

StorHouse/RM Features

Customers typically use StorHouse/RM as an active archive for scientific, or research, data or as a relational database extension system for aged data with mid-to-low-level access requirements.

StorHouse/RM:

- Executes on HP-UX, Solaris, and AIX platforms to accommodate different customer standards
- Scales to multiple petabytes with no performance degradation through a patented automatic partitioning, or segmentation, scheme
- Supports industry-standard interfaces such as ODBC, ESQL, and SQL
- Provides extremely fast load and index rates, including building 200K index entries per second per stream
- Utilizes unique range indexes on RAID for quicker access to relational data
- Supports FTP and SQL for bulk loading and bulk query unloading
- Provides an EXPLAIN facility for analyzing query operation
- Journals and restores database activities that affect metadata
- Optimizes performance by minimizing MAID spin-cycles, serializing tape accesses, and minimizing volume mounts
- Merges, or coalesces, small segments into larger ones to enhance query performance
- Supports a graphical user interface (GUI) for easy database administration

StorHouse/RM, FileTek's relational database management system (RDBMS) software component, works in conjunction with StorHouse/SM to administer the storage, access, and movement of relational data. Together, these integrated storage management and RDBMS technologies out-perform conventional data management systems that merely mix traditional database software with third-party commodity HSM tools and automated storage devices.

Executing on the same platform as StorHouse/SM, StorHouse/RM optimizes direct row-level access to data regardless of its location in the StorHouse/SM-managed virtual storage hierarchy. This feature eliminates the hours to days it takes traditional data warehouse systems to restore archive data from tape to disk for analysis and decision-making.

StorHouse/RM has a comprehensive architecture that provides:

- Structured Query Language (SQL), embedded SQL (ESQL), and ODBC interfaces
- High-speed data loaders and unloaders for moving information to and from StorHouse
- RDBMS gateways to Oracle, IBM DB2 Universal Database (UDB), Microsoft SQL Server, and Sybase
- An SAP-integrated interface to SAP Business Warehouse (BW) for offloading accessible aged data to StorHouse.

Figure 1 illustrates the StorHouse/RM architecture.

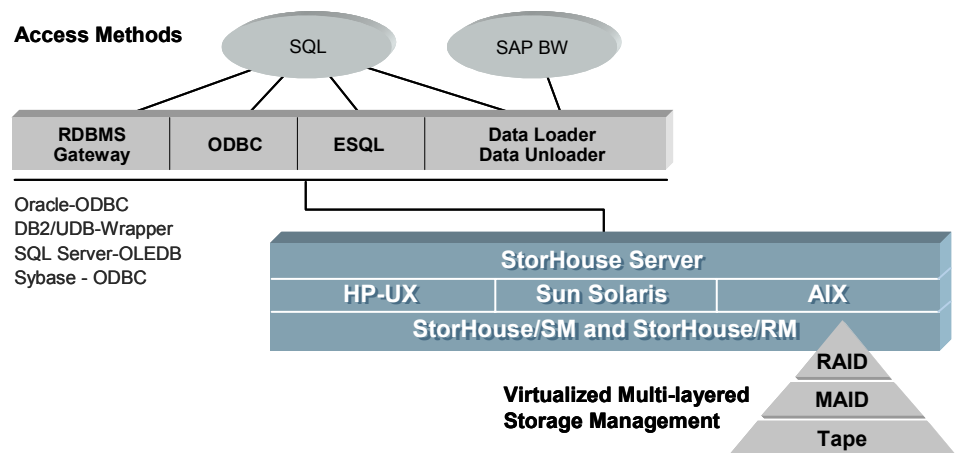


Figure 1: StorHouse/RM Architecture

All StorHouse/SM features are available with StorHouse/RM, including an automated server architecture; a virtual storage hierarchy; and scalable processors, storage devices, and user components.

StorHouse/RM Benefits

Enabling High-Availability Relational Archives

As a database extension system for critical aged data, StorHouse/RM enables highly available, cost-effective relational archives that satisfy enterprise-wide ILM and compliance requirements.

