



QMDS: A File System Metadata Service Supporting a Graph Data Model-Based Query Language

Sasha Ames

University of California, Santa Cruz & LLNL
sasha@soe.ucsc.edu

Maya B. Gokhale

Lawrence Livermore National
Laboratory

Carlos Maltzahn

University of California, Santa Cruz



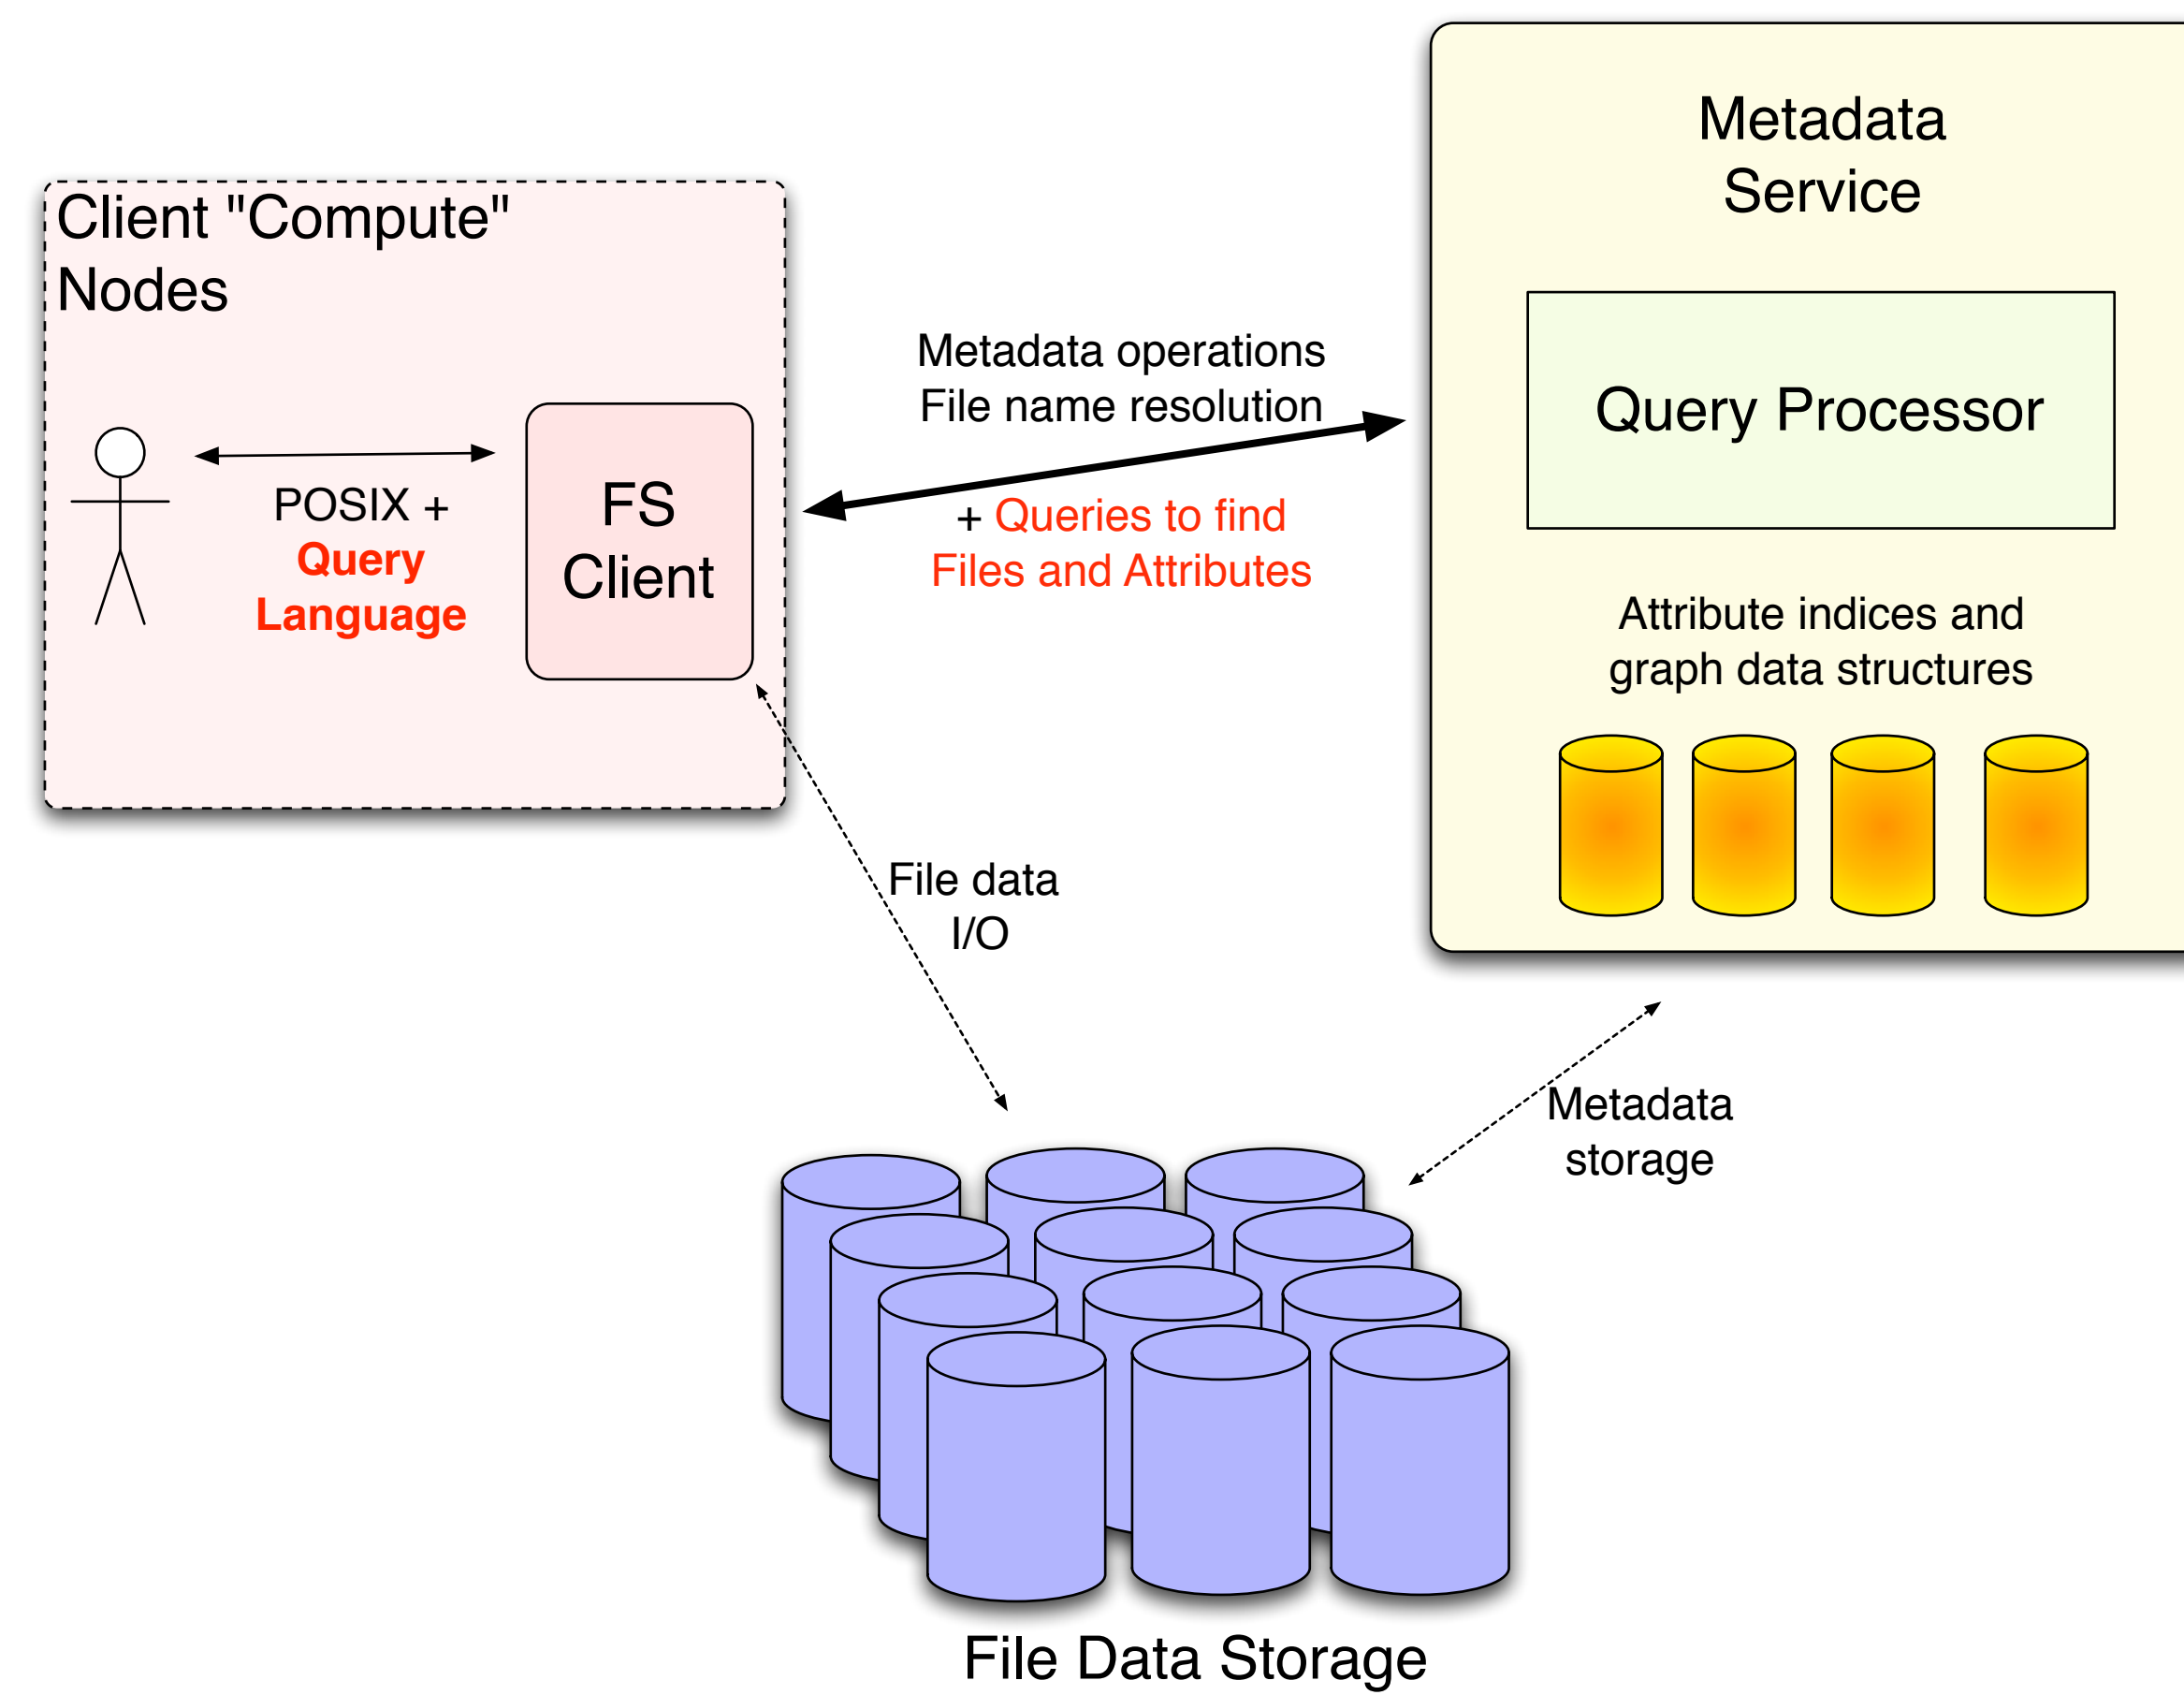
Introduction

File system metadata management has become a bottleneck for many data-intensive applications that rely on high-performance file systems. Part of the bottleneck is due to the limitations of an almost 50 year old interface standard with metadata abstractions that were designed at a time when high-end file systems managed less than 100MB.

