An In-depth Technical Guide to the ndbm Database

Author: BenchChem Technical Support Team. **Date**: December 2025

Compound of Interest	
Compound Name:	NDBM
Cat. No.:	B12393030 Get Quote

For Researchers, Scientists, and Drug Development Professionals

This guide provides a comprehensive technical overview of the **ndbm** (New Database Manager) library, a key-value store database system. It is designed for an audience with a technical background, such as researchers, scientists, and drug development professionals, who may encounter or consider using **ndbm** for managing experimental data, metadata, or other forms of key-addressable information.

Core Concepts of **ndbm**

ndbm is a high-performance, file-based database system that stores data as key-value pairs. It is an evolution of the original dbm (Database Manager) and offers enhancements such as the ability to have multiple databases open simultaneously. At its core, **ndbm** is a library of functions that an application can use to manipulate a database.

The fundamental principle of **ndbm** is its use of a hashing algorithm to quickly locate data on disk. When a key-value pair is stored, a hash of the key is calculated, which determines the storage location of the corresponding value. This allows for very fast data retrieval, typically in one or two disk accesses, making it suitable for applications requiring rapid lookups of relatively static data.[1]

An **ndbm** database is physically stored as two separate files:

- .dir file: This file acts as a directory or index, containing a bitmap of hash values.[1]

- .pag file: This file contains the actual data, the key-value pairs themselves.[1]

This two-file structure separates the index from the data, which can contribute to efficient data retrieval operations.

Data Presentation: Quantitative Analysis

The following tables summarize key quantitative aspects of **ndbm** and related dbm-family databases.

Table 1: Key and Value Size Limitations

Database Implementation	Typical Key Size Limit	Typical Value Size Limit	Notes
Original dbm	~512 bytes (total for key-value pair)	~512 bytes (total for key-value pair)	Considered obsolete.
ndbm	Varies by implementation	Varies by implementation	Often cited with a combined key-value size limit around 1008 to 4096 bytes.[2]
gdbm (GNU dbm)	No limit	No limit	Offers an ndbm compatibility mode that removes the size limitations.
Berkeley DB	No practical limit	No practical limit	Also provides an ndbm emulation layer with enhanced capabilities.

Table 2: Performance Benchmarks of dbm-like Databases

The following data is based on a benchmark test storing 1,000,000 records with 8-byte keys and 8-byte values.

Database	Write Time (seconds)	Read Time (seconds)	File Size (KB)
NDBM 5.1	8.07	7.79	814,457
GDBM 1.8.3	14.01	5.36	82,788
Berkeley DB 4.4.20	9.62	5.62	40,956
SDBM 1.0.2	11.32	N/A*	606,720
QDBM 1.8.74	1.89	1.58	55,257

*Read time for SDBM was not available due to database corruption during the test.

Source: Adapted from Huihoo, Benchmark Test of DBM Brothers.

Internal Mechanics: The Hashing Algorithm

ndbm and its derivatives often employ a variant of the sdbm (Sedgewick's Dynamic Bit Manipulation) algorithm for hashing keys. This algorithm is known for its good distribution of hash values, which helps in minimizing collisions and ensuring efficient data retrieval.[3][4][5]

The core of the sdbm algorithm is an iterative process that can be represented by the following pseudo-code:

This simple yet effective algorithm contributes to the fast lookup times characteristic of **ndbm** databases.

Experimental Protocols: Core **ndbm** Operations

The following section details the standard procedures for interacting with an **ndbm** database using its C-style API. These protocols are fundamental for storing and retrieving experimental data.

Data Structures

The primary data structure for interacting with **ndbm** is the datum, which is used to represent both keys and values. It is typically defined as:

- dptr: A pointer to the data.

- dsize: The size of the data in bytes.

Key Experimental Steps

- Opening a Database Connection:

  - Protocol: Use the dbm_open() function.

  - Synopsis: DBM *dbm_open(const char *file, int flags, mode_t mode);

  - Description: This function opens a connection to the database specified by file. The flags argument determines the mode of operation (e.g., O_RDWR for read/write, O_CREAT to create the database if it doesn't exist). The mode argument specifies the file permissions if the database is created.[6][7]

  - Returns: A pointer to a DBM object on success, or NULL on failure.

- Storing Data:

  - Protocol: Use the dbm_store() function.

  - Synopsis: int dbm_store(DBM *db, datum key, datum content, int store_mode);

  - Description: This function stores a key-value pair in the database. The store_mode can be DBM_INSERT (insert only if the key does not exist) or DBM_REPLACE (overwrite the value if the key exists).[6][8]

  - Returns: 0 on success, a non-zero value on failure.

