

Makers of **Berkeley DB**

SLEEPYCAT
SOFTWARE



The B in LAMP

Fourth Annual Southern California Linux Expo

David Schachter
Performance Engineer
Sleepycat Software

Agenda

- 1. Introductions: Who, What, Why?**
- 2. BDB and Linux**
- 3. BDB and Apache**
- 4. BDB and/or MySQL**
- 5. BDB and {Perl, PHP, Python}**
- 6. How to Get the Latest, Greatest Stuff**
- 7. Berkeley DB in Action**

Makers of **Berkeley DB**



Who, What, Why?

The Company and Its Customers

Did you know that Berkeley DB..?

- **is everywhere:** runs on 200 million machines across the world
- **scales up, scales down:** powers everything from mobile phones to stock exchanges
- **is proven:** runs inside every copy of Linux (and MacOS X and Solaris)
- **powers the Internet:** extensively used by Google, Yahoo, AOL, Ask Jeeves, Amazon, eBay and many others
- **is highly available:** handles 70% of the global daily email traffic

About Sleepycat

- Founded in 1996 by original architects of BSD Unix
- Executives include industry luminaries
 - **Dr. Margo Seltzer** (Harvard Professor, 4.4 BSD file system author)
 - **Keith Bostic** (2.10 BSD architect, 4.4 BSD developer)
 - **Michael Ubell** (Ingres, DEC, Britton-Lee [founder], Illustra, Informix)
 - **Michael Olson** (Britton-Lee, Illustra, Informix)

“Open source isn't some magic pixie dust that just makes products great. It's not a business model by itself. It is a tool that companies can use to leverage popular projects and put them into production in profitable and useful ways.”

Mike Olson, Sleepycat CEO

In “What open source software buyers can learn from Google”, SearchOpenSource.com, 16-Nov-05

Berkeley DB in Open Source



- Linux
- BSD Unix, MacOS X Server



- Apache



- OpenLDAP



- Sendmail



- Kerberos



- Zope



- OpenOffice



- GNU C library

- Movable Type blog



- Subversion



- Perl



- Python



- PHP



- Eiffel

Language Bindings



- Tcl



- Ruby



...and many, many more

Berkeley DB Java Edition new API discussion:

http://www.theserverside.com/news/thread.tss?thread_id=38916

Public License

Permits use in open source projects or applications that are not distributed to third parties.

Commercial License

Permits closed-source distribution of an application to third parties and provides business assurance.

Advantages

- Very broad, low cost distribution
- Very high code quality
- Barrier to entry

- High margin, licensing-based revenue stream
- Profitable, sustainable business model

Customer Benefits

- Huge developer community for support
- Peer-reviewed design for ease of use
- Less debugging, easy integration
- Easy download, trial, and distribution
- No escrow issues
- Freedom from vendor lock-in

- Rights to redistribute proprietary code
- Representations, warranties and indemnification
- Full-time, dedicated development team (continuity, roadmap, documentation)
- Single vendor to hold accountable

Makers of Berkeley DB



LAMP:

Linux

Apache

MySQL

Perl/Python/PHP

Makers of Berkeley DB



LAMP:

Linux

Apache

MySQL

Perl/Python/PHP

Berkeley DB in Linux

- RPM
 - OpenLDAP
 - Login
 - Lots more...
- ```
“rpm -q -whatrequires
libdb-4.0.so” on Redhat 7:
```
- apr-util-0.9.4-4
  - cyrus-sasl-2.1.10-4
  - httpd-2.0.48-4
  - pam-0.75-48
  - webalizer-2.01\_10-11

# Berkeley DB 2 in Debian Linux 3.1

- zmailer-ssl
- pavuk
- libapache-mod-fastcgi
- snes9express
- skksearch
- qtstalker
- pavuk
- htdig

# Berkeley DB 3 in Debian Linux 3.1

- gsnes9x
- yank
- xsitecopy
- xgsmlib
- xemacs21-nomule
- xemacs21-mule-canna-wnn
- xemacs21-mule
- xemacs21-gnome-nomule
- xemacs21-gnome-mule-canna-wnn
- xemacs21-gnome-mule
- xemacs21-bin
- xcin
- gpredict
- gpr
- gphotocoll
- gpgp
- godbcconfig
- gnucash-hbci
- gnucash
- gnotepad+
- gnomp3
- gnomermind
- gnomekiss
- gnome-think

