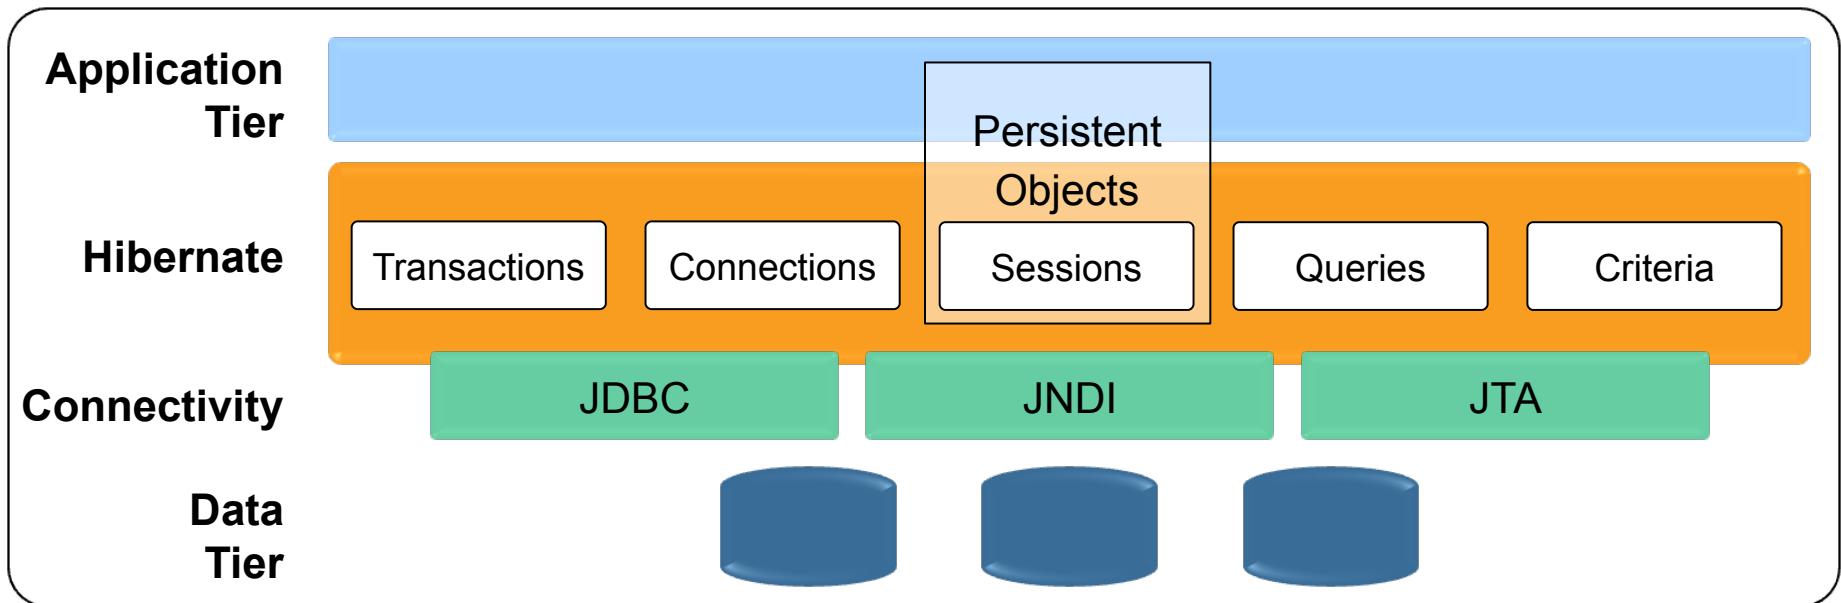
Chief Architect
EnterpriseDB Corporation

About EnterpriseDB

- Award-winning open source database company
- World's largest concentration of PostgreSQL expertise
- Headquartered in Westford, MA, USA
- 300+ customers including Sony, FTD, hi5 Networks, NetApp, FAA, LexisNexis, St. Jude Children's Hospital
- 35+ partners including Red Hat, IBM, Compiere



- Object-relational persistence framework
- Supports collections, object relations, composite types
- HQL query language, caching, JMX support



About PostgreSQL and Postgres Plus **EnterpriseDB®**



- Powerful open source object-relational database
- 20+ years of global community development
- Enterprise-class functionality

