

An In-depth Technical Guide to NDBM Data Structures

Author: BenchChem Technical Support Team. **Date:** April 2026

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 New Database Manager (**NDBM**), a foundational key-value pair data structure. While largely superseded by more modern libraries, understanding **NDBM**'s core principles offers valuable insight into the evolution of database technologies and the fundamental concepts of on-disk hash tables. This document is intended for researchers and professionals who require a deep understanding of data storage mechanisms for managing scientific and experimental data.

Core Concepts of NDBM

NDBM is a library of subroutines that provides a simple and efficient interface for managing key-value databases stored on disk.[1] It was developed as an enhancement to the original DBM library, offering improvements such as the ability to have multiple databases open simultaneously.[2] The primary function of **NDBM** is to store and retrieve arbitrary data based on a unique key, making it an early example of a NoSQL data store.

At its core, **NDBM** implements an on-disk hash table. This structure allows for fast data retrieval, typically in one or two file system accesses, without the overhead of a full relational database system.[1] Data is organized into key-value pairs, where both the key and the value

can be arbitrary binary data. This flexibility is particularly useful for storing heterogeneous scientific data.

NDBM On-Disk Structure

An **NDBM** database consists of two separate files:

- The Directory File (.dir): This file acts as an index or a bitmap for the data file.^[1] It contains a directory that maps hash values of keys to locations within the page file.
- The Page File (.pag): This file stores the actual key-value pairs.^[1]

This two-file structure separates the index from the data, which can improve performance by allowing the potentially smaller directory file to be more easily cached in memory. It's important to note that modern emulations of **NDBM**, such as those provided by Berkeley DB, may use a single file with a .db extension.^{[2][3][4]}

The Hashing Mechanism: Extendible Hashing

NDBM utilizes a form of extendible hashing to dynamically manage the on-disk hash table.^[5] This technique allows the hash table to grow as more data is added, avoiding the need for costly full-table reorganizations.

The core components of the extendible hashing mechanism in **NDBM** are:

- Directory: An in-memory array of pointers to data buckets on disk. The size of the directory is a power of 2.
- Global Depth (d): An integer that determines the size of the directory (2^d). The first 'd' bits of a key's hash value are used as an index into the directory.
- Buckets (Pages): Fixed-size blocks in the .pag file that store the key-value pairs.
- Local Depth (d'): An integer stored with each bucket, indicating the number of bits of the hash value shared by all keys in that bucket.

Data Insertion and Splitting Logic:

- A key is hashed, and the first d (global depth) bits of the hash are used to find an entry in the directory.
- The directory entry points to a bucket in the .pag file.
- The key-value pair is inserted into the bucket.
- If the bucket is full:
 - If the bucket's local depth d' is less than the directory's global depth d , the bucket is split, and its contents are redistributed between the old and a new bucket based on the $d'+1$ -th bit of the keys' hashes. The directory pointers are updated to point to the correct buckets.
 - If the bucket's local depth d' is equal to the global depth d , the directory itself must be doubled in size. The global depth d is incremented, and the bucket is then split.

This dynamic resizing of the directory and splitting of buckets allows **NDBM** to handle growing datasets efficiently.

Experimental Protocols: Algorithmic Procedures

While specific experimental protocols from scientific literature using the original **NDBM** are scarce due to its age, we can detail the algorithmic protocols for the primary **NDBM** operations. These can be considered the "experimental" procedures for interacting with the data structure.

Protocol for Storing a Key-Value Pair

- Initialization: Open the database using `dbm_open()`, specifying the file path and access flags (e.g., read-write, create if not exists). This returns a database handle.
- Data Preparation: Prepare the key and content in datum structures. A datum is a simple struct containing a pointer to the data (`dptr`) and its size (`dsize`).
- Hashing: The **NDBM** library internally computes a hash of the key.
- Directory Lookup: The first d (global depth) bits of the hash are used to index into the in-memory directory.

- **Bucket Retrieval:** The directory entry provides the address of the data bucket in the .pag file. This bucket is read from disk.
- **Insertion and Overflow Check:** The new key-value pair is added to the bucket. If the bucket exceeds its capacity, the bucket splitting and/or directory doubling procedure (as described in Section 3) is initiated.
- **Write to Disk:** The modified bucket(s) and, if necessary, the directory file are written back to disk.
- **Return Status:** The dbm_store() function returns a status indicating success, failure, or if an attempt was made to insert a key that already exists with the DBM_INSERT flag.[6][7]

Protocol for Retrieving a Value by Key

- **Initialization:** Open the database using dbm_open().
- **Key Preparation:** Prepare the key to be fetched in a datum structure.
- **Hashing and Directory Lookup:** The key is hashed, and the first d bits are used to find the corresponding directory entry.
- **Bucket Retrieval:** The directory entry's pointer is used to locate and read the appropriate bucket from the .pag file.
- **Key Search:** The keys within the bucket are linearly scanned to find a match.
- **Data Return:** If a matching key is found, a datum structure containing a pointer to the corresponding value and its size is returned. If the key is not found, the dptr field of the returned datum will be NULL.[6]


Quantitative Data Summary

Due to the age of **NDBM**, direct and recent performance benchmarks are not readily available. However, we can summarize the known characteristics and limitations in a structured format for comparison.

