An In-depth Technical Guide to the ndbm File Format

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 who rely on robust data storage, understanding the underlying architecture of database systems is paramount. This guide provides a detailed technical exploration of the **ndbm** (new database manager) file format, a foundational key-value store that has influenced numerous subsequent database technologies.

Core Concepts of **ndbm**

The **ndbm** library, a successor to the original dbm, provides a simple yet efficient method for storing and retrieving key-value pairs. It was a standard feature in early Unix-like operating systems, including 4.3BSD.[1][2][3] Unlike modern database systems that often use a single file, **ndbm** utilizes a two-file structure to manage data: a directory file (.dir) and a data page file (.pag).[1] This design is predicated on a hashing algorithm to provide fast access to data, typically in one or two file system accesses.[1]

The fundamental unit of data in **ndbm** is a datum, a structure containing a pointer to the data (dptr) and its size (dsize). This allows for the storage of arbitrary binary data as both keys and values.[3][4]

The On-Disk File Format: A Deep Dive

The **ndbm** file format is intrinsically tied to its implementation of extendible hashing. This dynamic hashing scheme allows the database to grow gracefully as more data is added, without requiring a complete reorganization of the file.

Tech Support

The Directory File (.dir)

The .dir file acts as the directory for the extendible hash table. It does not contain the actual key-value data but rather pointers to the data pages in the .pag file. The core of the .dir file is a hash table, which is an array of page indices.

A simplified view of the .dir file's logical structure reveals its role as an index. It contains a bitmap that is used to keep track of used pages in the .pag file.[1]

The Page File (.pag)

The .pag file is a collection of fixed-size pages, where each page stores one or more key-value pairs. The structure of a page is designed for efficient storage and retrieval. Key-value pairs that hash to the same logical bucket are stored on the same page.

When a page becomes full, a split occurs. A new page is allocated in the .pag file, and some of the key-value pairs from the full page are moved to the new page. The .dir file is then updated to reflect this change, potentially doubling in size to accommodate a more granular hash function.

The Hashing Mechanism

The efficiency of **ndbm** hinges on its hashing algorithm, which determines the initial placement of keys within the .pag file. While the original **ndbm** source code from 4.3BSD would provide the definitive algorithm, a widely cited and influential hashing algorithm comes from sdbm, a public-domain reimplementation of **ndbm**.

The sdbm hash function is as follows:

This simple iterative function was found to provide good distribution and scrambling of bits, which is crucial for minimizing collisions and ensuring efficient data retrieval.

Collision Resolution: In **ndbm**, collisions at the hash function level are handled by storing multiple key-value pairs that hash to the same bucket on the same data page. When a page overflows due to too many collisions, the page is split, and the directory is updated. This is a form of open addressing with bucket-level collision resolution.

Tech Support

Key Operations and Experimental Protocols

The **ndbm** interface provides a set of functions for interacting with the database. Understanding these is key to appreciating its operational workflow.

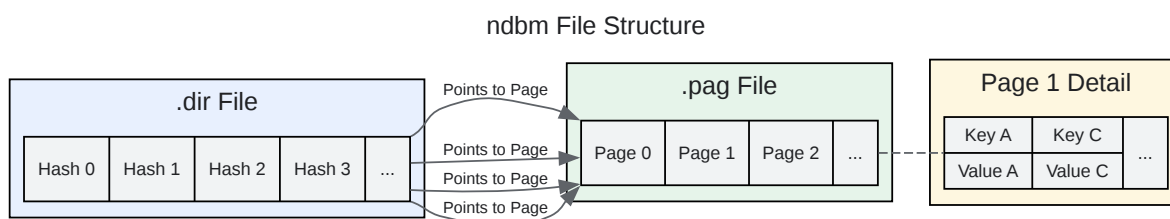| Function	Description
dbm_open()	Opens or creates a database, returning a handle to the two-file structure.[4][5]
dbm_store()	Stores a key-value pair in the database.[4][5]
dbm_fetch()	Retrieves the value associated with a given key. [4][5]
dbm_delete()	Removes a key-value pair from the database.
dbm_firstkey()	Retrieves the first key in the database for iteration.
dbm_nextkey()	Retrieves the next key in the database for iteration.
dbm_close()	Closes the database files.[4][5]

Experimental Protocol for a dbm_store Operation:

- Key Hashing: The key is passed through the **ndbm** hash function to generate a hash value.

