

# history of ndbm and dbm libraries

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

## Compound of Interest

Compound Name: NDBM

Cat. No.: B12393030

[Get Quote](#)

An In-depth Technical Guide to the Core of dbm and **ndbm**

## Introduction

In the history of Unix-like operating systems, the need for a simple, efficient, and persistent key-value storage system led to the development of the Database Manager (dbm). This library and its successors became foundational components for various applications requiring fast data retrieval without the complexity of a full-fledged relational database. This document provides a technical overview of the original dbm library, its direct successor **ndbm**, and the subsequent evolution of this database family, tailored for an audience with a technical background.

## Historical Development and Evolution

The dbm family of libraries represents one of the earliest forms of NoSQL databases, providing a straightforward associative array (key-value) storage mechanism on disk.<sup>[1]</sup>

### The Genesis: dbm

The original dbm library was written by Ken Thompson at AT&T Bell Labs and first appeared in Version 7 (V7) Unix in 1979.<sup>[1][2][3]</sup> It was designed as a simple, disk-based hash table, offering fast access to data records via string keys.<sup>[1][3]</sup> A dbm database consisted of two files:

- **.dir** file: A directory file containing the hash table indices.
- **.pag** file: A data file containing the actual key-value pairs.<sup>[2][4][5]</sup>

This initial implementation had significant limitations: it only allowed one database to be open per process and was not designed for concurrent access by multiple processes.<sup>[2][4]</sup> The pointers to data returned by the library were stored in static memory, meaning they could be overwritten by subsequent calls, requiring developers to immediately copy the results.<sup>[2]</sup>

## The Successor: ndbm

To address the limitations of the original, the New Database Manager (**ndbm**) was developed and introduced with 4.3BSD Unix in 1986.<sup>[2][3]</sup> While maintaining compatibility with the core concepts of dbm, **ndbm** introduced several crucial enhancements:

- Multiple Open Databases: It modified the API to allow a single process to have multiple databases open simultaneously.<sup>[1][2]</sup>
- File Locking: It incorporated file locking mechanisms to enable safe, concurrent read access.<sup>[2]</sup> However, write access was still typically limited to a single process at a time.<sup>[6]</sup>
- Standardization: The **ndbm** API was later standardized in POSIX and the X/Open Portability Guide (XPG4).<sup>[2]</sup>

Despite these improvements, **ndbm** retained the two-file structure (.dir and .pag) and had its own limitations on key and data size.<sup>[4][5]</sup>

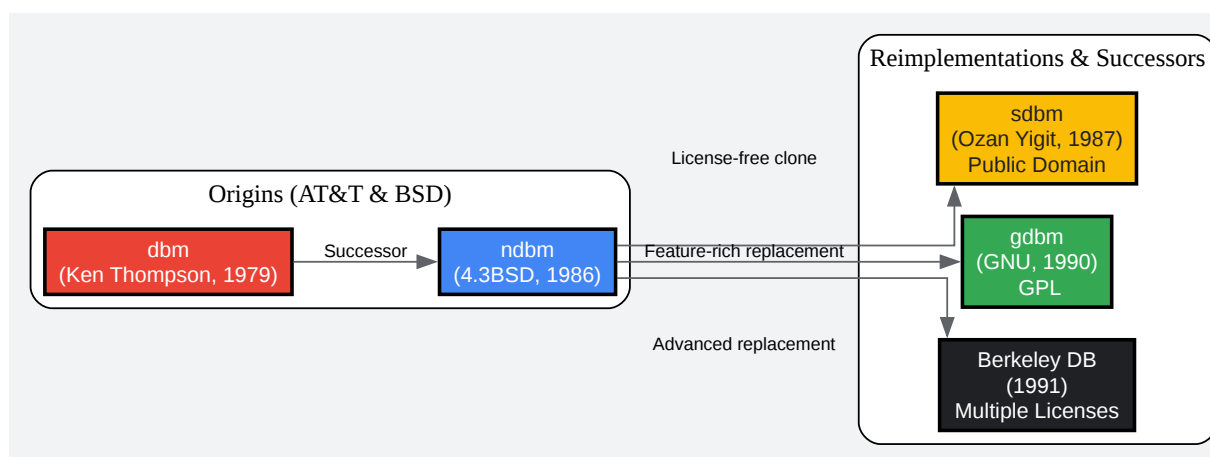
## The Family Expands

The influence of dbm and **ndbm** led to a variety of reimplementations, each aiming to improve upon the original formula by removing limitations or changing licensing.

