Introduction to Apache Derby

Dan Debrunner
Senior Technical Staff Member
IBM

David Van Couvering
Senior Staff Engineer
Sun

TS-3154

Where Would You Like Your Data?

Learn more about Apache Derby,
the database that can go anywhere

Agenda

Overview

What Is Derby Good For?

Derby Architecture

Overview

- Derby enables data anywhere
- A complete database in a small package
- Mature, robust, performant, supported
- Community-based open source

With Apache Derby, Your Data Can Go Anywhere

In a browser,
memory stick,
laptop,
server machine,
mainframe,
PDA...

—anywhere Java™ technology goes

How Does Derby Do This?

- Pure Java technology
- Easy to use
- Embeddable and client/server
- Small footprint
- Secure

A Complete Database in a Small Package

- JDBC™ Optional Package for CDC/Foundation Profile (JSR 169)
- SQL92 and SQL99
- Java based procedures, triggers, referential constraints, fully transactional with recovery
- Online backup/restore
- Database encryption
- 2 MB runtime footprint

Java Technology in Your Database

```
CREATE FUNCTION SEND_MAIL(
  TO_ADDRESS VARCHAR(320),
  SUBJECT VARCHAR(320),
  BODY VARCHAR(32000)) RETURNS INT
LANGUAGE JAVA PARAMETER STYLE JAVA NO SQL
EXTERNAL NAME 'testing.MailTest.sendSMTP_F';

-- Send a Welcome e-mail when new customers are added.
CREATE TRIGGER WELCOME_CUSTOMER
AFTER INSERT ON CUSTOMER REFERENCING new_table AS newtab
FOR EACH STATEMENT MODE DB2SQL
SELECT SEND_MAIL(c.email, 'Welcome to AcmeWidgets',
M.email_text)
FROM newtab C, MAILINGS M
WHERE C.TYPE = M.CUST_TYPE AND M.OFFER_TYPE = 'welcome'
```

A SQL Function in Java Technology

```
public static int sendSMTP_F
   (String toAddress, String subject, String content)
{
   recipient = new InternetAddress(toAddress);
   ...
   msg = new MimeMessage(session);
   msg.setFrom(from);
   msg.setSubject(subject);
   msg.setText(content);
   msg.addRecipient(Message.RecipientType.TO, recipient);

   javax.mail.Transport.send(msg);
   return 0;
}
```

See http://wiki.apache.org/db-derby/SendEmailRoutine

Mature and Supported

- In production since 1997
- In use by numerous products
  - http://wiki.apache.org/db-derby/UsesOfDerby
- Supported versions from Sun and IBM

Derby Performs

- Comparable with other open source databases
  - Outperforms MySQL on large databases
  - MySQL better on small main-memory database
  - No significant performance loss with client/server, except for SELECT operations
  - See http://tinyurl.com/nxmj8
- Ongoing improvements
- Tips
  - Prepared statements, cache size, indices

Derby Holds a Lot of Data

- Theoretical limit is 2^64 bytes per table and 2^64 tables
- Testing up to 250GB in 10.2 release
- Near-term goal: test up to 1 TB

You Can Get Involved!

- Apache Community
- Anyone can contribute
- http://db.apache.org/derby

What Is Derby Good For?

- Browser local store
- Middle-tier cache
- Database on the go
- Read-only database
- Easy-to-manage departmental database
- Development and unit testing

