

Using Apache Derby

The Open Source Database
Embedded in ColdFusion 8

Charlie Arehart

Tues, Nov 18



- Why are we here?
 - ColdFusion 8 includes the open source Apache Derby database
 - Easy to misconstrue what it does, doesn't do; why/why not use it
 - Will try to show here why you should consider it, even for production, transactional, multi-user applications
 - Will also discuss several topics of typical interest when considering/using it



- Audience For This Talk
- Quick Demonstration
- Fundamentals
 - Embedded vs Client/Server Mode
 - Dashing a Common Myth
 - Some Natural Questions: What is it? Why use it? And more
- Using Derby, And Tools
 - How to Use/Enable it in ColdFusion
 - Connecting to Derby Databases Using Tools
- Miscellany
 - Derby SQL Topics/Issues
 - Some Common Challenges
 - Resources for Learning More

- 11 yrs CF experience (26 in Enterprise IT)
 - Independent consultant since April 2006
 - Member, Adobe Community Experts
 - Certified Adv CF Developer (4 - 8), Certified Adobe Instructor
 - Writer in FAQU, Adobe DevCenter, CommunityMX, CFDJ, more
 - Contributor to all three Ben Forta CF8 books
 - Frequent speaker to user groups, conferences worldwide including
 - Adobe Max, cf.Objective, CFUnited, Scotch on the Rocks, webDU, webManiacs
 - Run the Online ColdFusion Meetup (coldfusionmeetup.com)
 - Living in Alpharetta, Georgia (north of Atlanta)

- Web home at www.carehart.org
 - Hosts ~300 blog entries, 70+ articles, 80+ presentations, more
 - UGTV: recordings of 300+ presentations by ~200 CFUG speakers
 - CF411.com - 700+ tools/resources in 100+ categories
 - Consulting: available for troubleshooting, tuning, training
 - For as few as days, hours; remote or on-site


Audience For This Talk

- Those just curious about it
- Those already using it
- Those skeptical about it


QUICK DEMONSTRATION

FUNDAMENTALS

- Comes in two flavors: “Embedded” and “Network” modes
 - Both support multi-user, transactional applications built with CF/other servers
- Embedded
 - Runs in-process (within CF/JRun),
 - Low overhead
 - Conservative defaults (minimal memory use)
 - Accepts connections only from threads within the embedding process (CF/JRun)



- Networked (or client/server)
 - Separately started, TCP-based network server
 - Accepts connections from other processes or hosts (outside of CF)
 - Lightweight client driver



- Distinction should become more clear in remainder of presentation

- I've heard that Derby is a single-user DBMS
 - Again, no. The misconception is due to misunderstanding that the Embedded edition accepts connections only from within the *embedding application*
 - This does NOT mean only **one user of that application!**
 - Derby, as embedded within CF, is certainly multi-threaded, supporting multiple users and databases at one time

Some Natural Questions: What is Derby?

- Let's answer several natural questions ...




- Is Derby a “real” DBMS? Absolutely

- Open source relational database implemented entirely in Java (<http://db.apache.org/derby/>)
 - Implements JDBC and SQL standards, all in a tiny footprint (~2 MB)
- Yet has advanced features you'd expect in a quality DBMS:
 - Stored procedures, triggers, views, referential integrity constraints (including cascading deletes), localization/unicode
 - Cost-based query optimization, transactions, savepoints, checkpoint/recovery, transaction logs, backup, schemas, authentication (including LDAP), encryption, and much more
 - Row-level locking with available escalation to table level when needed, deadlock detection, and more

- Is Derby something new? Not really
 - Actually started life as Cloudscape (from Cloudscape) in 1997
 - Cloudscape Inc was acquired by Informix in 1999
 - IBM acquired Informix DB assets in 2001, continued development as IBM Cloudscape
 - IBM contributed it to Apache in 2004 (continues contributing, with Sun)
 - Now known also as Sun Java DB and **included in Sun Java 1.6 SDK**
 - More at <http://developers.sun.com/javadb/>
- Why has Adobe included it in CF8?
 - Provides a built-in DBMS for use by customers, that is portable
 - One can create a scalable database without any add'l software installation
 - And one can rely on it being there if bundling/distributing a CF8 application
 - Is used for example databases (portability really helped there)



Some Natural Questions: Why Use It?

