

# Active Storage Hierarchy, Database Systems and Applications – Socratic Exegesis

**Felipe Cariño Jr.**

FileTek. Inc.  
360 N. Sepulveda Blvd, Suite 1080  
El Segundo, CA 90245  
**FCARINO@filetek.com**

**William O'Connell**

IBM Toronto Laboratory  
1150 Eglinton Avenue East  
Toronto, Ontario M3C 1H7  
**BOCONNEL@ca.ibm.com**

**John Burgess**

FileTek, Inc.  
9400 Key West Avenue  
Rockville, MD 20850  
**JGB@filetek.com**

**Joel Saltz**

Computer Science Department  
University of Maryland College Park, MD 20742  
Joint Appointment with  
Johns Hopkins University Medical  
**saltz@hyena.cs.umd.edu**

## Panel Abstract

This panel addresses a very important area that is often neglected or overlooked by database systems, database applications developers and data warehouse designers, namely storage. We propose to inform, discuss and debate the use of “*Active Storage Hierarchy*” in database systems and applications. By active storage hierarchy we mean a database system that uses all storage media (i.e. optical, tape, and disk) to store and retrieve data and not just disk. We will examine, discuss and debate how active storage compares and/or complements what is known in the database research community as “*Active Disks*” [RGF 98] and other emerging disk-centric storage paradigms. The presentations and analysis will span current real products, emerging technology to active (and visionary) research in several related areas, like storage technology, storage systems, federated databases and database system uses of storage.

## Panel Format

- Overview of Storage Technology, Current and Future Commercial Products.
- Overview of Database Research and Commercial Database Product Plans
- Overview and sample case studies of current and emerging applications that do and/or will in the future exploit cost-effective storage hierarchy.
- Discussion and debate on feasibility and future (visionary) use of storage systems in database applications.

---

*Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the VLDB copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Very Large Data Base Endowment. To copy otherwise, or to republish, requires a fee and/or special permission from the Endowment*

**Proceedings of the 25th VLDB Conference, Edinburgh, Scotland, 1999.**

## Why This Panel?

Figure 1 shows the storage hierarchy, capacity, performance and cost-benefit tradeoffs. The importance of storage cost and performance can be seen by projected future database applications sizes (see Table1). A panel discussion and debating point will be that even if visionary disk-based research and proposals [RGF 98] [UAS 99] [Gray 98] become reality and commercially viable, then projected disk storage costs may/will require the use of the complete active storage hierarchy (i.e. tape, optical, disk, and so on).

## Enabling and Emerging Technologies

There are several existing, evolving and emerging technologies (and products) that when combined and integrated will open the door to “new” uses and applications.

1. The first obvious one is storage itself. An exegesis of storage technologies (commercial or research work) will provide a foundation to discuss future storage uses. A key new storage evolution is the emergence and deployment of Storage Area Network (SAN) technology across enterprises. A major database challenge will be to integrate data sources and provide the same transparent access to data regardless where the data is stored. They can operate on SANs just like they operate on direct attached disks, so the question is will SANs enable new database architectures.
2. The second major panel item is how real commercial database systems currently exploit storage. StorHouse/Relational Manager (RM) [CB 99] approach (see Figure 2) is to store data on active storage and issue queries that retrieve data directly from data on all storage types. We will describe how the major database

vendors utilize what is commonly referred to as Hierarchical Storage Management (HSM).

3. The third major (large storage) technology enabler is the emergence and deployment of object/relational databases. As Table 1 shows these applications will require orders of magnitude more storage than current application utilize.
4. Federated databases systems (like DataJoiner) present another interesting opportunity or alternative to the StorHouse/RM direct approach. With respect to large applications that require cost-effective storage (again a discussion point), then federating data sources or moving inactive or large objects to lower-cost storage may be essential to deploy applications. The panel will discuss how federated data sources can exploit storage options.

## Panel Discussion & Debating Points

1. The panel will be informative in that we will get major storage vendors to come and/or share their product capabilities and future plans.
2. We plan to have the major database vendors represented or present their product plans in this area. Partisan and parochial interests and instincts of commercial vendors always makes for lively (informative) discussion.
3. The panel will also describe emerging federated, object/relational, legacy and atomic data applications that currently under-utilize storage options (another discussion and debate point).
4. Finally, leading academic researchers will provide their vision of the future.

## Panelists

Panel organizers plus industrial and researchers have been invited to participate.