Deploy Derby Invisibly in a Browser

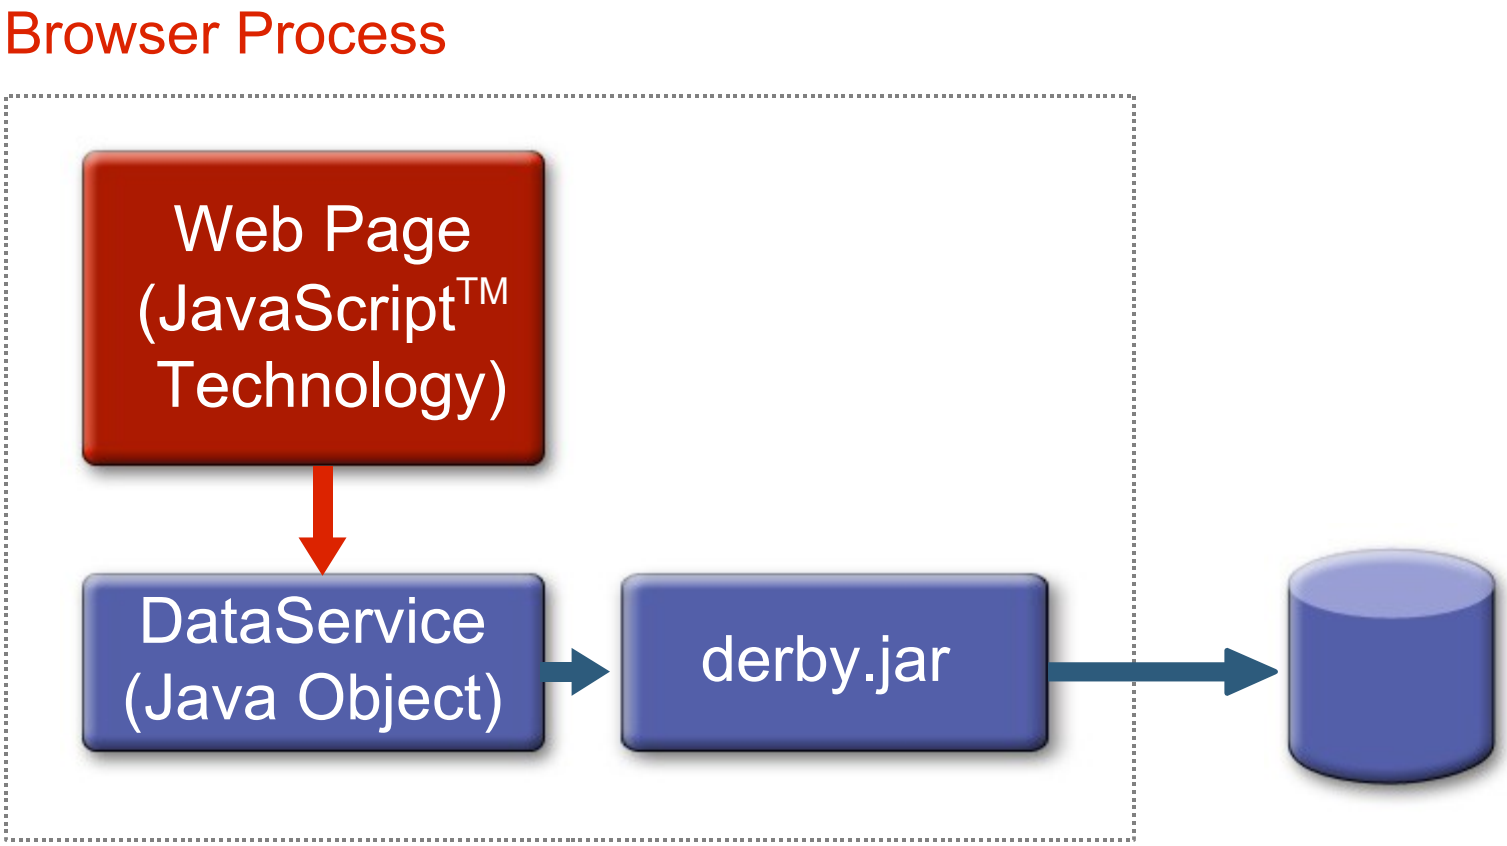

Example Use Cases

- Store sensitive data on user machine
- Offline web mail or blogs
- Local cache
- Retain state even if browser exits

How Derby Makes This Work

- Embedded
- Invisible
- 2 MB Footprint
  - Compressible to 600k, see http://blogs.sun.com/roller/page/FrancoisOrsini/20060118
- Secure
- Fast

