



Table of Contents

I Reference Guide 1

1 Introduction 3

Mapping the Terrain: Theory and Practice 4

What Is Berkeley DB? 8

What Is Berkeley DB Not? 11

Do You Need Berkeley DB? 14

What Other Services Does Berkeley DB
Provide? 15

What Does the Berkeley DB Distribution
Include? 16

Where Does Berkeley DB Run? 16

Sleepycat Software's Berkeley DB Products 16

2 Getting Started: A Simple Tutorial 19

Key/Data Pairs 20

Object Handles 20

Error Returns 20

Opening a Database 21

Adding Elements to a Database 22

Retrieving Elements from a Database 25

Removing Elements from a Database 26

Closing a Database 28

3 Access Method Configuration 31

Selecting an Access Method 32

Logical Record Numbers 34

Selecting a Page Size 35

Selecting a Cache Size 36

Selecting a Byte Order 37

Non-Local Memory Allocation 37

Btree Comparison	37
Btree Prefix Comparison	39
Minimum Keys Per Page	40
Retrieving Btree Records by Number	40
Page Fill Factor	41
Specifying a Database Hash	41
Hash Table Size	42
Managing Record-Based Databases	42
Selecting a Queue Extent Size	43
Flat-Text Backing Files	43
Logically Renumbering Records	42
4 Access Method Operations	47
Opening a Database	47
Opening Multiple Databases in a Single File	47
Upgrading Databases	49
Retrieving Records	50
Storing Records	50
Deleting Records	50
Flushing the Database Cache	51
Database Statistics	51
Closing a Database	51
Database Cursors	52
Retrieving Records with a Cursor	52
Storing Records with a Cursor	53
Deleting Records with a Cursor	54
Duplicating a Cursor	54
Logical Join	54
Data Item Count	58
Closing a Cursor	58
Cursor Stability	58
Database Verification and Salvage	59

viii Contents

5 Access Method Wrap-Up	61
Retrieved Key/Data Permanence for C/C++	61
Database Limits	62
Disk Space Requirements	62
Partial Record Storage and Retrieval	65
Error Support	67
6 Berkeley DB Architecture	70
Programming Model	72
Programmatic APIs	72
Scripting Languages	74
Supporting Utilities	74
7 Berkeley DB Environment	77
Creating a Database Environment	78
File Naming	79
Filename Resolution in Berkeley DB	80
Security	82
Shared Memory Regions	83
Remote Filesystems	84
Opening Databases Within the Environment	84
Error Support	85
8 Berkeley DB Concurrent Data Store Applications	87
9 Berkeley DB Transactional Data Store Applications	91
Why Transactions?	92
Terminology	92
Application Structure	93
Opening the Environment	95
Opening the Databases	98
Recoverability and Deadlock Avoidance	100
Atomicity	103
Repeatable Reads	106
Transactional Cursors	107
Nested Transactions	110
Environment Infrastructure	111

Deadlock Detection	111
Performing Checkpoints	113
Database and Log File Archival Procedures	115
Log File Removal	117
Recovery Procedures	119
Recovery and Filesystem Operations	120
Berkeley DB Recoverability	121
Transaction Throughput	123
10 XA Resource Manager	127
Configuring Berkeley DB with the Tuxedo System	128
Frequently Asked Questions	129
11 Programmer Notes	131
Error Returns to Applications	132
Environmental Variables	134
Building Multithreaded Applications	134
Berkeley DB Handles	135
Name Spaces	136
Library Version Information	138
Compatibility with Historic UNIX Interfaces	138
Recovery Implementation	139
Application-Specific Logging and Recovery	139
Run-Time Configuration	144
12 The Locking Subsystem	147
Page Locks	149
Standard Lock Modes	150
Locking Without Transactions	151
Locking with Transactions: Two-Phase Locking	151
Access Method Locking Conventions	152
Berkeley DB Concurrent Data Store Locking Conventions	154
Deadlocks and Deadlock Avoidance	155
Configuring Locking	156
Locking and Non-Berkeley DB Applications	158

x Contents

13	The Logging Subsystem	161
	Configuring Logging	162
	Log File Limits	162
14	The Memory Pool Subsystem	165
	Configuring the Memory Pool	166
15	The Transaction Subsystem	167
	Transaction Limits	168
	Configuring Transactions	169
	Transactions and Non-Berkeley DB Applications	170
16	RPC/Client Server	173
	Client Program	174
	Server Program	175
17	Java API	177
	Compatibility	178
	Java Programming Notes	179
	Java FAQ	180
18	Tcl API	181
	Loading Berkeley DB with Tcl	181
	Using Berkeley DB with Tcl	182
	Tcl API Programming Notes	183
	Tcl Error Handling	183
	Tcl FAQ	185
18	Dumping and Reloading Databases	187
	The db_dump and db_load Utilities	187
	Dump Output Formats	188
	Loading Text into Databases	189
20	Debugging Applications	191
	Compile-time Configuration	192
	Run-time Error Information	193
	Reviewing Berkeley DB Log Files	193
	Common Errors	197

21 Building Berkeley DB for UNIX and QNX Systems 201

- Building for UNIX 201
- Configuring Berkeley DB 202
- Changing Compile or Load Options 205
- Installing Berkeley DB 206
- Dynamic Shared Libraries 207
- Running the Test Suite Under UNIX 208
- Architecture-Independent FAQs 209
- AIX 212
- FreeBSD 213
- HP-UX 214
- IRIX 215
- Linux 216
- OSF/1 216
- SCO 216
- Solaris 216
- SunOS 218
- Ultrix 218

22 Building Berkeley DB for Win32 Platforms 219

- Building for Win32 219
- Running the Test Suite Under Windows 222
- Windows Notes 223
- Windows FAQ 224

23 Building Berkeley DB for VxWorks Systems 225

- VxWorks Notes 227
- VxWorks FAQ 228

24 Upgrading Berkeley DB Applications 231

- Upgrading Berkeley DB Installations 231

xii Contents

25 Test Suite 235

Running the Test Suite 235

Test Suite FAQ 237

26 Distribution 239

27 Additional References 241

Technical Papers on Berkeley DB 241

Background on Berkeley DB Features 242

Database Systems Theory 242

II API Manual 243

28 C API 245

29 C++ API 373

30 Java API 489

31 Tcl API 591

32 Supporting Utilities 617