- Directory Lookup: The hash value is used to calculate an index into the directory table in the .dir file.

- Page Identification: The entry in the directory table provides the page number within the .pag file where the key-value pair should be stored.

- Page Retrieval: The corresponding page is read from the .pag file into memory.

- Key-Value Insertion: The new key-value pair is appended to the data on the page.

- Overflow Check: If the insertion causes the page to exceed its capacity, a page split is triggered.

- Page Split (if necessary): a. A new page is allocated in the .pag file. b. The key-value pairs on the original page are redistributed between the original and the new page based on a refined hash. c. The directory in the .dir file is updated to point to the new page. This may involve doubling the size of the directory.

- Page Write: The modified page(s) are written back to the .pag file.

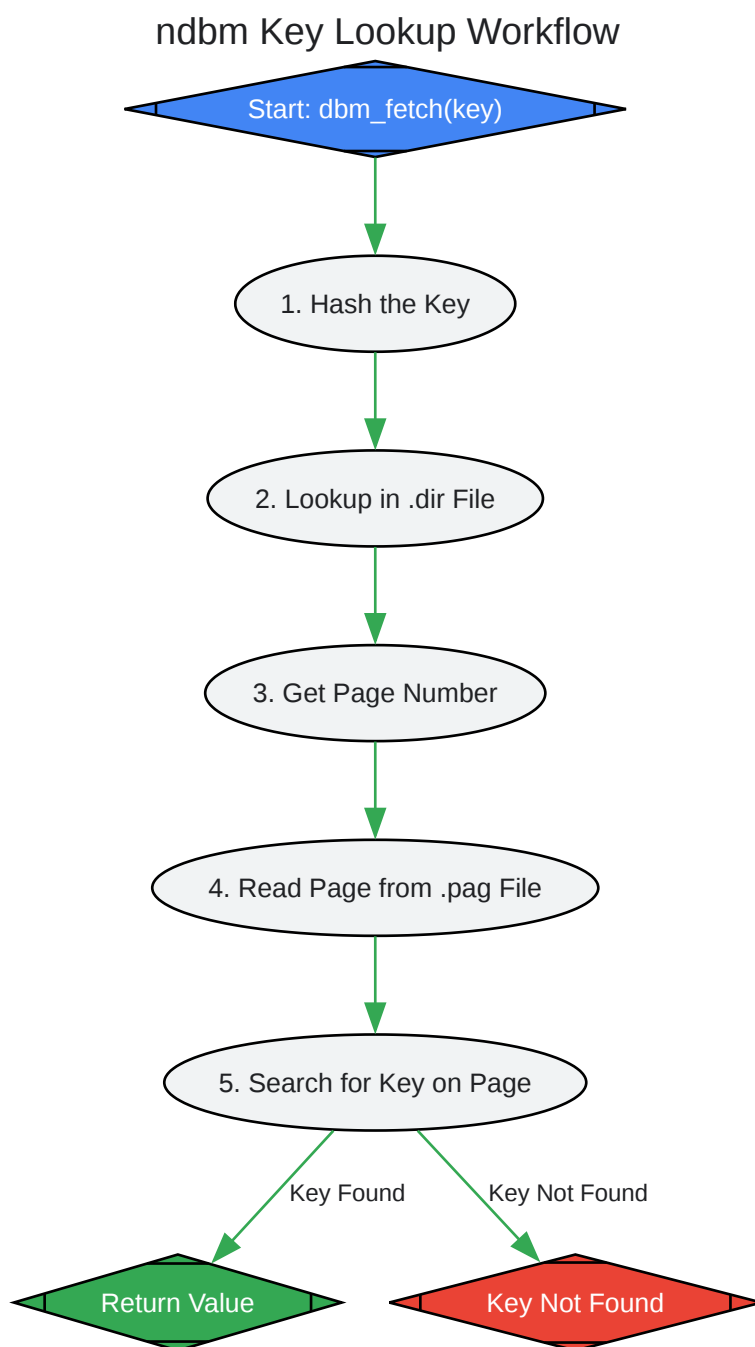# Visualizing the **ndbm** Architecture and Workflow