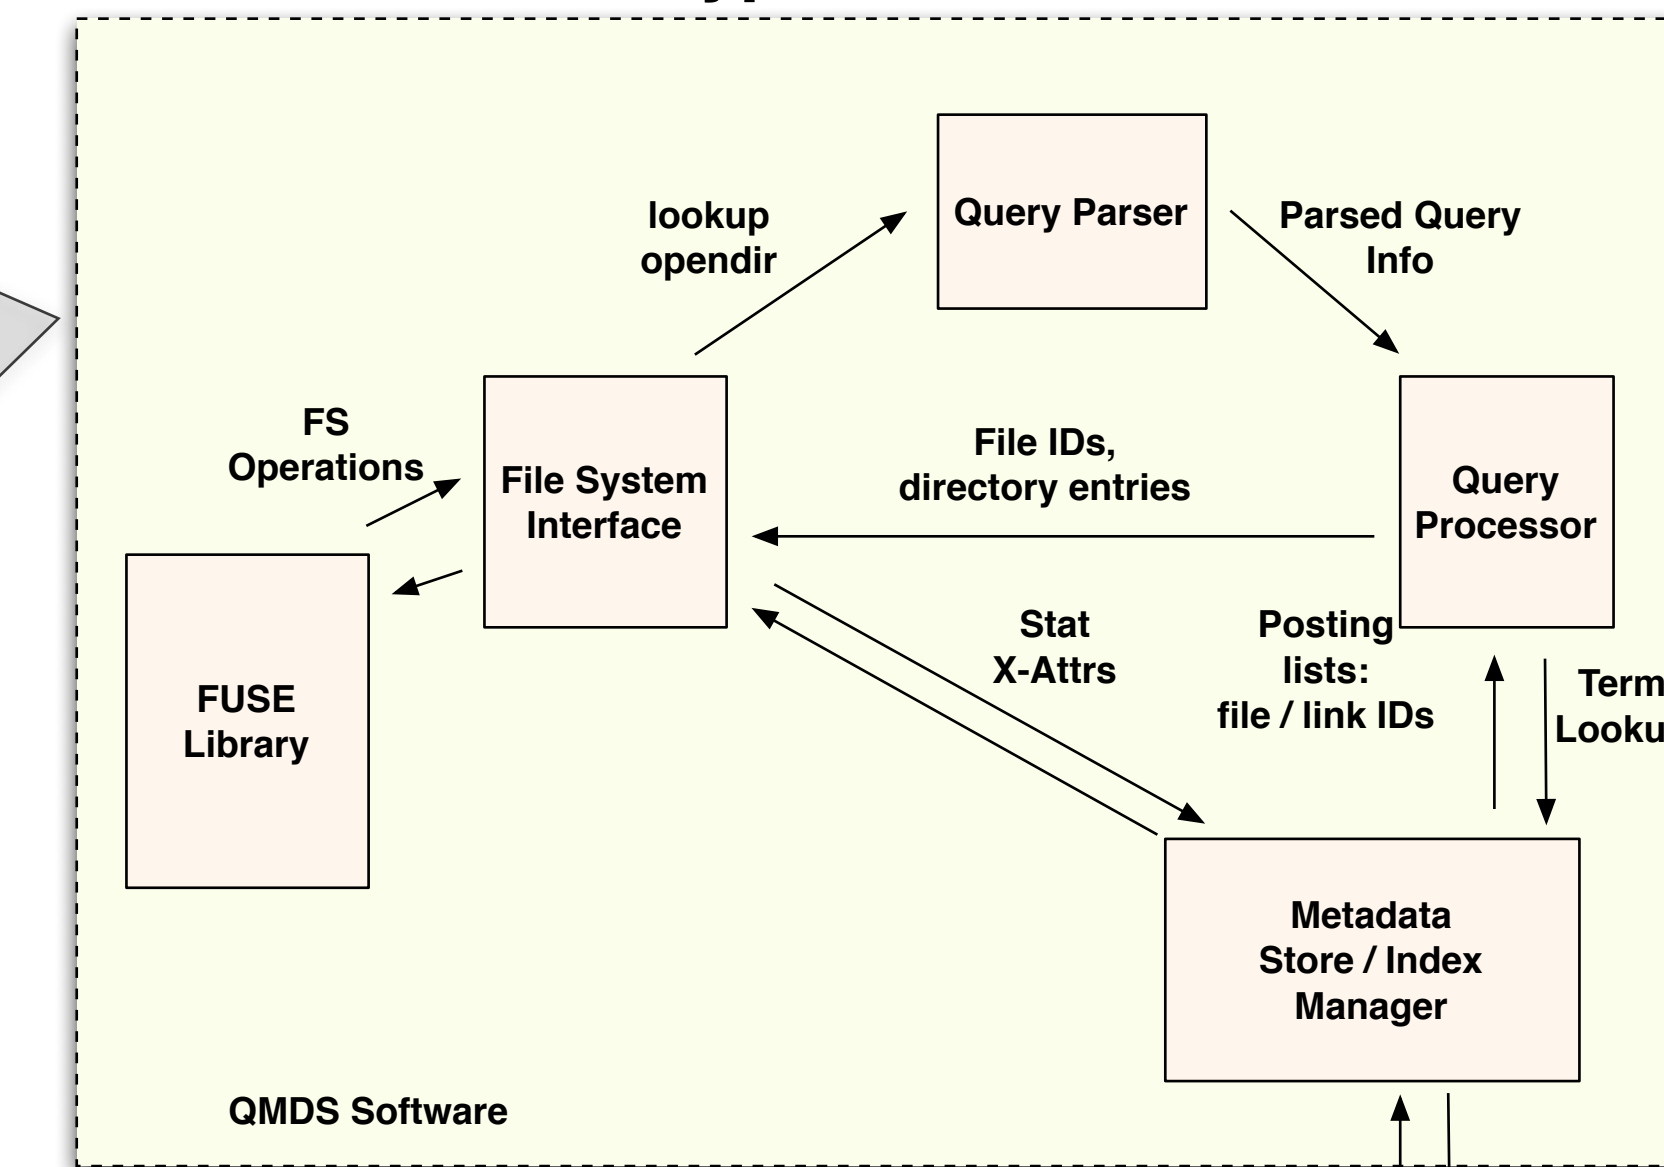
Today's high-performance file systems store 7 to 9 orders of magnitude more data, resulting in numbers of data items for which these metadata abstractions are inadequate, such as directory hierarchies unable to handle complex relationships among data.

Users of file systems have attempted to work around these inadequacies by moving application-specific metadata management to relational databases to make metadata searchable. Splitting file system metadata management into two separate systems introduces inefficiencies and systems management problems.