- Retrieving Data:

  - Protocol: Use the dbm_fetch() function.

  - Synopsis: datum dbm_fetch(DBM *db, datum key);

  - Description: This function retrieves the value associated with a given key.[6][8]

- Returns: A datum structure containing the value. The dptr field will be NULL if the key is not found.

- Deleting Data:

  - Protocol: Use the dbm_delete() function.

  - Synopsis: int dbm_delete(DBM *db, datum key);

  - Description: This function removes a key-value pair from the database.[9]

  - Returns: 0 on success, a non-zero value on failure.

- Closing the Database Connection:

  - Protocol: Use the dbm_close() function.

  - Synopsis: void dbm_close(DBM *db);

  - Description: This function closes the database connection and ensures that all changes are written to disk.[6][7]

Mandatory Visualizations
ndbm High-Level Data Flow

The following diagram illustrates the basic workflow of storing and retrieving data using the **ndbm** library.
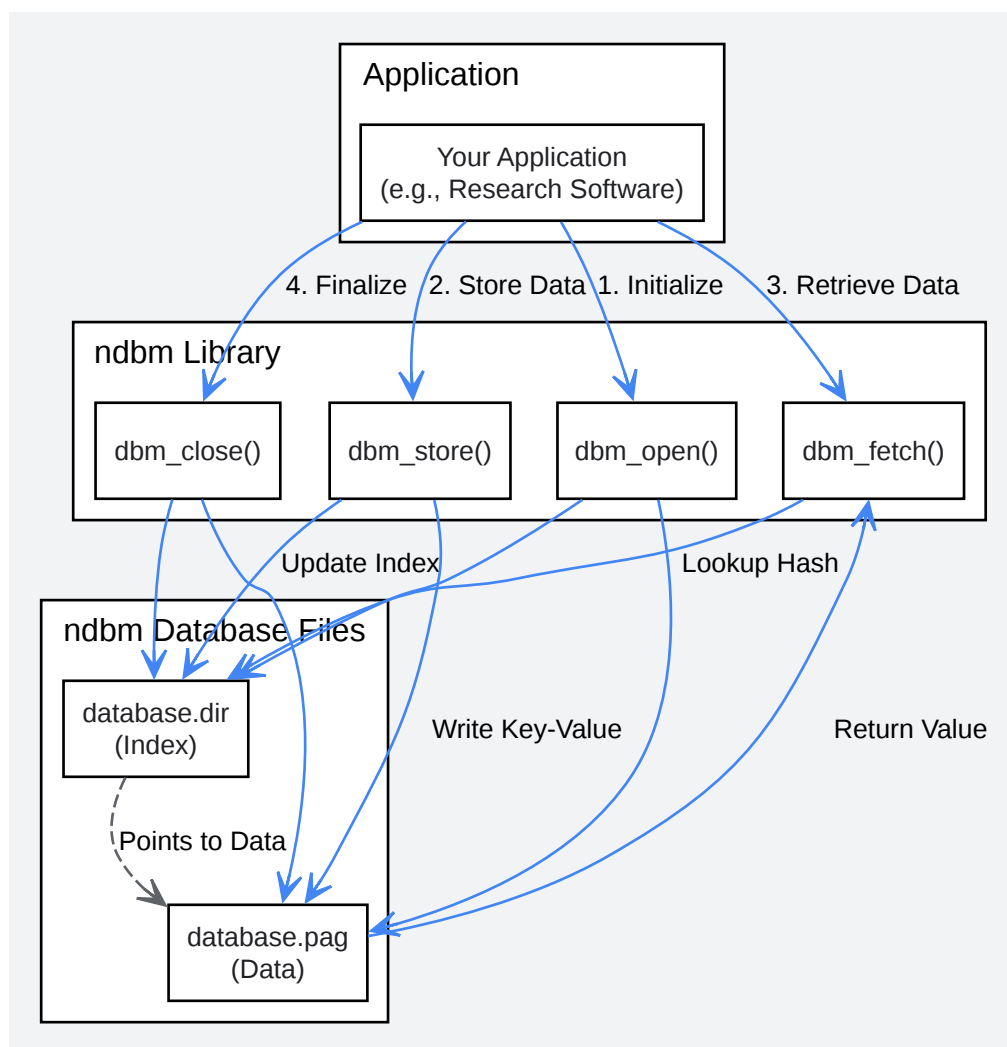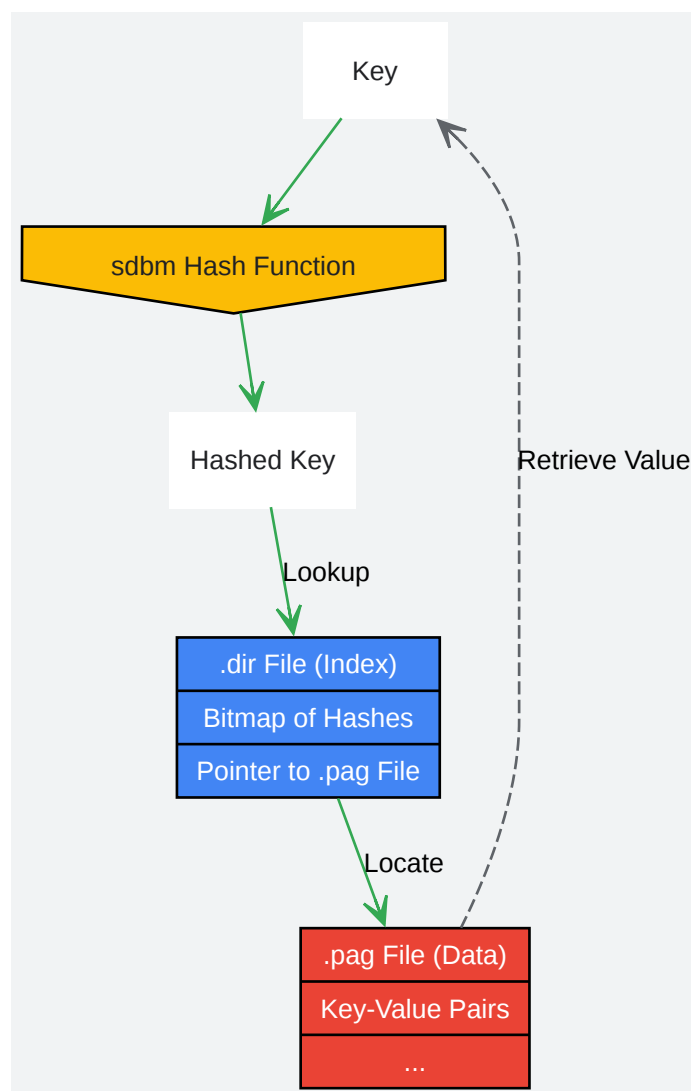
A high-level overview of the **ndbm** data storage and retrieval workflow.