# Berkeley DB 4 in Debian Linux 3.1

- apt-utils
- arla
- bogofilter
- caudium-php4
- cfengine
- cfengine2
- drac
- dsniff
- etpan-ng
- evolution
- evolution-data-server
- evolution-exchange
- libapache2-mod-xmlrpc2
- libapache2-request-perl
- Libapr0
- libchise1
- libedataserver3
- libetpan1
- libhdb7-heimdal
- libnss-ldap
- libopenft-gift
- libotp0-kerberos4kth
- librdf-perl
- librdf-ruby
- librdf0



Makers of Berkeley DB



# LAMP:

Linux

Apache

MySQL

Perl/Python/PHP

# Berkeley DB in Apache

- Mod4DB
- SpamAssassin
- libapache-modxslt
- libapache-mod-ssl
- libapache-mod-php4
- libapache-mod-iptables
- apache2-utils
- apache2-mpm-worker
- apache2-mpm-prefork
- apache2-mpm-perchild
- apache2-common
- apache-ssl
- apache-perl
- apache-dbg
- apache-common
- apache

Makers of **Berkeley DB**



# LAMP:

Linux

Apache

**M**ySQL

Perl/Python/PHP

# Berkeley DB in MySQL

- MySQL storage engine (first one!)
- Helped MySQL start
- Cooperators building open source

# Comparison of BDB and MySQL

## Berkeley DB

- Fast
- Small
- Simple
- Proven

## MySQL

- Full-featured
- Standards Compliant
- Flexible Data Model
- Powerful, Flexible Queries

Makers of Berkeley DB



# LAMP:

Linux

Apache

MySQL

Perl/Python/PHP

# Berkeley DB for Perl, Python, PHP

Adapters for:

- Eiffel
- Perl
- PHP
- Python
- Ruby
- TCL

Makers of **Berkeley DB**

SLEEPYCAT  
SOFTWARE



# How to Get Berkeley DB

**Hint: You already have it.**

## Download Berkeley DB now from our new Developer Zone ▶

### Now available: Berkeley DB XML Version 2.2 ▶

[Read press release](#)

#### White Papers, Videos and Case Studies:

- Berkeley DB Video**  
Margo Seltzer, founder and CTO of Sleepycat Software, provides a conceptual introduction, technical overview and programming examples for Berkeley DB.
- Google Case Study**  
Learn about how Berkeley DB High Availability powers all of Google's single sign-on and personalization services.
- Berkeley DB Performance Metrics & Benchmarks**  
Performance measurements for Berkeley DB on various platforms.
- Anatomy of a Native XML Database**  
Learn about the design and architecture of a native XML database.


#### Consulting

Let Sleepycat experts help make your project successful. [Find out more.](#)

#### Sign up for the Sleepycat newsletter

Email  [Subscribe](#)

#### Making News

- Sleepycat Ships New Version of Berkeley DB XML with 500 Percent Performance Increase
- Sleepycat Launches New Consulting Packages for Berkeley DB
- Berkeley DB XML Speeds Up Service Oriented Architecture at Starwood Hotels
- Sleepycat's Java Database Powers General Dynamics CoMotion Visualization Software for Battlefield Analysis
-  InfoWorld Test Center Berkeley DB XML scores 9.3

#### Featured Customers



Ford uses Berkeley DB to authenticate partners who access Ford's Web applications.



Hewlett Packard uses Berkeley DB in several products, including storage, security and wireless software.



RSA Security relies on Berkeley DB for digital certificate management in its Keon product line.

# Downloading Berkeley DB

- <http://dev.sleepycat.com/downloads/latestreleases.html>
  - tar, zip and VxWorks CD image
  - Windows installer with SDK and HTML documentation
- Release change logs
- Patches
- Historic distributions

# Berkeley DB Documentation

- Included in distributions
- Available online  
<http://dev.sleepycat.com/documentation/bdb.html>
- “Getting Started” Guides for C, C++, Java
- “Writing Transactional Applications” for C, C++, Java
- API Guides for C, C++, Java
- Javadoc for Java API, Java Collections Tutorial
- Guides for Building, Troubleshooting, Reference

Welcome to Sleepycat. [Sign in](#) or [register](#).

## Welcome to the Sleepycat Developer Zone

We are committed to making software developers successful. Here you will find all the resources you will need to create applications using Sleepycat products.

Help us improve the Developer Zone. We'd love your suggestions or feedback. Send email to [feedback@sleepycat.com](mailto:feedback@sleepycat.com)



**Downloads** of our full software, including source code, test code, sample code and documentation, as well as all of our historical releases.



Documentation for [Berkeley DB](#), [Berkeley DB Java Edition](#) and [Berkeley DB XML](#), including comprehensive reference guides, getting started guides and tutorials.