Feature	NDBM	GDBM (GNU DBM)	Berkeley DB
Primary Use	Simple key-value storage	A more feature-rich replacement for NDBM	High-performance, transactional embedded database
File Structure	Two files (.dir, .pag)	Can emulate the two-file structure but is a single file internally	Typically a single file
Concurrency	Generally not safe for concurrent writers	Provides file locking for safe concurrent access	Full transactional support with fine-grained locking
Key/Value Size Limits	Limited (e.g., 1018 to 4096 bytes)[2]	No inherent limits	No inherent limits
API	dbm_open, dbm_store, dbm_fetch, etc.	Native API and NDBM compatibility API	Rich API with support for transactions, cursors, etc.
In-memory Caching	Basic, relies on OS file caching	Internal bucket cache	Sophisticated in-memory cache management
Crash Recovery	Not guaranteed	Offers some crash tolerance	Full ACID-compliant crash recovery

Visualizations

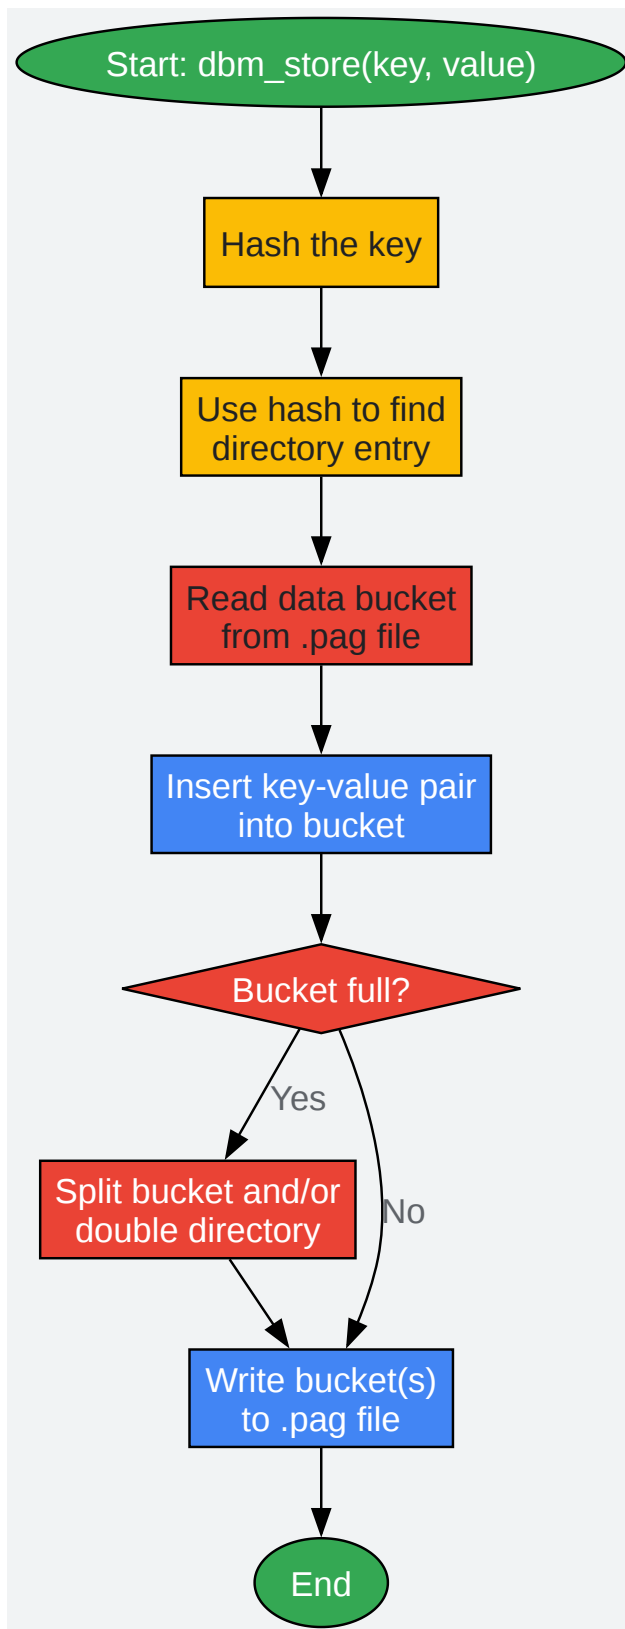
NDBM File Structure



[Click to download full resolution via product page](#)

Caption: The two-file architecture of an **NDBM** database.