- Why use Derby versus using something like MySQL? *Apache Derby* 
 - Significant: you can know that the DB server is installed in CF8.
 - With MySQL and most other DBMS's, you must install them in addition to CF
 - Benefit both for users and those deploying apps onto CF (whether as standalone or J2EE deployment)
- Is it a pure development only DB?
 - No, it is fully intended for production. Recall the long list of features mentioned
- How does it perform and scale?
 - Demonstrated to scale very well for large volumes and loads, in various scenarios
 - Clustering support available via open source Sequoia project
 - <http://community.continuent.com/community/sequoia>
 - See *Additional Notes* (downloadable at Max Presentation Library) for links to several other resources discussing Derby performance

- Is this the same database engine that's embedded in Adobe Air?
 - No, that's SQLite: <http://www.sqlite.org>
 - Then again, Derby can compete with it, running on small devices
- Is it in both CF8 Standard and Enterprise?
 - Yes, there's no difference in use of the DB in the two editions
- Can you use it with CF7 or 6?
 - Yes, but you would need to install Derby yourself, and you need to use the "Other" database driver type to create a DSN
 - See resources listed later under "Resources on using Derby with CF, prior to CF8"

USING DERBY, AND TOOLS

- Configured in CF 8 Admin, like other DBMS's
- Observing the Example Derby Databases preconfigured in CF8
 - What they are: artgallery, bookclub; cfcodeexplorer; cfdocexamples
 - Each stored as a directory per database, within:
 - `[coldfusion8]\db`
 - `[jrun4]\servers\[instance]\cfusion-ear\cfusion-war\WEB-INF\cfusion\db\`

- **Configuring a New Derby Datasource in CF**
 - Choose Embedded or “Client” (Network) mode
 - Embedded mode runs Derby as embedded within CF
 - Client/Network mode requires running a separate Server instance of Derby
 - Choice is driven by whether multiple applications need to access DB
 - Again, not whether multiple users need to access it
 - Setting up the optional Network Server
 - See various resources offered at end of talk
 - Choose whether you’re creating new database (or pointing to an existing one)
 - Check box option in CF Admin, if creating “Embedded” DSN
 - or can indicate via connectionstring: Create=true
 - Can also create programmatically with CF Admin API or DerbyCFC (<http://derbycfc.riaforge.org/>)

- What actions are you looking for?
 - Access from a favored IDE? Or any IDE that helps? Or the command line?
 - Ability to view data? Create SQL? View/change database schema/relationships?
 - Run test SQL statements?
 - Generate CFC/CRUD code? Build complete CRUD applications?
- Each of these is available in different forms
 - Tools that leverage CF's RDS feature and any CF DSN database type
 - Generic command line tools that support Derby (and don't know about CF)
 - IDEs that support Derby (and don't know about CF)

- Tools based on CF's RDS feature
 - Can access **Embedded** databases as defined within CF, using CF's RDS feature
 - For more on using RDS, see docs or my presentation (<http://carehart.org/presentations/#rds2>)
 - DreamWeaver, HomeSite+, CF Studio
 - Can look at a Derby database, see its tables and columns, view its data, and build SQL query/update statements
 - Eclipse (using Adobe CF8 extensions for Eclipse)
 - Adds ability to generate CFCs from tables to perform CRUD (create, read, update, delete) operations
 - <http://www.adobe.com/support/coldfusion/downloads.html#cfdevtools>
 - None of these offer visual tools to **create/modify** schema: tables/columns, etc (DDL).

- CF Tools NOT based on CF's RDS feature
 - CF-based DB tools specifically supporting CF/Derby Tools
 - DerbyCFC: <http://derbycfc.riaforge.org/>
 - CF-based DB tools that support Derby just as all other CF DSNs
 - CF DB Explorer: <http://cfdbexplorer.riaforge.org/>
 - DataMGR (CRUD creator): <http://datamgr.riaforge.org/>

- Other Derby Developer tools: command line
 - ij - SQL scripting tool (“eye-jay”)
 - `java -jar c:\coldfusion8\lib\derbyrun.jar ij`
 - dblook – schema dump/extraction tool
 - sysinfo – Derby system information
 - More on using these
 - http://www.java2s.com/Tutorial/Java/0340_Database/JavaDBTools.htm
- Other Derby Developer tools: IDEs
 - Available features to work with Derby from Eclipse and other GUI DB tools
 - Eclipse Derby and DTP plugins
 - Non-Eclipse tools like Squirrel, RazorSQL, Aqua Data Studio
 - Some support DDL as well as DML SQL
 - These do not rely on RDS connection
 - Cannot access an embedded database while CF is running. Use Network mode