**Options** for getting technical support and [requesting support](#) online.



Technical resources such as [FAQs](#), [technical articles](#), [videos](#), [white papers](#), and [performance benchmarks](#).



Information about how to participate in the Sleepycat community, including [mailing lists](#), [newsgroups](#), [blog](#) and [recognition program](#).

## What's New

- 2005.11.07 Now Available: [Berkeley DB XML 2.2](#)
- 2005.11.07 New technical [FAQ](#) system
- 2005.11.07 Online: [Sleepycat blog](#)
- 2005.11.07 Submit a [Support Request](#)
- 2005.11.07 Case Study: [Google](#)
- 2005.11.01 Sleepycat will deliver a Berkeley DB technical tutorial at [USENIX LISA](#) on December 6, 2005 in San Diego, CA.
- 2005.10.27 Now Available: [Berkeley DB Java Edition 2.0.90](#)
- 2005.10.25 New [consulting](#) packages
- 2005.09.30 Now Available: [Berkeley DB 4.3.29](#)

## Direct Download

**BERKELEYDB**

**Berkeley DB 4.3.29**

With encryption

[.tar.gz](#) | [.zip](#) | [.msi](#) |

[VxWorks](#)

Without encryption

[NC.tar.gz](#) | [NC.zip](#) |

[VxWorks](#)

**BERKELEYDB**

**Berkeley DB JE 2.0.90**

[.tar.gz](#) | [.zip](#)

**BERKELEYDB**

**Berkeley DB XML 2.2.8**

[.tar.gz](#) | [.zip](#) | [.msi](#)

## Sleepycat Newsletter

Email

[Subscribe](#)

## Feedback

Help us improve the DevZone. The DevZone.

- Exceeded my expectations
- Met my expectations
- Could be improved by:

## Tag

**Sleepycat** is our tag for blogs and bookmarks.

See:

[Technorati](#)  
[del.icio.us](#)

# Technical Resources

- **New** web site: <http://dev.sleepycat.com>
  - Downloads, documentation, FAQs, videos, whitepapers, performance benchmarks, blog, newsgroups, ...
- Mail to [support@sleepycat.com](mailto:support@sleepycat.com)
  - Prioritized, logged, managed technical issues
- Technical community discussions
  - Archived, but not managed; Sleepycat participates
  - [comp.databases.berkeley-db](mailto:comp.databases.berkeley-db) newsgroup
  - [xml@sleepycat.com](mailto:xml@sleepycat.com), [bdbha@...](mailto:bdbha@...), [bdbje@...](mailto:bdbje@...)

# Technical Resources













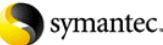














- Product announcement mailing lists
  - Very low traffic—Product announcements only
  - [sleepycat-announce@sleepycat.com](mailto:sleepycat-announce@sleepycat.com)
  - [xml-announce@sleepycat.com](mailto:xml-announce@sleepycat.com)
  - [bdbje-announce@sleepycat.com](mailto:bdbje-announce@sleepycat.com)

Makers of **Berkeley DB**



# Berkeley DB in Action

# Customers

| Devices/<br>Appliances                                                                        | Telecom<br>Infrastructure                                                           | Networking                                                                          | Security                                                                            | Enterprise<br>Infrastructure                                                         | Storage & Sys<br>Mgmt                                                                 | Enterprises                                                                           | E-Business                                                                            | Financial<br>Services                                                                 |
|-----------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|
| <br>MOTOROLA |    |    |    |    |    |    |    |    |
|              |    |    |    |    |    |    |    |    |
|            |  |  |  |  |  |  |  |  |

**Berkeley DB**

# Motorola Smart Phones



- **Requirement:** Motorola's Linux-based smart phones needed a data management component to be fast, flexible, reliable and

*“Berkeley DB was 20 times faster than other databases. It has the operational speed of a main memory database, the startup and shut down speed of a disk-resident database, and does not have the overhead of a client-server inter-process communication.”*

Ray van Tassle, Senior Staff Engineer, Motorola



Motorola A780



Motorola E680

- **Solution:** Motorola selected best-of-breed open source components including Sleepycat's Berkeley DB, Montavista Linux and Trolltech Qt. Over 3 million of these devices have been shipped.