Fully ACID compliant	Views
Foreign Keys	Triggers
Joins	Stored Procedures (multiple languages)
Multi-version Concurrency Control	Point-in-Time Recovery
Tablespaces	Asynchronous Replication
Nested transactions	Heap Only Tuples
Online/hot backups	Write-ahead logging (fault tolerance)

- Mature Hibernate dialect



- Application Server overloaded
- Database overloaded
- Slow database queries




Developers:
*“this database
is a dog”*



DBAs:
*“quit sending all
that inefficient
SQL”*

Tweakers.net Database-simulatie
Overzicht - MySQL 5.0.20a



Three popular scaling strategies:

1. Add more hardware to scale horizontally or vertically
2. Improve the performance of individual operations
 - The faster an operation completes, the sooner the system is available for other work
3. Eliminate non-critical operations to reserve resources
 - Reduce the overhead of business logic execution to keep resources available for handling user loads

- Many database administrators have never heard of it
 - ‘Isn't that what bears do?’
- When they do hear of it, they rarely like it
- The business and development benefits are not always readily apparent to database administrators




- When loading objects with a parent child relationship, there are a large number of selects performed on the database
 - The N+1 problem

```
LOG:  execute S_1: BEGIN
LOG:  execute <unnamed>: select fdgroup0_.fdgrp_cd as fd...
DETAIL:  parameters: $1 = '0100'
LOG:  execute <unnamed>: select fooddeses0_.fdgrp_cd as ...
DETAIL:  parameters: $1 = '0100'
LOG:  execute S_2: COMMIT
LOG:  execute S_1: BEGIN
LOG:  execute <unnamed>: select fdgroup0_.fdgrp_cd as fd...
DETAIL:  parameters: $1 = '0100'
LOG:  execute <unnamed>: select fooddeses0_.fdgrp_cd as ...
DETAIL:  parameters: $1 = '0100'
LOG:  execute S_2: COMMIT
DETAIL:  parameters: $1 = '0100'
LOG:  execute <unnamed>: select fooddeses0_.fdgrp_cd as ...
DETAIL:  parameters: $1 = '0100'
LOG:  execute S_2: COMMIT
```


- In many cases, the default fetching settings of “lazy” and “select” work well
 - When working with related tables that have many child records that are needed, this is inefficient

```
@Entity
@Table(name = "fd_group")
public class FdGroup implements Serializable {
    private static final long serialVersionUID = 1L;
    @Id
    @Basic(optional = false)
    @Column(name = "fdgrp_cd")
    private String fdgrpCd;
    ...
    @OneToMany(cascade = CascadeType.ALL, mappedBy = "fdgrpCd")
    @org.hibernate.annotations.Fetch(org.hibernate.annotations.FetchMode.JOIN)
    private Collection<FoodDes> foodDesCollection;
```

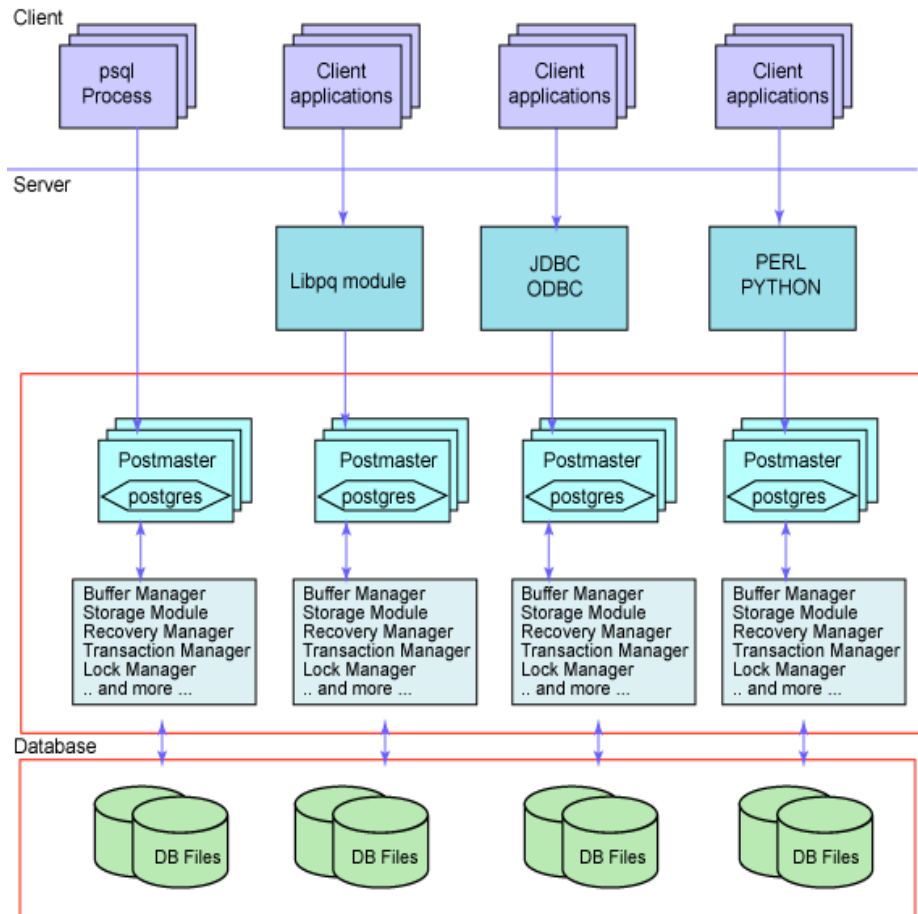
- The number of reads hitting the database are exceeding what a single server can handle.