DEMO

Running Derby embedded in a browser

For code and details, see
http://developers.sun.com/prodtech/javadb/

Derby as a Local Cache

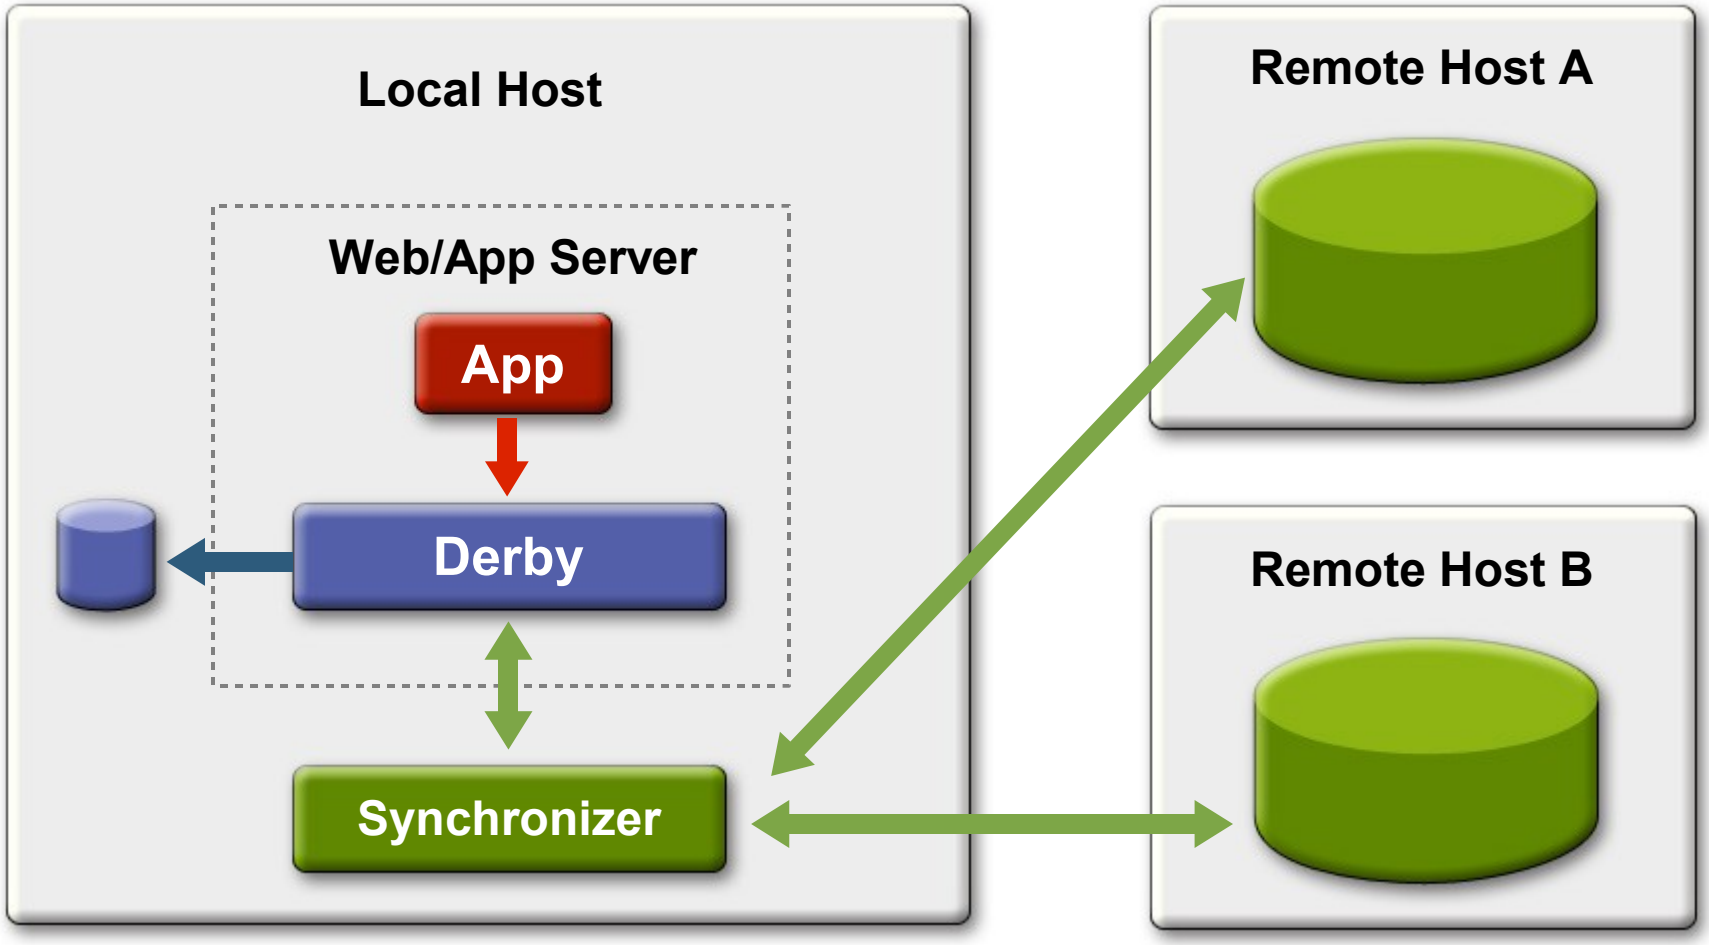

Embedding Derby

```
Connection conn =
   DriverManager.getConnection
    ("jdbc:derby:sample");
```

Enabling Embedded Derby For Clients

- In derby.properties file:

```
derby.drda.startNetworkServer=true
```

A Database That Can Go Anywhere

- Laptop, PDA, memory stick
- Small footprint
- No DBA
- Secure
- Java technology

Example Use Cases

- Metadata store for digital media collection
- Integrate with Open Office
- Store sensitive data off disk
- Retain critical data in emergency
- Up-to-date conference schedule on a stick

DEMO

Derby on a Stick

Read-Only Database

- Can be placed onto CD-ROM
- Can put inside your JAR/EAR/WAR file

Example Use Cases

- Sales or product catalog on CD-ROM
- Auto-updated local cache of sales price list on sales laptop

Read-Only Database in a JAR

- Add database directory to your JAR

```
jar uvf myapp.jar salesCatalog
```

- Put JAR file on classpath
- Connect using standard JDBC

```
Connection conn =
  DriverManager.getConnection
    ("jdbc:derby:/salesCatalog");
```