- With the SAP interface to StorHouse, StorHouse/RM is rapidly becoming the standard for relational archives. StorHouse/RM improves how SAP/BW users access aged data and the overall performance of their primary database systems.
- StorHouse/RM connects to Oracle, DB2/UDB, and Microsoft SQL Server databases through link interfaces. Once links are defined, users can direct queries to these databases and, when required, direct sub-queries to StorHouse to federate (UNION) the results. This capability provides improved performance and faster access to critical data.

In addition, StorHouse/RM:

- Transcends scalability limitations of traditional data warehouses
- Distributes data cost-effectively throughout an automatically managed virtual storage hierarchy
- Simplifies implementation by combining multiple technologies into one integrated product
- Provides the tools technologists need to build, store, and manage petabytes of data in a format easily accessible by a worldwide database community

Innovative Technology – Patented Segmentation

StorHouse/RM's innovative technology incorporates a patented segmentation scheme that enables unlimited, yet manageable, database structural growth. Segmentation makes it possible to use removable storage technologies or massive arrays of idle disks (MAID) for relational database storage.

In StorHouse/RM, applications produce relational tables through batch loads that create segments, or partitions, consisting of database user file components (tables, indexes, and large object data). As Figure 2 illustrates, each segment consists of one table data partition file, one index file for each value index and hash index, and one or more LOB subsegment files for a LOB column.

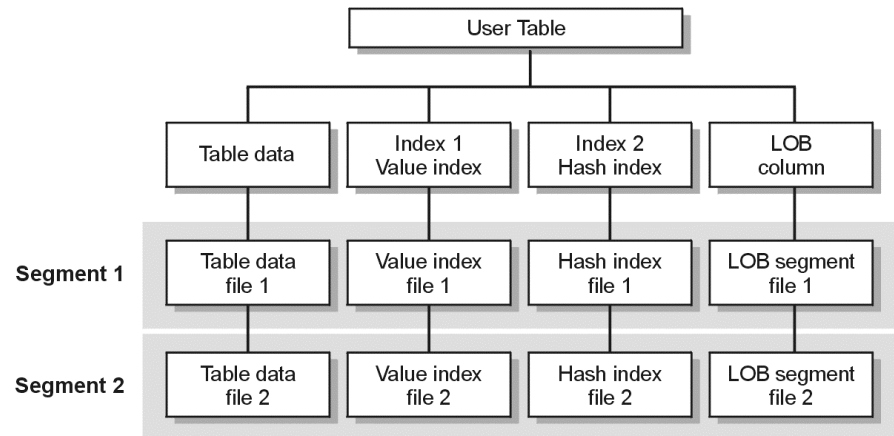


Figure 2: Segmentation

Segment components can be stored separately throughout the virtualized hierarchy to optimize the performance of each storage layer.

Segmentation provides numerous benefits.

- Segments easily map to StorHouse files, volume sets, and file sets, which StorHouse/SM automatically manages. This simplifies system and database administration tasks and keeps storage and administration costs down.
- StorHouse/RM maps index files to individual table data files for faster load and delete processing. There's no need to build one large global index for all segments and then issue time-consuming updates against that index.
- Users can run loads and queries against the same table in parallel.

Queries process more quickly because index files and the most recent table data files can reside on faster access media.

For More Information

To learn how StorHouse/RM can protect and manage your critical relational data, or for information about other FileTek products, contact a FileTek sales representative. You can also e-mail your questions to info@filetek.com.

©2006 FileTek, Inc. All rights reserved. StorHouse is a registered U.S. trademark of FileTek, Inc. Other trademarks included herein are the property of their respective owners. StorHouse is protected by the following U.S. patents: 4,864,572; 5,247,660; 5,727,197; and 6,049,804.

FileTek

Corporate Headquarters:

FileTek, Inc.

9400 Key West Avenue
Rockville, MD 20850
Phone: 301.251.0600
info@filetek.com
www.filetek.com

International Headquarters:

FileTek Ltd

1 Northumberland Ave.
London WC2N 5BW
Ph: +44 (0) 207.872.5583
intsales@filetek.com
www.filetek.com