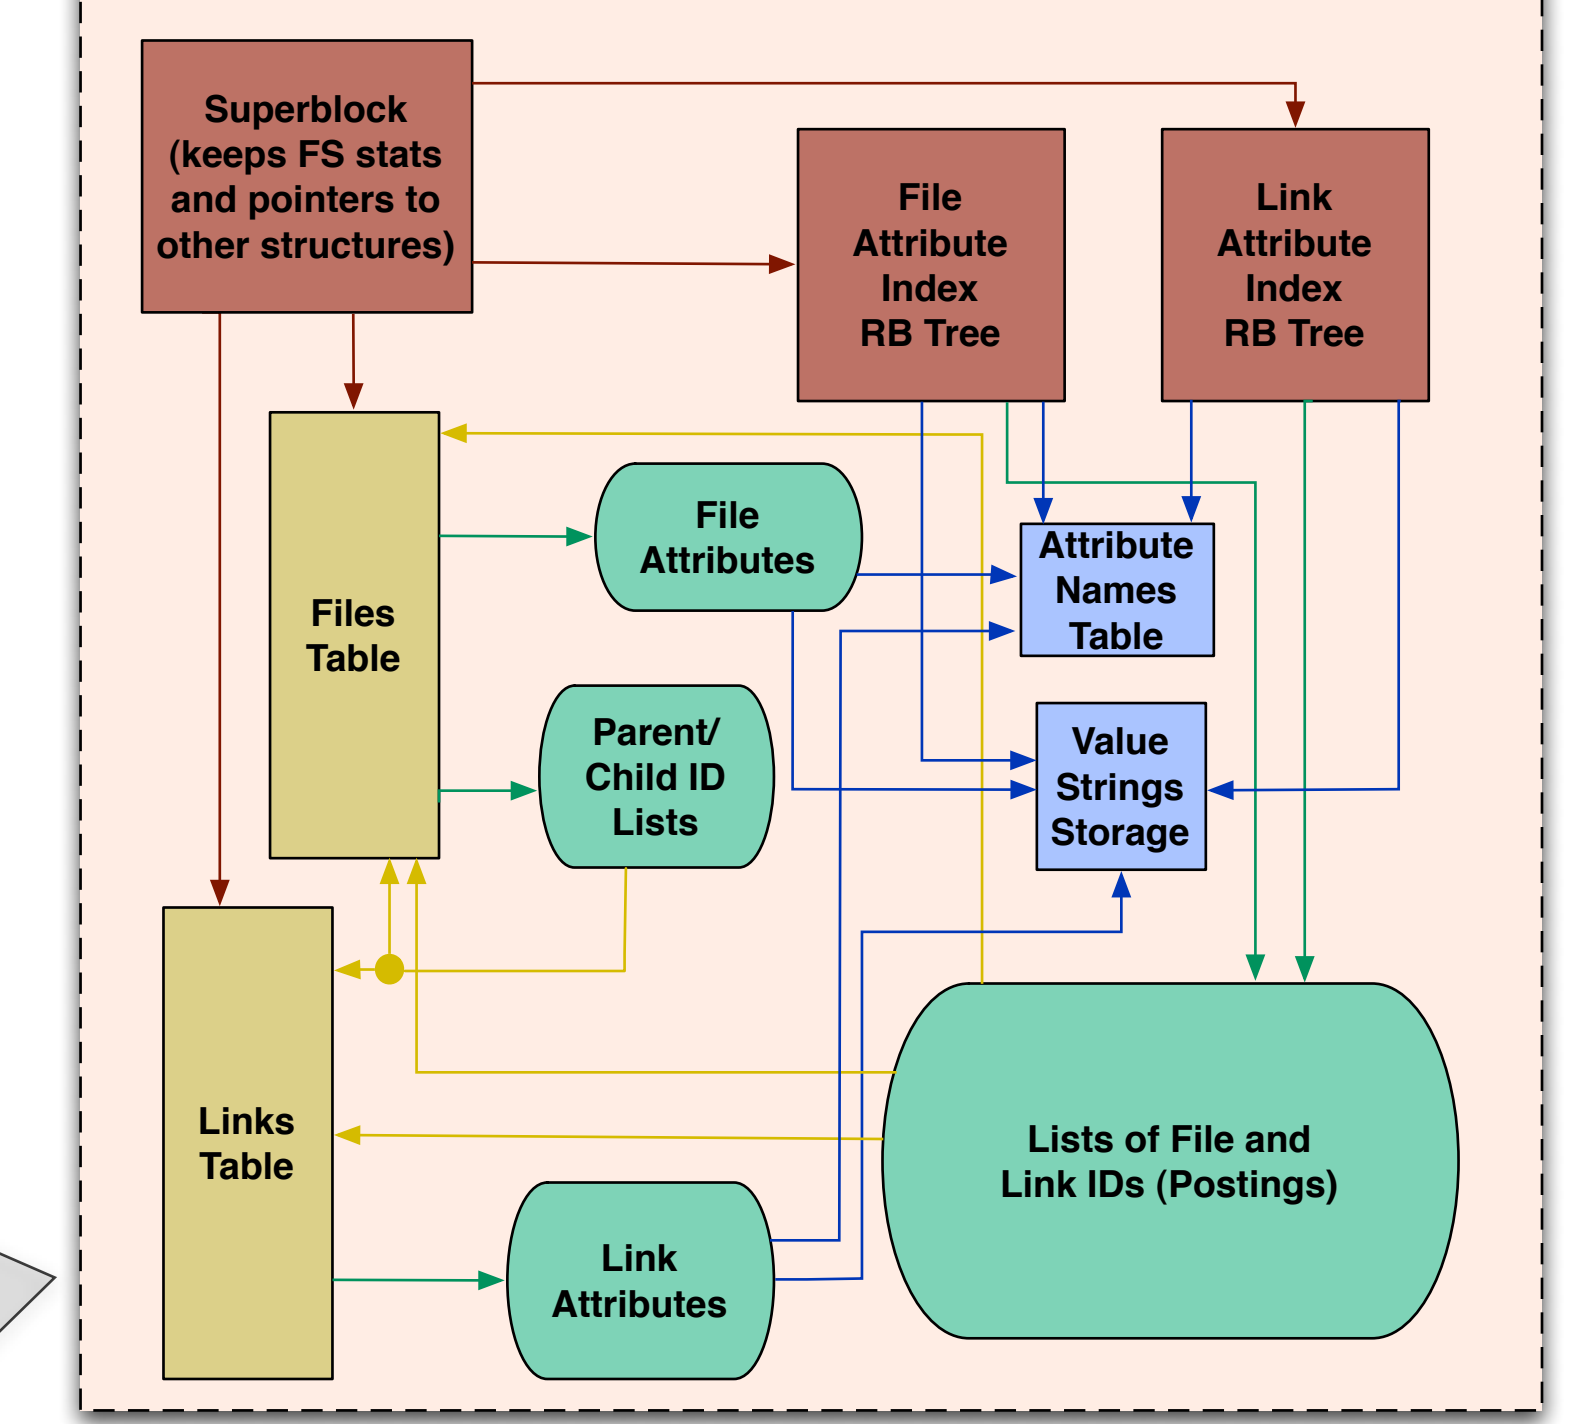
To address this problem, we propose **QMDS**: a file system metadata management service that integrates all file system metadata and uses a graph data model with attributes on nodes and edges. Our service uses a query language interface for file identification and attribute retrieval.