Makers of **Berkeley DB**

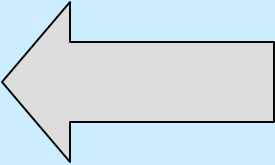
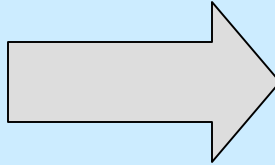
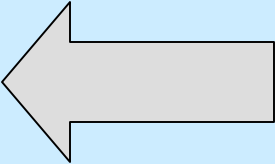
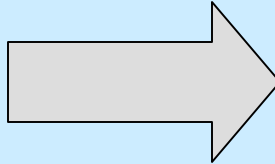


# Thank you!

We acknowledge and appreciate Bill Weinberg's article  
*The "B" in LAMP: How Berkeley DB Helps LAMP Shine*  
in the Enterprise OpenSource Journal, Jan/Feb 2006.  
[www.linuxpundit.com/cv/docs/EOSJ\\_the\\_B\\_in\\_LAMP.pdf](http://www.linuxpundit.com/cv/docs/EOSJ_the_B_in_LAMP.pdf)

David Schachter  
ds@sleepycat.com

# Support Options

| Open Source Licensees                                                                                                                    | Commercial Licensees                                                                                                                                                                                                                        |                                                                                       |
|------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|
| <p>Production Support</p> <ul style="list-style-type: none"> <li>• Production Systems</li> <li>• Guaranteed Response, 8 hours</li> </ul> | <p>Annual Maintenance</p> <ul style="list-style-type: none"> <li>• Upgrades and Updates</li> </ul> <p>Production Support</p> <ul style="list-style-type: none"> <li>• Production Systems</li> <li>• Guaranteed Response, 8 hours</li> </ul> |                                                                                       |
|                                                        | <p>Premium Support</p> <ul style="list-style-type: none"> <li>• Guaranteed Response, 4 hours</li> <li>• Access to Named Engineer</li> <li>• Help with Design and Tuning</li> </ul>                                                          |   |
|                                                       | <p>Critical Systems Support</p> <ul style="list-style-type: none"> <li>• Guaranteed Response, 1 hour</li> <li>• Unlimited points of contacts</li> <li>• 7 * 24 * 365 coverage</li> </ul>                                                    |  |

# When to use Berkeley DB

- Proven, fast and cost-effective data management component for application developers
- Unlike client-server and relational databases, *Berkeley DB* is:
  - Blindingly fast: no client/server, SQL, or data translation overhead
  - Zero install, zero maintenance, no DBA needed
  - Small footprint (350-650 kB)
  - Reliable
- You should use Berkeley DB when...
  - Performance and reliability are key
  - Client/server is not appropriate
  - Relational and/or SQL add unnecessary complexity
  - Cost of ownership is important

# Common Usage Scenarios

1. Database of Record (Google)
2. Cache front-end to database of record (Amazon)
3. Data on the Move (Motorola Linux Phones)
4. Internal (embedded) database (Cisco)
  1. Embedded on a device
5. Transactional database (vs corporate database)
  1. Lookup and pointer tables
  2. Indexing
  3. Session/state management
  4. Metadata management
  5. etc

# Summary

- In today's highly complex, interdependent, global networks, Berkeley DB is an ideal fit for data management to support
  - Realtime processing,
  - High speed access to data, on disk and in-memory
  - Increased distribution, scalability and performance
  - Efficiency and reduction
- Ideal for embedded, transactional data management
- Sleepycat support is enterprise strength and highly responsive
  
- Berkeley DB's unique business model ensures
  - Total Cost of Ownership is significantly reduced
  - Highest possible code quality