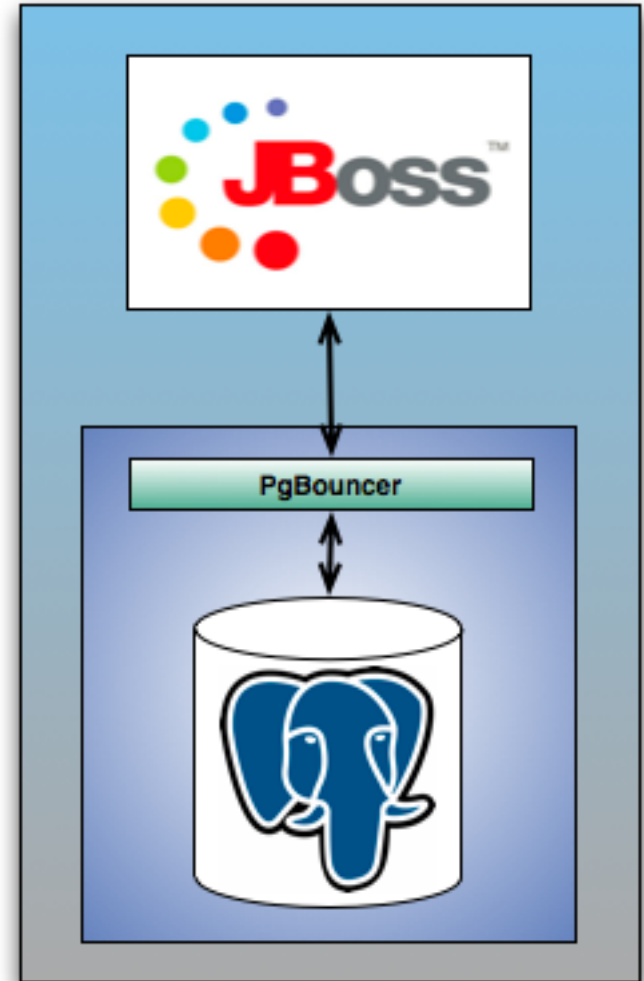
- A second level cache eliminates a lot of interaction with the database
- There are many options depending on the needs on the application
 - EhCache
 - OSCache
 - Terracotta
 - Infinispan
- Be aware of direct database changes causing cache invalidations



- As the number of concurrent users increase, the database has trouble dealing with all of the new connections
 - Postgres is process based so a new database connection is very expensive
 - This expense is magnified in a Windows environment




- PgBouncer provides an easy way to administer connection pool for Postgres
 - This will be transparent to all Hibernate configurations
- Unfortunately, all data must pass through an additional layer adding overhead
- Ideal for use when clustering the application servers and a common pool across all servers is needed