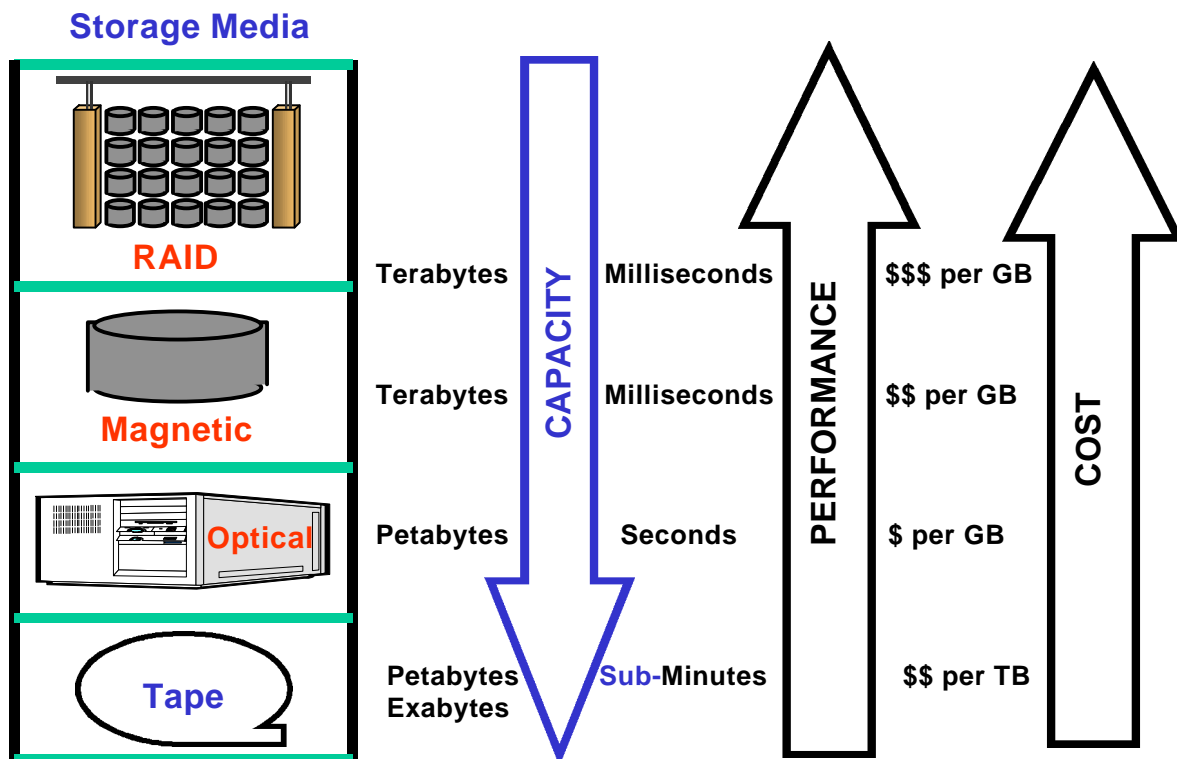
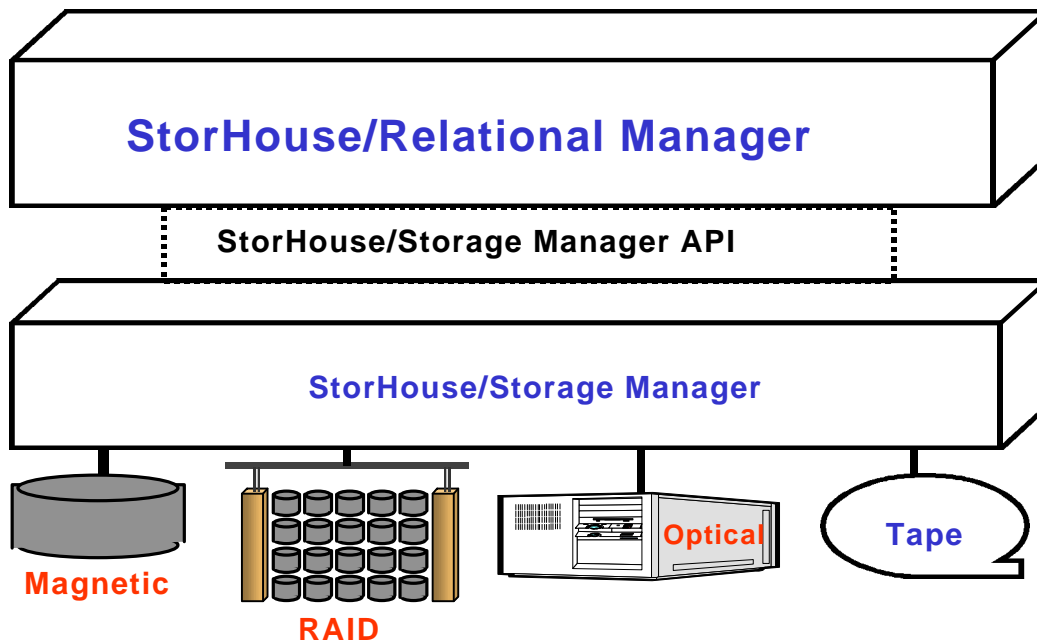


Figure 1: Storage hierarchy - Capacity, Performance and Cost-benefit Tradeoffs

Database Size	Application
Terabyte $10^{12}$ bytes	National Retail Point Sale Data
Petabyte $10^{15}$ bytes	Text and Images Product Descriptions
Exabyte $10^{18}$ bytes	National Medical Insurance Records
Zettabyte $10^{21}$ bytes	Spatial and Terrestrial Data
Yottabyte $10^{24}$ bytes	Large Video Archives

Table 1: Database Sizes for Diverse Applications



**Figure 2: StorHouse/Relational Manager – Direct Storage & Querying**

## References

[AUS98] Acharya, A., Uysal, M. and Saltz, J., "Structure and Performance of Decision Support Algorithms on Active Disks", UCSB TRCS98-28, October 1998.

[AUS99] Acharya, A., Uysal, M. and Saltz, J., "Active Disks: Programming Model, Algorithms and Evaluation", ASPLOS-VIII.

[BG 88] Bitton, D. and Gray, J., "Disk Shadowing", VLDB 1988.

[CB 99] Cariño, F. and Burgess, J., "Lost in Active Storage Space - New Frontier for IT Managers", To Appear in Intelligent Enterprise Magazine August 1999.

[Gray98] Jim Gray's 10-year best VLDB paper [DG 88] keynote address at VLDB '98.

[RGF98] Erik Riedel, Garth Gibson and Christos Faloutsos, "Active Storage For Large-Scale Data Mining and Multimedia, VLDB 1998.