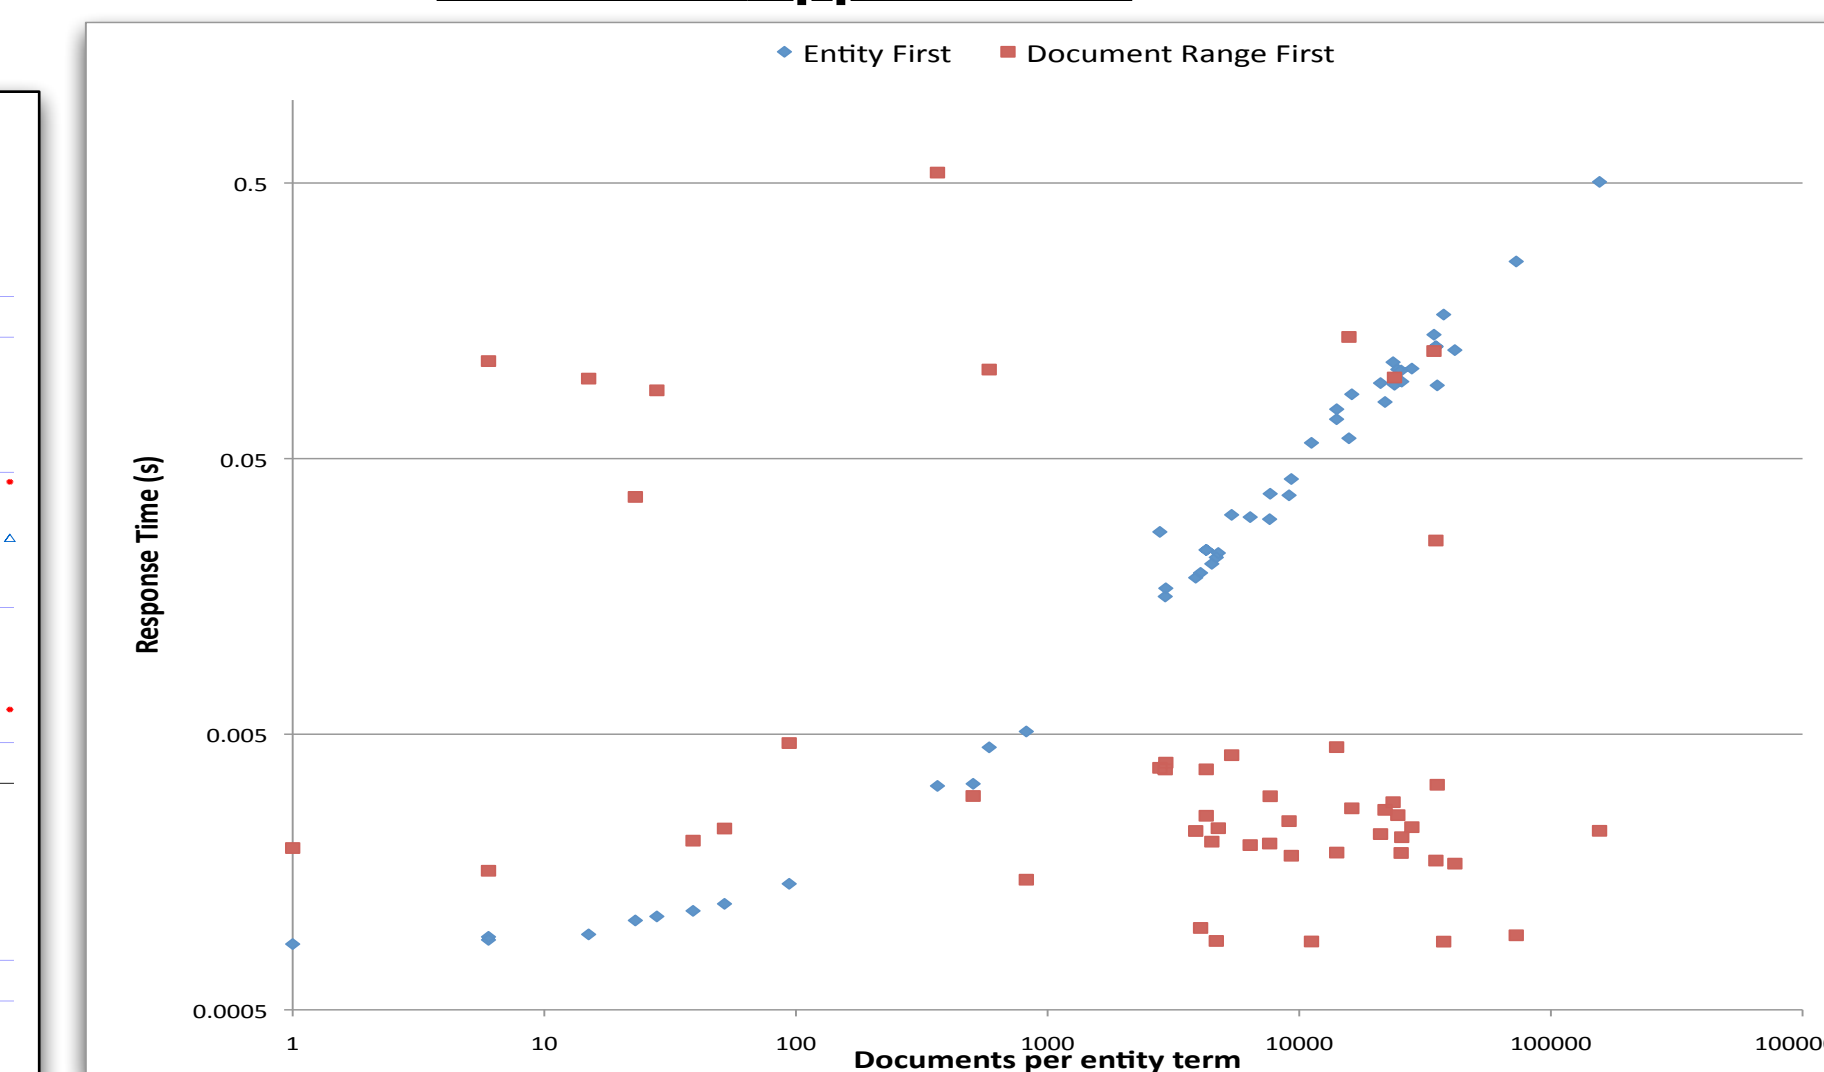
QMDS Prototype Software Architecture