- A connection pool within the application server eliminates the overhead of a layer
 - C3P0, DBCP, and Proxool can all serve this purpose
 - Easily setup using the standard configuration files

```
hibernate.connection.driver_class = org.postgre...
hibernate.connection.url = jdbc:postgresql://lo...
hibernate.connection.username = postgres
hibernate.connection.password = password
hibernate.c3p0.min_size=5
hibernate.c3p0.max_size=20
hibernate.c3p0.timeout=1800
hibernate.c3p0.max_statements=50
hibernate.dialect = org.hibernate.dialect.Postg...
```



- The update statements send all of the columns even if they did not change
 - The Postgres logs show the following lines when just updating the column “filler” to “HelloWorld”

```
LOG:  execute <unnamed>: update tellers set bid=$1,  
      filler=$2, tbalance=$3 where tid=$4  
DETAIL:  parameters: $1 = '1', $2 = 'HelloWorld',  
                    $3 = '0', $4 = '1'  
LOG:  execute S_2: COMMIT
```

- This causes several problems...

- Why is this a problem?
 - The application server and the database are doing more work than necessary
 - A much longer SQL statement must be constructed by Hibernate
 - There is more network traffic as this is sent to the database
 - The database needs to parse a much bigger SQL statement
- Effect of curing this symptom (Minor)
 - Modern hardware has the CPU and network bandwidth to handle this additional overhead
 - Probably only noticeable with a large number of transactions

- Why is this still a problem?
 - Foreign Key Integrity checks are fired unnecessarily
 - Adds significantly more load to the database for each update
 - This grows for each Foreign Key on the table
 - The amount of work increases as the size of the related table increases
- Effect of curing this symptom (Moderate)
 - If there are only a few Foreign Keys, this may not affect performance
 - If the related tables are small and updates are somewhat frequent, the buffers will be in the database cache

- Really, why is this a problem?
 - None of the updates can leverage HOT if there are indexes on the table
 - This can cause table and index bloat for frequently updated tables
 - Adds additional maintenance overhead to the database to clean up the bloat
- Effect of curing this symptom (Major)
 - Performance steadily degrades on frequently updated tables

- How is this solved?
 - Use the Dynamic Update annotation

```
@Entity
@Table(name = "tellers")
@org.hibernate.annotations.Entity(
    selectBeforeUpdate = true,
    dynamicInsert = true, dynamicUpdate = true)
@NamedQueries(
    {@NamedQuery(name = "Tellers.findAll",
        query = "SELECT t FROM Tellers t"),
    @NamedQuery(name = "Tellers.findById",
        query = "SELECT t FROM Tellers t WHERE t.tid = :tid"),
    @NamedQuery(name = "Tellers.findById",
        query = "SELECT t FROM Tellers t WHERE t.bid = :bid"),
    @NamedQuery(name = "Tellers.findByTbalance",
        query = "SELECT t FROM Tellers t
            WHERE t.tbalance = :tbalance"))}
public class Tellers implements Serializable {
```

- A table is getting very large and the database administrator just partitioned the table, but now the application can no longer insert into the table.

```
Could not synchronize database state with session
org.hibernate.StaleStateException: Batch update returned unexpected
row count from update [0]; actual row count: 0; expected: 1
    at org.hibernate.jdbc.Expectations$BasicExpectation.checkBatched(Expecta...
    at org.hibernate.jdbc.Expectations$BasicExpectation.verifyOutcome(Expect...
    at org.hibernate.jdbc.BatchingBatcher.checkRowCounts(BatchingBatcher.java:68)
    at org.hibernate.jdbc.BatchingBatcher.doExecuteBatch(BatchingBatcher.java:48)
    at org.hibernate.jdbc.AbstractBatcher.executeBatch(AbstractBatcher.java:246)
    at org.hibernate.engine.ActionQueue.executeActions(ActionQueue.java:237)
    at org.hibernate.engine.ActionQueue.executeActions(ActionQueue.java:141)
```

- What is the problem?
 - From the database command line, psql, a row can be inserted
 - Hibernate is expecting the row count to be 1, but Postgres is returning 0
- How can the row be inserted into the database but return a message that 0 rows have be updated?
 - Postgres uses inherited tables and triggers to implement table partitioning
 - No rows are actually inserted into the base table, so the database is behaving correctly

- How is this solved?
 - Change the Postgres trigger function to return a row

```
CREATE OR REPLACE FUNCTION accounts_insert_trigger()  
    RETURNS trigger AS  
$BODY$  
DECLARE  
    ret accounts%ROWTYPE;  
BEGIN  
    ret.aid = -1;  
    .  
    .  
    .  
    RETURN ret;  
END;  
$BODY$  
LANGUAGE 'plpgsql' VOLATILE
```