Easy to Manage Departmental DB

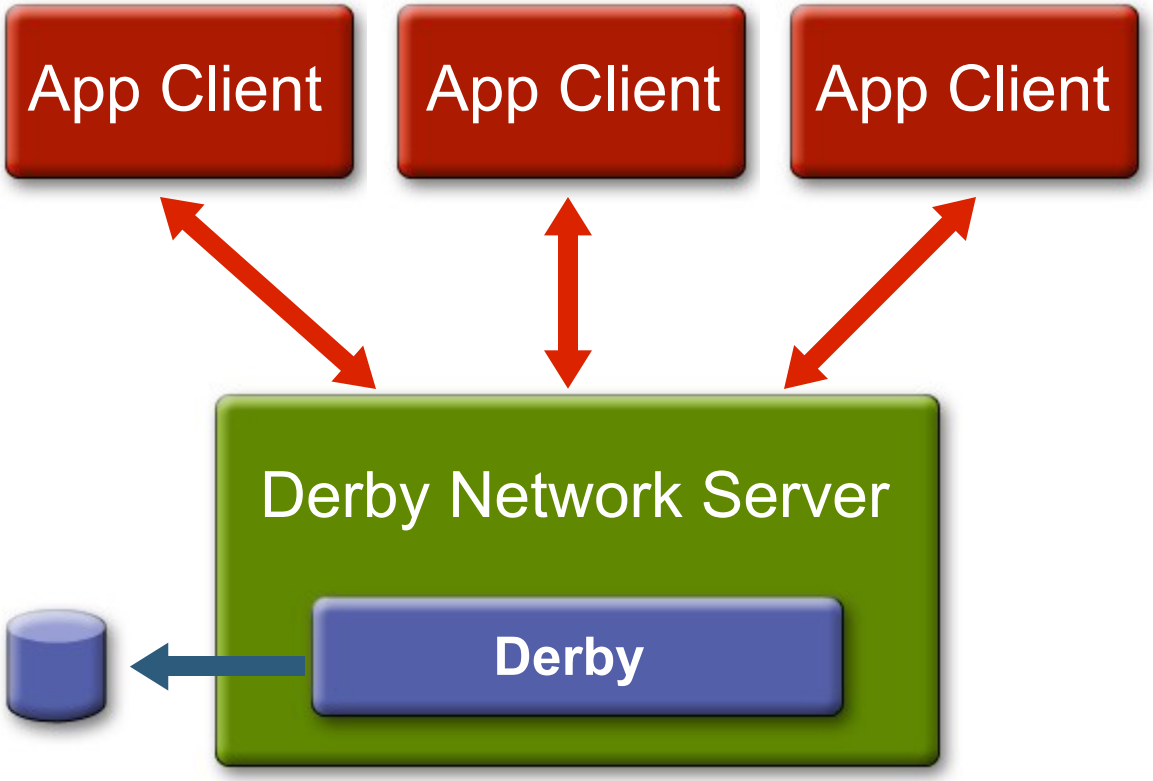

Starting Derby Network Server

now

```
startnetworkserver.[bat|ksh]
```

future

```
java -jar derbynet.jar start
```

Java Technology Development and Unit Testing

- Embedded
- JDBC 3 and 4 software, SQL92 and SQL99
- Full multi-user support
- Integrated plugin for Eclipse, NetBeans™ software

Engine Architecture

- Layered
- Modular
- Cost-based optimizer
- Queries compiled into byte-code
- Fully transactional and recoverable