- sdbm: Written in 1987 by Ozan Yigit, sdbm ("small dbm") was a public-domain clone of **ndbm**, created to avoid the AT&T license restrictions.<sup>[1][3]</sup>
- gdbm: The GNU Database Manager (gdbm) was released in 1990 by the Free Software Foundation.<sup>[2][3]</sup> It implemented the **ndbm** interface but also added features like crash tolerance, no limits on key/value size, and a different, single-file database format.<sup>[1][3][7]</sup>
- Berkeley DB (BDB): Originating in 1991 to replace the license-encumbered BSD **ndbm**, Berkeley DB became the most advanced successor.<sup>[1]</sup> It offered significant enhancements,

including transactions, journaling for crash recovery, and support for multiple access methods beyond hashing, all while providing a compatibility interface for **ndbm**.<sup>[4]</sup>

The evolutionary path of these libraries shows a clear progression towards greater stability, fewer limitations, and more flexible licensing.



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

Evolution of the dbm library family.

## Core Technical Details and Methodology

The fundamental principle behind dbm and its variants is the use of a hash table stored on disk. This allows for very fast data retrieval based on a key.

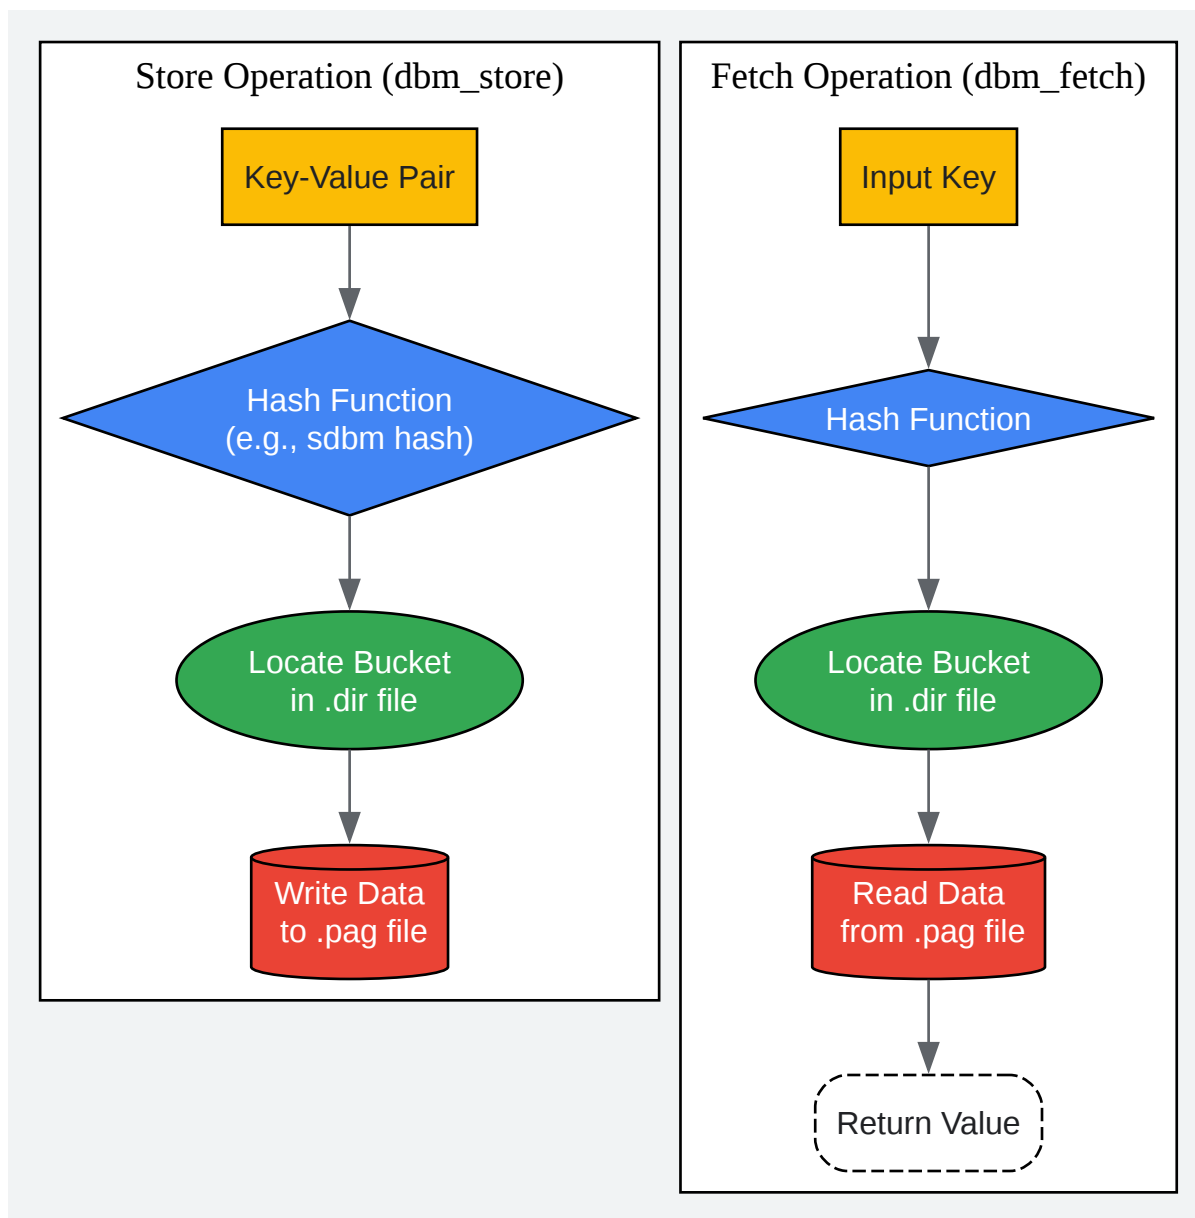
## Data Structure and Hashing

The core methodology involves a hashing function to map a given key to a specific location ("bucket") within the database files.<sup>[1][3]</sup>

- Hashing: When a key-value pair is to be stored, the library applies a hash function to the key, which computes an integer value.

- **Bucket Location:** This hash value is used to determine the bucket where the key-value pair should reside.
- **Storage:** The key and its associated data are written into the appropriate block in the .pag file. An index pointing to this data is stored in the .dir file.
- **Collision Handling:** Since different keys can produce the same hash value (a "collision"), the library must handle this. The dbm implementation uses a form of extendible hashing.<sup>[1]</sup> If a bucket becomes full, it is split, and the hash directory is updated to accommodate the growing data.

This approach ensures that, on average, retrieving any value requires only one or two disk accesses, making it significantly faster than sequentially scanning a flat file.<sup>[5]</sup>



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