- But there are some side affects to this approach

- Returning a row from the trigger function leaves a row in the base table which is not wanted
 - Create an After trigger to remove the row

```
CREATE OR REPLACE FUNCTION accounts_insert_clean_trigger()  
  RETURNS trigger AS  
  $BODY$  
BEGIN  
  DELETE FROM accounts WHERE aid = -1;  
  RETURN NULL;  
END;  
$BODY$  
LANGUAGE 'plpgsql' VOLATILE
```

- This is not very scalable. 2 inserts and 1 delete to create a single row is not practical

- What is a better way to solve this?
 - Use the SQLInsert annotation to suppress the row count check when inserting the row

```
@Entity
@Table(name = "accounts")
@SQLInsert(
    sql="INSERT INTO accounts (abalance, bid, filler, aid)VALUES (?, ?, ?, ?)",
    check=ResultCheckStyle.NONE)
@NamedQueries({
    @NamedQuery(name = "Accounts.findAll",
        query = "SELECT a FROM Accounts a"),
    @NamedQuery(name = "Accounts.findByAid",
        query = "SELECT a FROM Accounts a WHERE a.aid = :aid"),
    @NamedQuery(name = "Accounts.findByBid",
        query = "SELECT a FROM Accounts a WHERE a.bid = :bid")})
public class Accounts implements Serializable {
```

- The diagram illustrates a database schema with the following tables and their attributes:

 - public.nutr_def**
 - nutr_no BPCHAR(3) (PK)
 - units TEXT(2147483647) (FK)
 - tagname TEXT(2147483647) (FK)
 - nutrdesc TEXT(2147483647) (FK)
 - num_dec INT(2)(5)
 - sr_order INT(4)(10)
 - Index nutr_def_pkey(nutr_no)
 - public.nutr_data**
 - ndb_no BPCHAR(5) (FK)
 - nutr_no BPCHAR(3) (FK)
 - nutr_val FLOAT8(17) (FK)
 - num_data_pts FLOAT8(17) (FK)
 - std_error FLOAT8(17) (FK)
 - src_cd INT(4)(10) (FK)
 - deriv_cd TEXT(2147483647) (FK)
 - ref_ndb_no BPCHAR(5) (FK)
 - add_nutr_mark BPCHAR(1) (FK)
 - num_studies INT(4)(10)
 - min FLOAT8(17) (FK)
 - max FLOAT8(17) (FK)
 - d1 INT(4)(10) (FK)
 - low_eb FLOAT8(17) (FK)
 - up_eb FLOAT8(17) (FK)
 - stat_cmt TEXT(2147483647) (FK)
 - cc BPCHAR(1) (FK)
 - Index nutr_data_pkey(ndb_no, nutr_no)
 - Index nutr_data_deriv_cd_idx(deriv_cd)
 - Index nutr_data_nutr_no_idx(nutr_no)
 - Index nutr_data_src_cd_idx(src_cd)
 - public.food_des**
 - ndb_no BPCHAR(5) (FK)
 - fdgrp_cd BPCHAR(4) (FK)
 - long_desc TEXT(2147483647) (FK)
 - short_desc TEXT(2147483647) (FK)
 - commname TEXT(2147483647) (FK)
 - manufacturer TEXT(2147483647) (FK)
 - survey BPCHAR(1) (FK)
 - ref_desc TEXT(2147483647) (FK)
 - refuse INT(4)(10) (FK)
 - sdname TEXT(2147483647) (FK)
 - n_factor FLOAT8(17) (FK)
 - pro_factor FLOAT8(17) (FK)
 - fat_factor FLOAT8(17) (FK)
 - cho_factor FLOAT8(17) (FK)
 - Index food_des_pkey(ndb_no)
 - Index food_des_fdgrp_cd_idx(fdgrp_cd)
 - public.weight**
 - ndb_no BPCHAR(5) (FK)
 - seq BPCHAR(2) (FK)
 - amount FLOAT8(17) (FK)
 - memo_desc TEXT(2147483647) (FK)
 - gm_wgt FLOAT8(17) (FK)
 - num_data_pts INT(4)(10) (FK)
 - std_dev FLOAT8(17) (FK)
 - Index weight_pkey(ndb_no, seq)
 - public.fid_group**
 - fdgrp_cd BPCHAR(4) (FK)
 - tdgrp_desc TEXT(2147483647) (FK)
 - Index fid_group_pkey(fidgrp_cd)
 - public.src_cd**
 - src_cd INT(4)(10) (FK)
 - srcdesc TEXT(2147483647) (FK)
 - Index src_cd_pkey(src_cd)
 - public.deriv_cd**
 - deriv_cd TEXT(2147483647) (FK)
 - derivod_desc TEXT(2147483647) (FK)
 - Index deriv_cd_pkey(deriv_cd)
 - public.data_src**
 - datsrc_id BPCHAR(6) (FK)
 - authors TEXT(2147483647) (FK)
 - title TEXT(2147483647) (FK)
 - year INT(4)(10) (FK)
 - journal TEXT(2147483647) (FK)
 - vol_city TEXT(2147483647) (FK)
 - issue_state TEXT(2147483647) (FK)
 - start_page TEXT(2147483647) (FK)
 - end_page TEXT(2147483647) (FK)
 - Index data_src_pkey(datsrc_id)
 - public.datsrcn**
 - ndb_no BPCHAR(5) (FK)
 - nutr_no BPCHAR(3) (FK)
 - datsrc_id BPCHAR(6) (FK)
 - FK (ndb_no, nutr_no)
 - FK (ndb_no, nutr_no)
 - FK (ndb_no, nutr_no)
 - Index datsrcn_pkey(ndb_no, nutr_no, datsrc_id)
 - Index datsrcn_datsrc_id_idx(datsrc_id)