To better illustrate the concepts described, the following diagrams are provided in the DOT language for Graphviz.

ndbm File Structure

.dir File					Points to Page	.pag File				Page 1 Detail		
Hash 0	Hash 1	Hash 2	Hash 3	...	Points to Page / Points to Page / Points to Page	Page 0	Page 1	Page 2	...	Key A / Value A	Key C / Value C	...

Click to download full resolution via product page

Logical relationship between .dir and .pag files.

The diagram above illustrates the fundamental architecture of an **ndbm** database. The .dir file contains a hash table that maps hash values to page numbers in the .pag file. Multiple directory entries can point to the same page. The .pag file itself is a collection of pages, each containing the actual key-value pairs.

ndbm Key Lookup Workflow

Start: dbm_fetch(key)

1. Hash the Key

2. Lookup in .dir File

3. Get Page Number

4. Read Page from .pag File

5. Search for Key on Page

Key Found → Return Value

Key Not Found → Key Not Found

Workflow for a key lookup operation (dbm_fetch).

This workflow diagram shows the steps involved in retrieving a value for a given key. The process begins with hashing the key, followed by a lookup in the .dir file to identify the correct data page in the .pag file. The relevant page is then read and searched for the key.

Modern Implementations and Compatibility

While the original **ndbm** is now largely of historical and academic interest, its API has been preserved in modern database libraries such as GNU gdbm and Oracle Berkeley DB.[6][7] These libraries provide an **ndbm** compatibility interface, allowing older software to be compiled and run on modern systems. However, it is crucial to note that the underlying on-disk file formats of these modern implementations are different from the original **ndbm** format and are generally not compatible with each other.[7]

Feature	Original ndbm	GNU gdbm (in ndbm mode)	Berkeley DB (in ndbm mode)
File Structure	.dir and .pag files	.dir and .pag files (may be hard links)	Single .db file
On-Disk Format	Specific to the original implementation	gdbm's own format	Berkeley DB's own format
Data Size Limits	Key/value pair size limits (e.g., 1024 bytes)[2]	No inherent limits	No inherent limits
Concurrency	No built-in locking	Optional locking	Full transactional support

Conclusion

The **ndbm** file format represents a significant step in the evolution of key-value database systems. Its two-file, extendible hashing design provided a robust and efficient solution for data storage in early Unix environments. While it has been superseded by more advanced database technologies, its core concepts and API have demonstrated remarkable longevity, influencing and being preserved in modern database libraries. For professionals in data-intensive fields, understanding the principles of **ndbm** offers valuable insights into the foundational techniques of data management.

 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. grokipedia.com [grokipedia.com]

- 2. Introduction to dbm | KOSHIGOE.Write(something) [koshigoe.github.io]

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

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

- 5. RonDB - World's fastest Key-Value Store [rondb.com]

- 6. DBM (computing) - Wikipedia [en.wikipedia.org]

- 7. Unix Incompatibility Notes: DBM Hash Libraries [unixpapa.com]

- To cite this document: BenchChem. [An In-depth Technical Guide to the ndbm File Format]. BenchChem, [2025]. [Online PDF]. Available at: [https://www.benchchem.com/product/b12393030#ndbm-file-format-explained]

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