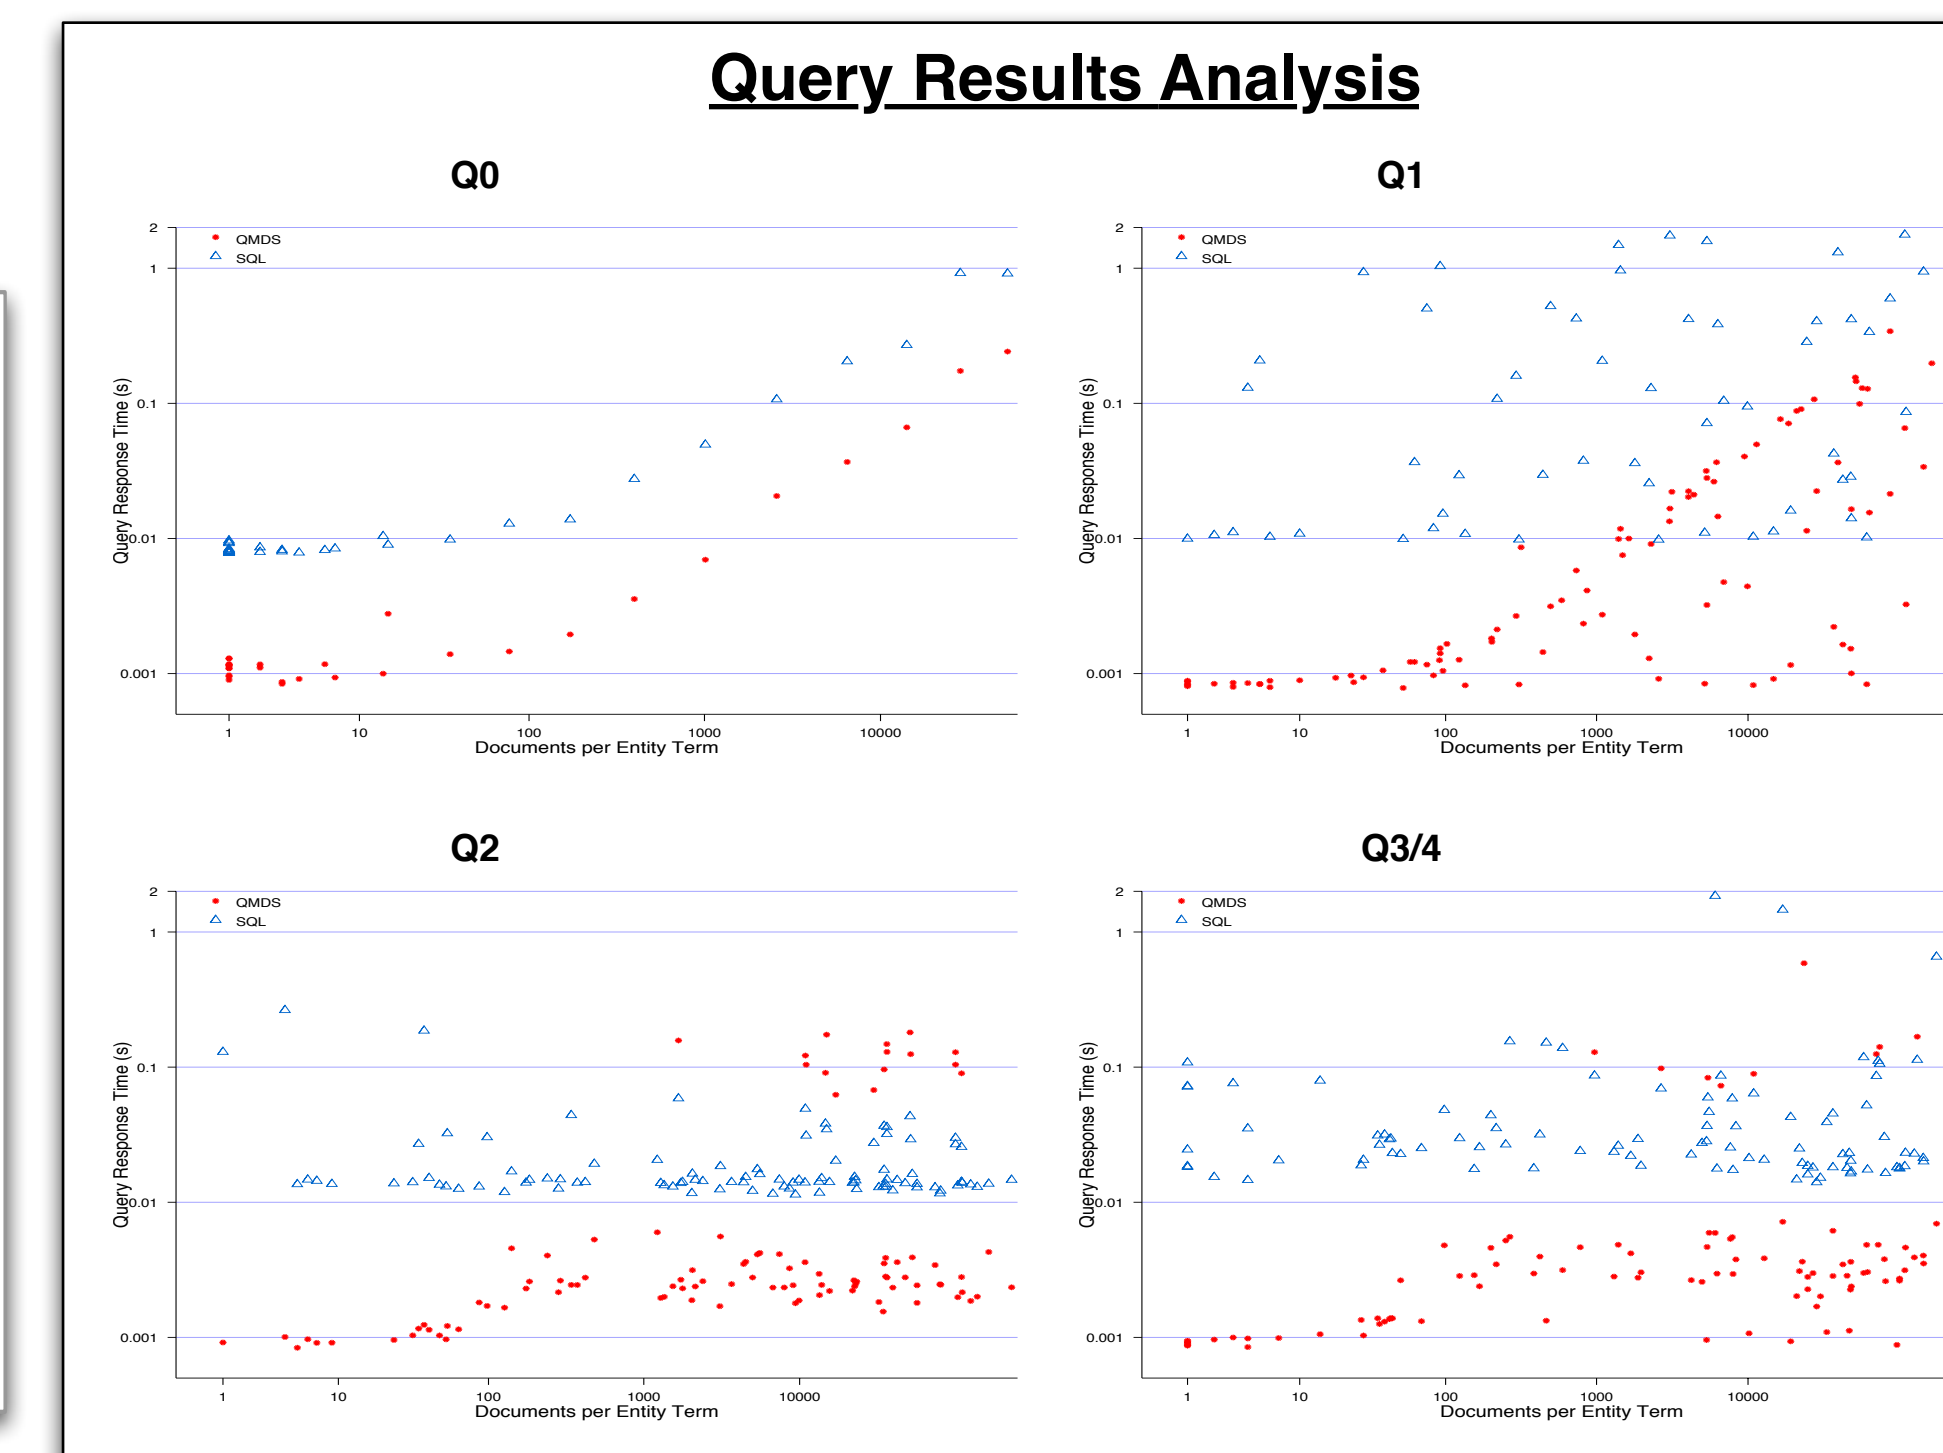
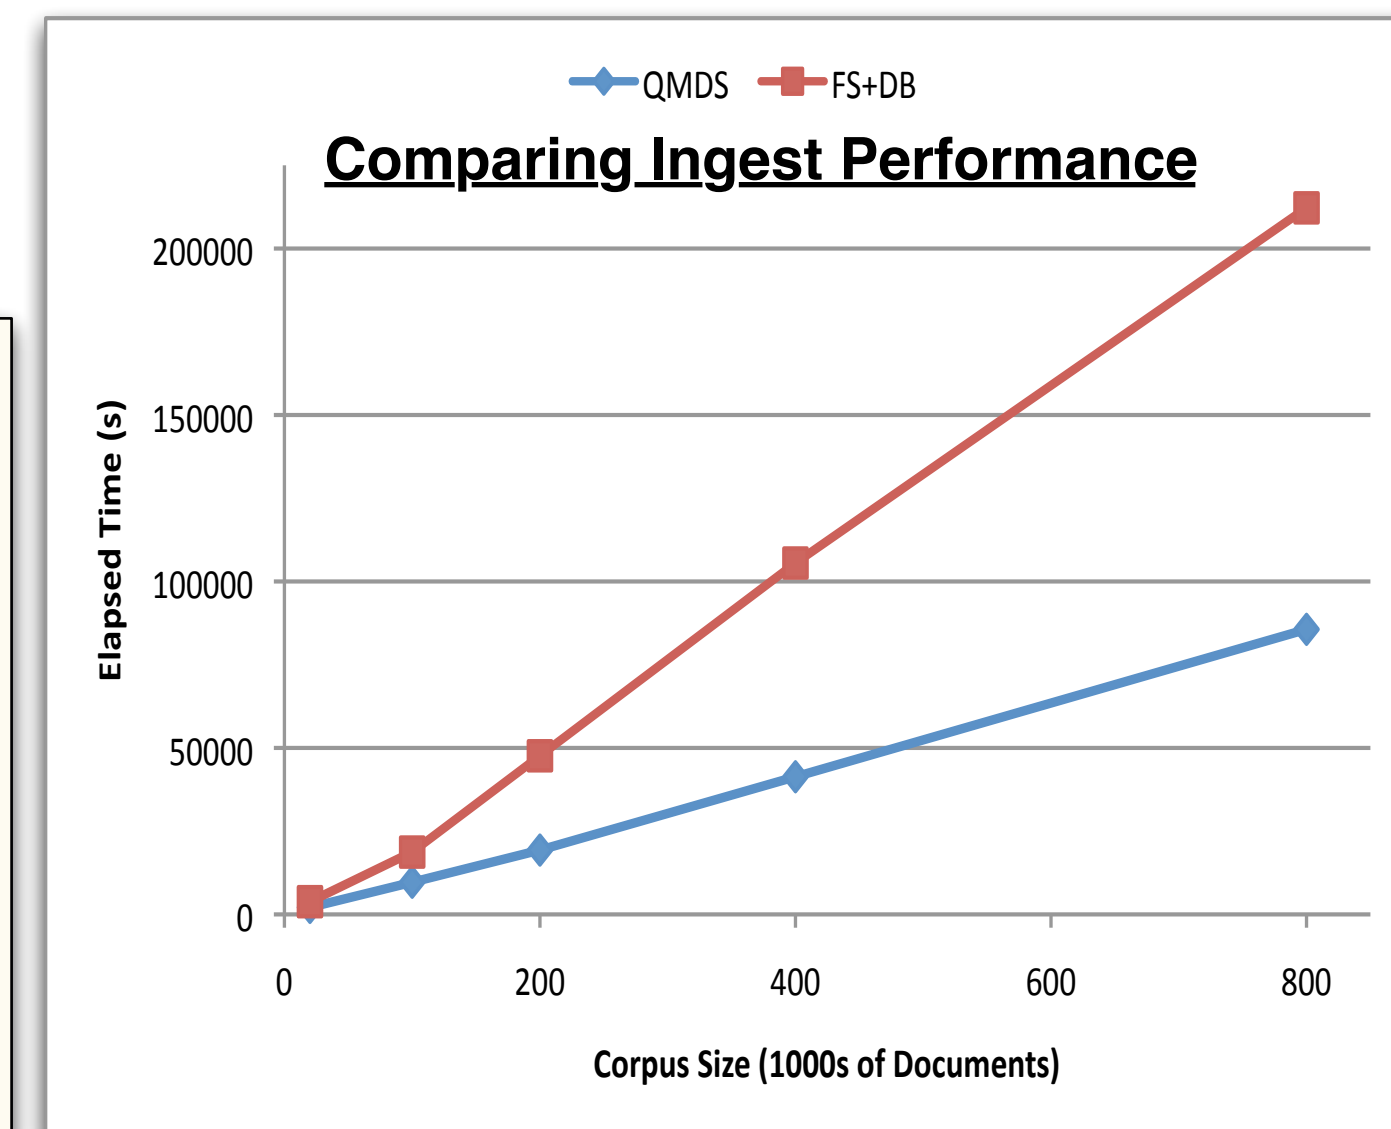