Layered Architecture

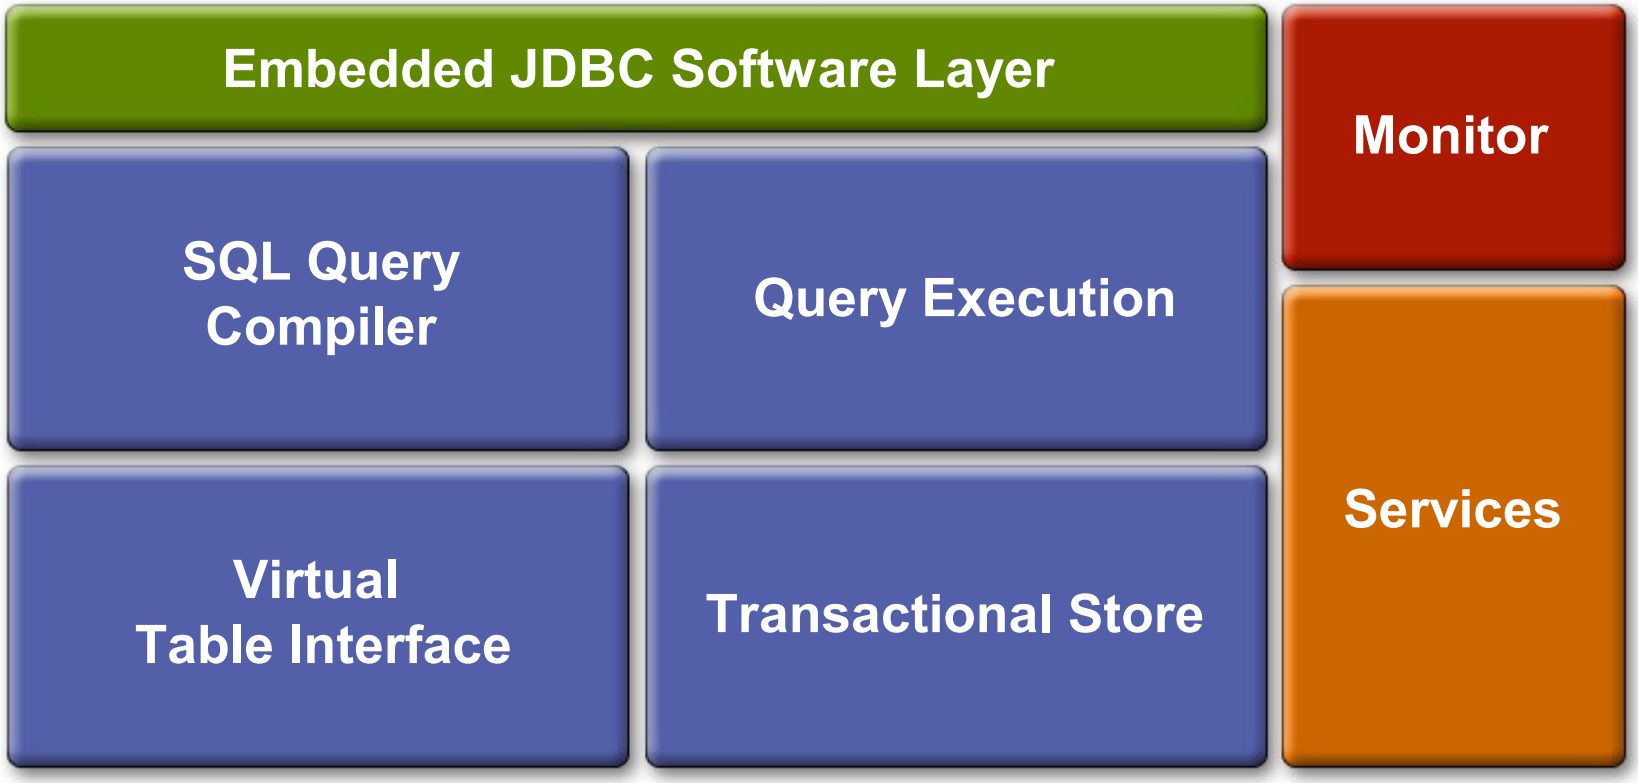

Pluggable Infrastructure

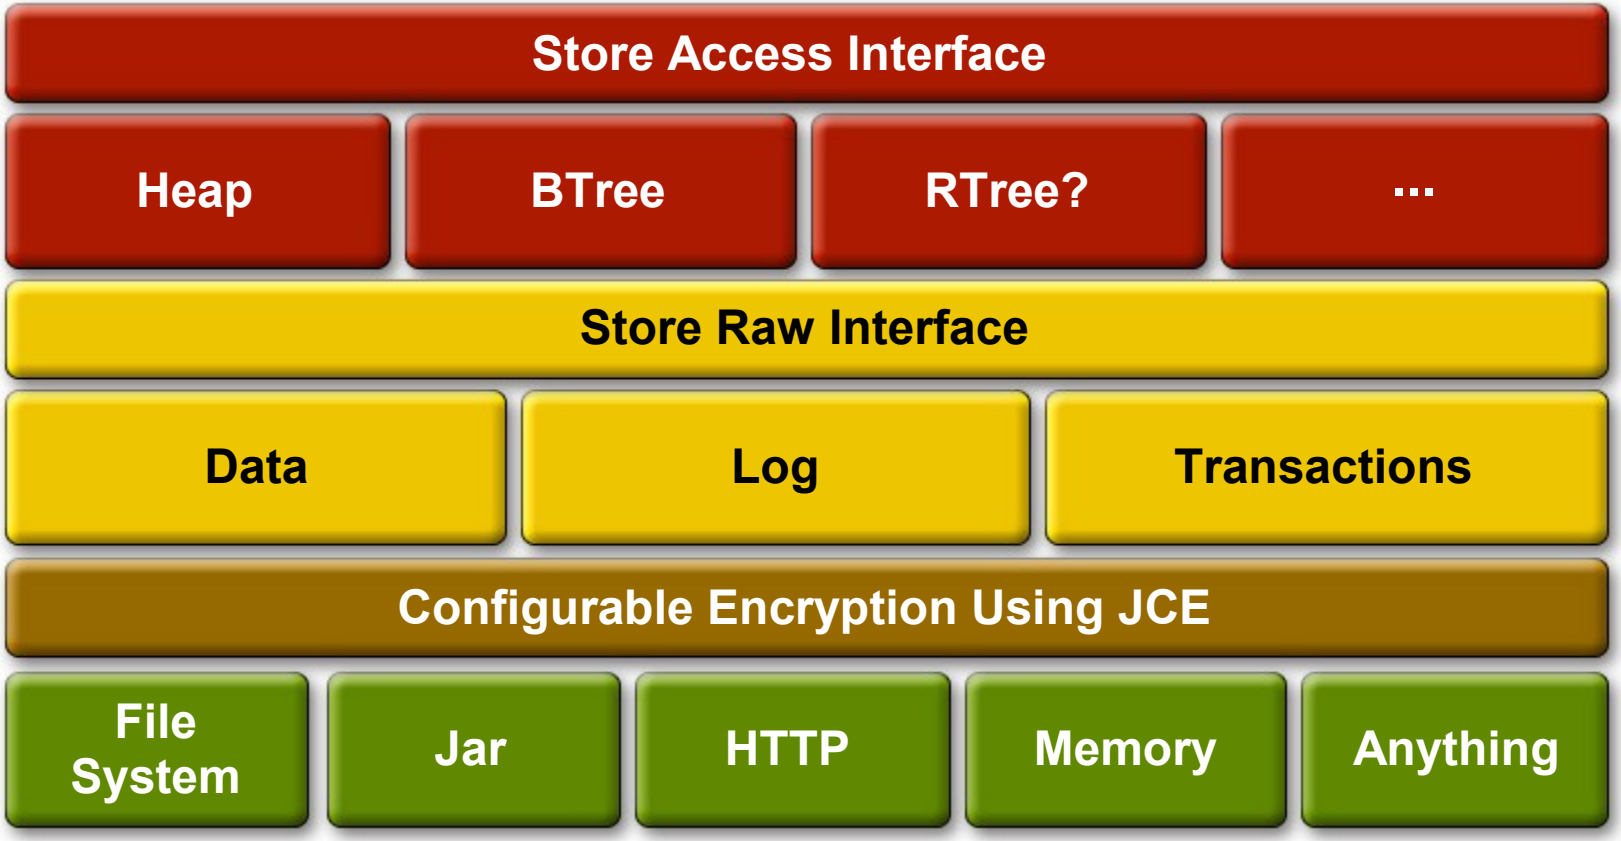

SQL Optimizer

- Preprocess compilation phase
  - Internal query modifications
  - Aimed at giving optimize phase more choices
- Plans selected by cost estimates
  - Data distribution statistics automatically maintained
  - ORDER-BY eliminations for index scans
- Optimizer Overrides
  -
```
SELECT * FROM
--DERBY-PROPERTIES joinOrder=FIXED
t1, t2 WHERE t1.c1=t2.c1
```

SQL Compiled to Java Byte Code

SELECT * FROM CUSTOMER WHERE ID = ?