Simplified workflow for dbm store and fetch operations.

## Quantitative Data and Specifications

The various dbm implementations can be compared by their technical limitations and features. While formal benchmarks of these legacy systems are scarce, their documented specifications provide a clear comparison.

Feature	dbm (Original v7)	ndbm (4.3BSD)	gdbm (GNU)	Berkeley DB (Modern)
Release Date	1979[1][2][3]	1986[2][3]	1990[2][3]	1991 (initial)[1]
File Structure	Two files (.dir, .pag)[2][4]	Two files (.dir, .pag)[4][5]	Single file[6]	Single file[4]
Key/Value Size Limit	~512 bytes (total per entry)[2][3]	~1024 - 4096 bytes (implementation dependent)[3][4]	No limit[3]	No practical limit
Concurrent Access	1 process max[2][4]	Multiple readers, single writer[2]	Multiple readers, single writer[6]	Full transactional (multiple writers)
Crash Recovery	None	None	Yes (crash tolerance)[1][7]	Yes (journaling, transactions)
API Header		[2]	[2]	[8]

## Conclusion

The dbm library and its direct descendant **ndbm** were pioneering technologies in the Unix ecosystem. They established a simple yet powerful paradigm for on-disk key-value storage that influenced countless applications and spawned a family of more advanced database engines. While modern applications often rely on more sophisticated systems like Berkeley DB, GDBM, or other NoSQL databases, the foundational concepts of hashing for fast, direct data access introduced by dbm remain a cornerstone of database design. Understanding their history and technical underpinnings provides valuable insight into the evolution of data storage technology.

### Need Custom Synthesis?

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

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

## References

- 1. DBM (computing) - Wikipedia [en.wikipedia.org]
- 2. grokipedia.com [grokipedia.com]
- 3. Introduction to dbm | KOSHIGOE.Write(something) [koshigoe.github.io]
- 4. Unix Incompatibility Notes: DBM Hash Libraries [unixpapa.com]
- 5. IBM Documentation [ibm.com]
- 6. gdbm [edoras.sdsu.edu]
- 7. dbm & Interfaces to Unix & databases — Python 3.14.2 documentation [docs.python.org]
- 8. dbm/ndbm [docs.oracle.com]
- To cite this document: BenchChem. [history of ndbm and dbm libraries]. BenchChem, [2025]. [Online PDF]. Available at: [https://www.benchchem.com/product/b12393030#history-of-ndbm-and-dbm-libraries]

---

### 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](mailto:info@benchchem.com)