- Denormalization is not always bad
 - But be careful not to have multiple versions of the truth
- Reduces the number of database calls or joins to get the data for the user
- A common technique is using Materialized Views or OLAP cubes
 - These are not native constructs in Postgres
 - But...they can be created with native Postgres constructs like Rule and Triggers

- Some of the advanced features in Postgres will speed up the queries
 - Windowing Functions, Hierarchical Queries, Spatial, etc
- HQL does not support it and throws an exception

```
unexpected token: OVER near line 1, column 121 [SELECT insurance.insur...
```

- Send native SQL to the database

```
String SQL_QUERY = "SELECT insurance_name, id, invested_amount, avg(i...  
    + "invested_amount - avg(invested_amount) OVER(PARTI...  
    + "FROM insurance ";  
Query query = session.createQuery(SQL_QUERY)  
    .addScalar("insurance_name", Hibernate.STRING)  
    .addScalar("id", Hibernate.LONG)  
    .addScalar("invested_amount", Hibernate.LONG)  
    .addScalar("a", Hibernate.DOUBLE)  
    .addScalar("diff", Hibernate.DOUBLE);
```


- This adds database specific code inside the application

- Use a database view


```
CREATE VIEW insurance_diff AS
  SELECT insurance_name, id, invested_amount,
         avg(invested_amount) OVER(PARTITION BY insurance_name),
         invested_amount - avg(invested_amount)
         OVER(PARTITION BY insurance_name) AS diff
FROM insurance
```

- Allows the use of standard HQL
 - Prevents database specific code inside the application
- Puts the database specific code inside the database

- The number of reads hitting the database are still exceeding what a single server can handle.
- Sometimes horizontal scaling is necessary



- Replicate using Slony to additional servers
 - Asynchronous single master multiple slave
- Load Balance reads using pgpool-II
- All writes go to the master node
- Connection pooling is also handled by pgpool-II



- The possibility of row version problems exists
 - The application can read a row from the slave server before a new version is replicated
- Optimistic Concurrency Control can be used to solve these problems
 - Row versions are handled by the application by adding a version column to the tables

```
public class Accounts implemen...
    private static final long ...
    @Id
    @Basic(optional = false)
    @Column(name = "aid")
    private Integer aid;
    @Column(name = "bid")
    private Integer bid;
    @Column(name = "abalance")
    private Integer abalance;
    @Column(name = "filler")
    private String filler;
    @Version
    @Column(name = "version")
    private Integer version;
```

- Fetching strategies
- Caching
- Setup connection pooling
- Reduce the overhead of updates
- Partition large tables
- Optimize the data model
- Use advanced database features
- Replication

Thank you.

Questions?