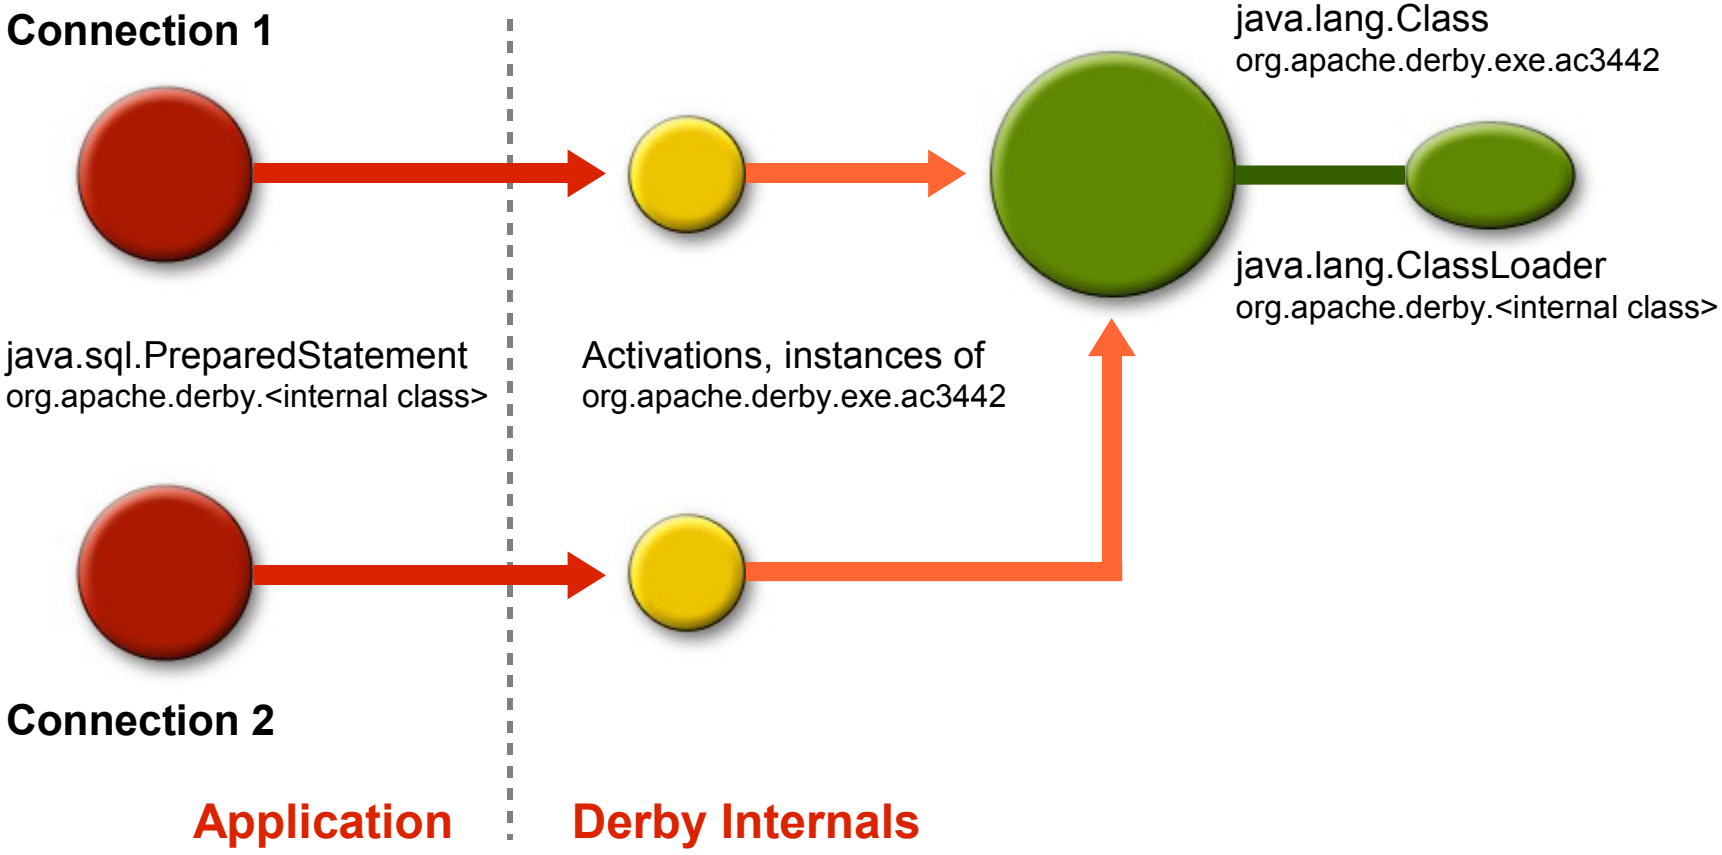

Transaction Logging

- Aries logging system for rollback and recovery
- Order of a change—write ahead logging
- Crash recovery automatically executed

Summary

- You can use Apache Derby to put your data anywhere
- Small, but fully functional engine
- Lots of great uses
- Community-based open source
- Open standards

For More Information

- Apache Derby JavaOne[SM] Conference Event
- http://db.apache.org/derby
- http://www.ibm.com/developerworks/cloudscape
- http://developers.sun.com/prodtech/javadb/
- Book: Apache Derby—Off to the Races

Q&A

Introduction to Apache Derby

Dan Debrunner
Senior Technical Staff Member
IBM

David Van Couvering
Senior Staff Engineer
Sun

TS-3154