# Enterprise Database Functionality

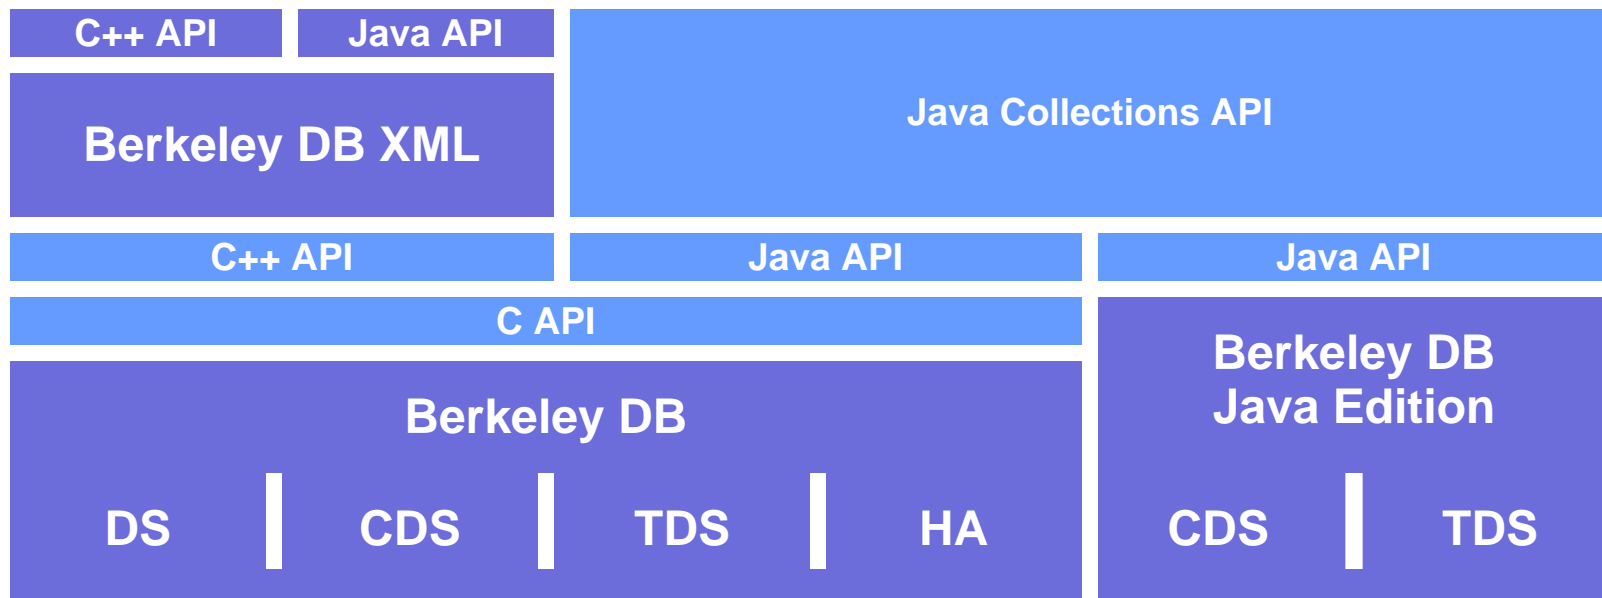
- General purpose, transactional, recoverable database engine
- High availability option for scalability
- Programmatic configuration and management or scriptable command-line administrative tools
- Security – data management completely hidden in your application
- Bytes to many terabytes
- Highly configurable feature set
  - Software toolkits for locking, logging, memory pool management

# The Berkeley DB Product Range

- **Berkeley DB**
  - General purpose transactional database engine
  - C, C++, Java, Tcl, Perl, Python, PHP, Ruby APIs
  - Java via JNI (Java Native Interface)
- **Berkeley DB XML**
  - XML layer on top of Berkeley DB
  - C++, Java, Tcl, Perl, Python, PHP APIs
  - XQuery
- **Berkeley DB Java Edition**
  - General purpose transactional database engine
  - Written in 100% Pure Java

# Product Relationships

Language bindings for C/C++, Java, Python, Perl, PHP, Ruby, Tcl, Eiffel, ...



More than 50 operating systems:  
UNIX, Linux, MacOS X, Windows, VxWorks, QNX, ...

# Layered Configurations

- Data Store (DS)
  - Key/data pair storage
  - No locking of any kind
- Concurrent Data Store (DS plus...)
  - Single writer at a time, multiple readers
  - Database- or environment-level locking
- Transactional Data Store (CDS plus...)
  - Many readers, many writers
  - Page- or record-level locking
  - ACID compliance (Durability, etc.)
- High Availability (TDS plus...)
  - Single-master, multiple client replication

# Key Features of TDS and HA

- Transactions
  - ACID semantics (Atomic, Consistent, Isolated, Durable)
  - Recover from system or application failure
  - Group multiple operations into a single logical operation
  - High-performance commits (lazy write, group commit)
- High Availability/Replication
  - Used for hot-standby failover and for load balancing in distributed applications
  - Single master (writer), multiple clients (readers)
  - Keep distributed copies of the database consistent

# SQL and Berkeley DB - Similarities

| SQL               | Berkeley DB                           |
|-------------------|---------------------------------------|
| Database          | Environment                           |
| Table             | Database                              |
| Row               | Key/Data pair (record)                |
| Boolean AND       | Join (DB->join)                       |
| Secondary Index   | Secondary Index or Secondary Database |
| ACID Transactions | ACID Transactions                     |
| Recovery          | Recovery                              |

# SQL and Berkeley DB - Differences

| SQL                                  | Berkeley DB                                                 |
|--------------------------------------|-------------------------------------------------------------|
| SQL Data Manipulation Language (DML) | Berkeley DB API                                             |
| Query planner and executor           | Application controlled behavior                             |
| Boolean OR                           | Multiple cursors                                            |
| Inner and Outer Joins                | Application controlled database lookups                     |
| Aggregates                           | Application controlled                                      |
| ORDER BY                             | Cursor over appropriate index or application sort in memory |