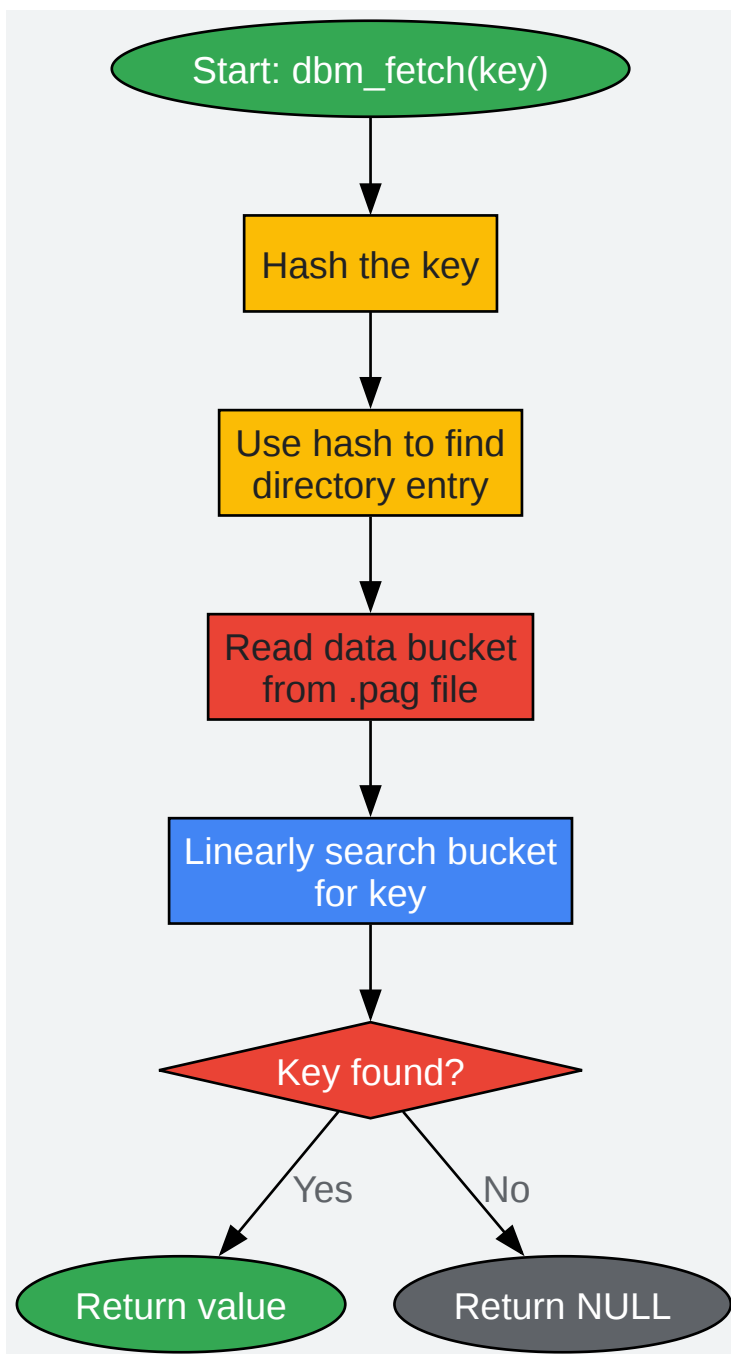
NDBM Data Storage Workflow



[Click to download full resolution via product page](#)

Caption: Logical workflow for storing data in an **NDBM** database.

NDBM Data Retrieval Workflow



[Click to download full resolution via product page](#)

Caption: Logical workflow for retrieving data from an **NDBM** database.

Conclusion

NDBM represents a significant step in the evolution of simple, efficient on-disk data storage. For researchers and scientists, understanding its architecture provides a solid foundation for appreciating the trade-offs involved in modern data management systems. While direct use of the original **NDBM** is uncommon today, its principles of key-value storage and extendible hashing are still relevant in the design of high-performance databases. When choosing a data storage solution for research applications, the principles embodied by **NDBM**—simplicity, direct key-based access, and predictable performance—remain valuable considerations. For new projects, however, modern libraries such as Berkeley DB or GDBM are recommended as they provide **NDBM**-compatible interfaces with enhanced features, performance, and robustness.

Need Custom Synthesis?

BenchChem offers custom synthesis for rare earth carbides and specific isotopic labeling.

Email: info@benchchem.com or [Request Quote Online](#).

References

- [1. IBM Documentation \[ibm.com\]](#)
- [2. Unix Incompatibility Notes: DBM Hash Libraries \[unixpapa.com\]](#)
- [3. dbm/ndbm \[docs.oracle.com\]](#)
- [4. Berkeley DB: dbm/ndbm \[ucalgary.ca\]](#)
- [5. DBM \(computing\) - Wikipedia \[en.wikipedia.org\]](#)
- [6. ndbm\(3\) - OpenBSD manual pages \[man.openbsd.org\]](#)
- [7. The NDBM library \[infolab.stanford.edu\]](#)
- To cite this document: BenchChem. [An In-depth Technical Guide to NDBM Data Structures]. BenchChem, [2026]. [Online PDF]. Available at: [\[https://www.benchchem.com/product/b12393030/docs#an-in-depth-technical-guide-to-ndbm-data-structures\]](https://www.benchchem.com/product/b12393030/docs#an-in-depth-technical-guide-to-ndbm-data-structures)

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?

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

[Contact our Ph.D. Support Team for a compatibility check](#)