ndbm Internal Hashing and Lookup

This diagram provides a conceptual view of how **ndbm** uses hashing to locate data within its file structure.

Tech Support

Conceptual diagram of the **ndbm** hashing and data lookup process.

Conclusion

The **ndbm** database provides a simple, robust, and high-performance solution for key-value data storage. While it has limitations in terms of data size in its original form, its API has been emulated and extended by more modern libraries like gdbm and Berkeley DB, which overcome these constraints. For researchers and scientists who need a fast, local, and straightforward database for managing structured data, **ndbm** and its successors remain a viable and relevant technology. Its simple API and file-based nature make it easy to integrate into various scientific computing workflows.

 Tech Support

> **Need Custom Synthesis?**
>
> BenchChem offers custom synthesis for rare earth carbides and specific isotopiclabeling.
>
> Email: info@benchchem.com or Request Quote Online.

References

- 1. IBM Documentation [ibm.com]

- 2. NDBM_File - Tied access to ndbm files - Perldoc Browser [perldoc.perl.org]

- 3. matlab.algorithmexamples.com [matlab.algorithmexamples.com]

- 4. cse.yorku.ca [cse.yorku.ca]

- 5. sdbm [doc.riot-os.org]

- 6. The NDBM library [infolab.stanford.edu]

- 7. ndbm (GDBM manual) [gnu.org.ua]

- 8. dbm/ndbm [docs.oracle.com]

- 9. NDBM Tutorial [franz.com]

- To cite this document: BenchChem. [An In-depth Technical Guide to the ndbm Database]. BenchChem, [2025]. [Online PDF]. Available at: [https://www.benchchem.com/product/b12393030#what-is-ndbm-database]

Disclaimer & Data Validity:

The information provided in this document is for Research Use Only (RUO) and is strictly not intended for diagnostic or therapeutic procedures. While BenchChem strives to provide accurate protocols, we make no warranties, express or implied, regarding the fitness of this product for every specific experimental setup.

Technical Support: The protocols provided are for reference purposes. Unsure if this reagent suits your experiment? [Contact our Ph.D. Support Team for a compatibility check]

Need Industrial/Bulk Grade?   Request Custom Synthesis Quote

BenchChem

Our mission is to be the trusted global source of essential and advanced chemicals, empowering scientists and researchers to drive progress in science and industry.

Contact

Address: 3281 E Guasti Rd

Ontario, CA 91761, United States

Phone: (601) 213-4426

Email: info@benchchem.com