MISCELLANY

- Supports JDBC 4, entry-level SQL-92, as well as some higher-level features of SQL-99 and newer
 - “SQL-99 and SQL-2003 features mapped to Derby”
 - <http://wiki.apache.org/db-derby/SQLvsDerbyFeatures>
- Some challenges encountered by CFers
 - Various challenges with identity columns, nulls, columns with spaces, more
 - http://www.bryantwebconsulting.com/blog/index.cfm/2007/6/8/using_the_derby_database
 - Lack of TOP keyword
 - See comments at <http://www.alagad.com/go/blog-entry/derby--sqlite-and-other-derby-myths>
 - Pagination (see comments at <http://www.alagad.com/go/blog-entry/derby--sqlite-and-other-derby-myths>)

- **10.2 – This is the version included in CF 8**
- 10.3
 - New features included Drop/Rename column, Blob/Clob support, more
 - <http://wiki.apache.org/db-derby/DerbyTenThreeRelease>
- 10.4
 - New features included asynchronous replication, JMX, more
 - <http://wiki.apache.org/db-derby/DerbyTenFourRelease>
- More details and additional older (and possibly newer) releases:
 - http://db.apache.org/derby/derby_downloads.html
- While you could download and implement updated versions, they would not be supported by CF8
 - Perhaps CF9 will include later updates

Some Common Challenges with Derby on CF



- A couple of known problems, some fixed in 8.01, some remain
- Also:
 - Creating a new datasource sometimes goes to “Other” driver screen
 - Problem with Unicode, easily solved
 - Renaming a column (via Alter Table)
- See *Additional Notes* for more

- If updating data in Derby databases, this calls for using orderly shutdown of CF (not just cancelling the jrun.exe task)
 - Otherwise Derby may spend extra time doing recovery at next startup
- To back up a database, you can use the online backup utility.
 - For information on this utility, see the *Derby Server and Administration Guide*
- Various Derby environment settings can be tweaked
 - See *Additional Notes* document for details, ideas, additional resources

- Other Derby Features
 - Zero Administration
 - available options, generally need not be tweaked
 - Hibernate support (<http://wiki.apache.org/db-derby/HibernateHelp>)
 - Spring support: (<http://www-128.ibm.com/developerworks/web/library/wa-spring2/>)
 - DB can be stored on CD
 - See “Creating Derby databases for read-only use” in *Derby Developer's Guide*
- Other possible Derby uses:
 - As a pre-configured demo database for your existing application
 - As a testing database and a way of verifying database independence
 - To replace existing file based storage

- Resources describing Derby in a CF Context
 - My “Getting Started with the Apache Derby Open Source Database in CF8”
 - Derby subsite at CFCommunity.org
 - “Working with the Apache Derby Database and ColdFusion”, by Chip Temm
 - “Using Apache Derby - The built-in CF8 Database”, by Scott Stroz
 - “Derby != SQLite and Other Derby Myths”, by Scott Stroz
 - “Using the Derby Database”, by Steve Bryant
 - “Creating Derby embedded databases in ColdFusion 8”, Kay Smoljak
 - Exploring Your DB With Eclipse 101 (Including Derby DBs!), Todd Sharp
- Resources on using Derby with CF, prior to CF8
 - “Installing Apache Derby Database on ColdFusionMX 7.0.2”
 - “ColdFusion Everywhere PART 2”
- See *Additional Notes* for links to all these

- General Interest Derby Resources
 - Derby manuals (<http://db.apache.org/derby/manuals/index.html>)
 - 10.2 manuals: http://db.apache.org/derby/manuals/index.html#docs_10.2
 - *Getting Started with Derby* (42pp)
 - *Derby Reference Manual* (291pp)
 - *Derby Developer's Guide* (135pp)
 - *Tuning Derby* (98pp)
 - *Derby Server and Administration Guide* (74pp)
 - *Derby Tools and Utilities Guide* (74pp)
 - Other Derby sites and resources
 - See *Additional Notes*

- Derby is a very robust, substantially complete DBMS
 - Not at all just a development database, nor for single-user use
- We've answered many common questions, misunderstandings
- Saw it can be used in either embedded or client/network mode
- Saw how to enable Derby DBs within CF
- Learned of many options for tools to access Derby DBs
 - Some built into normal CF editors (DW, HS, CFS, Eclipse)
 - Some through various other Derby-specific tools
- Discussed various SQL issues, challenges in using Derby with CF
- Pointed out many, many available resources to learn more

- Please fill out the evaluation
- Available for ColdFusion troubleshooting, setup, implementation consulting
 - Also other developer productivity coaching, system admin and tuning support, and more
 - Remote or on-site
 - Again, for as little as days, hours, even 15 minutes
 - <http://carehart.org/consulting/>