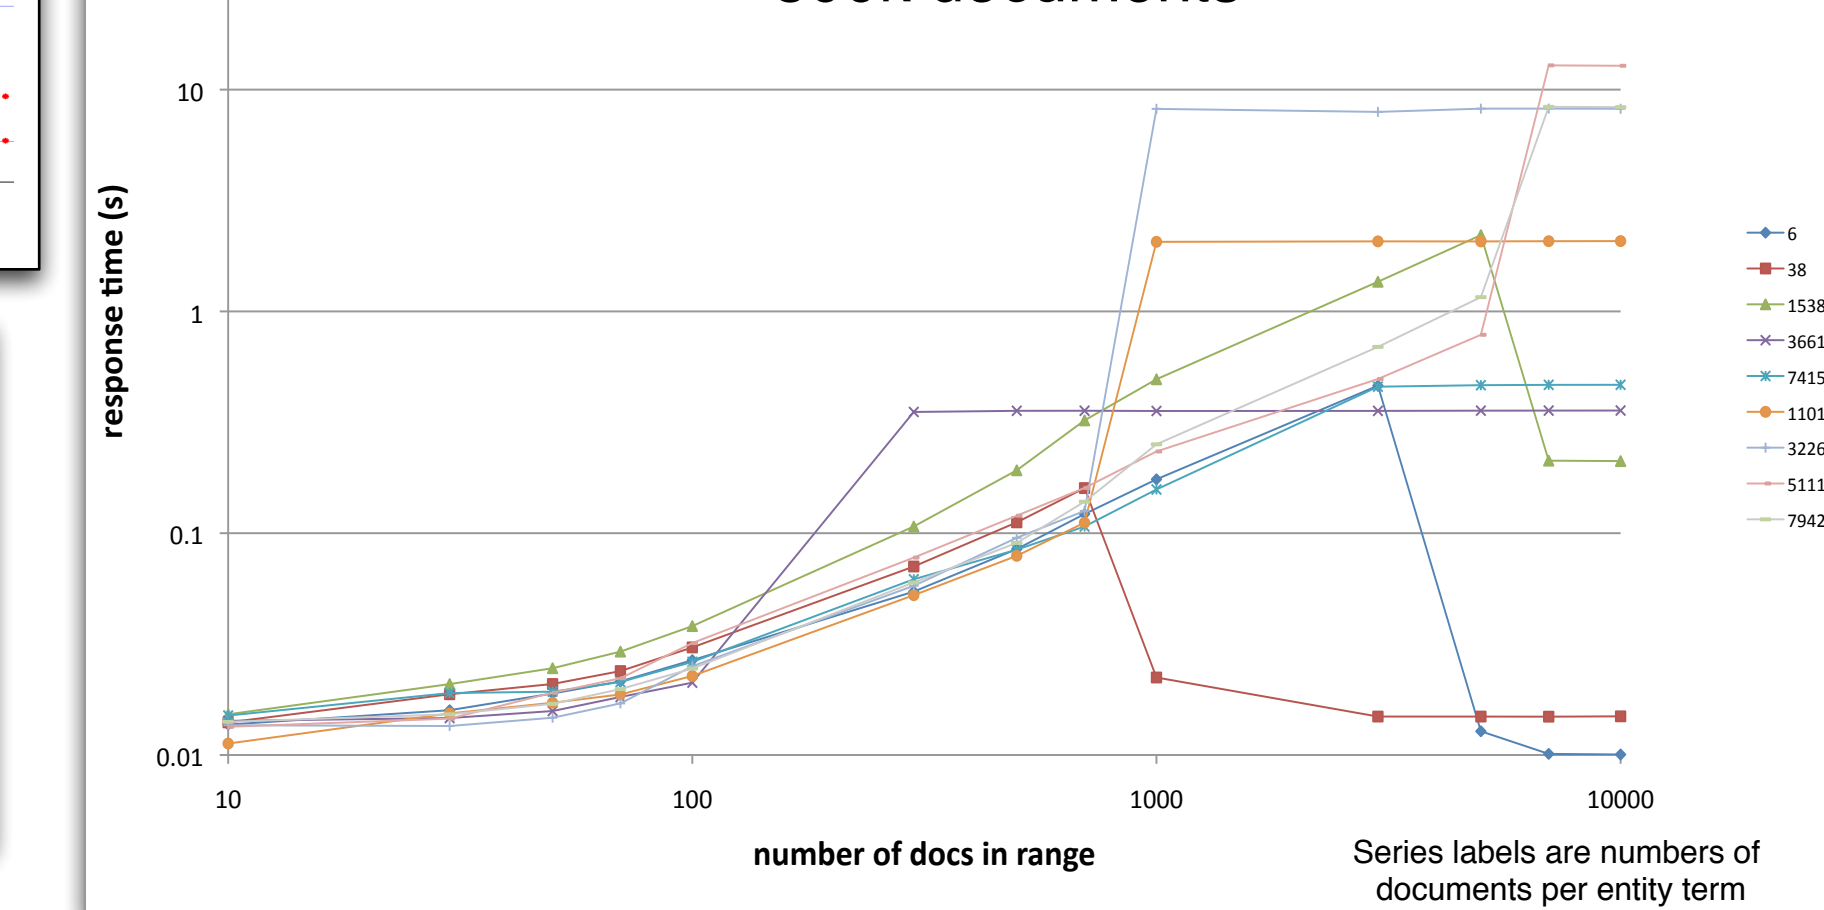
Metadata Storage Data Structure Layout



QMDS Q2 Approaches 400k documents



PostgreSQL Q2 - increasing range sizes 800k documents

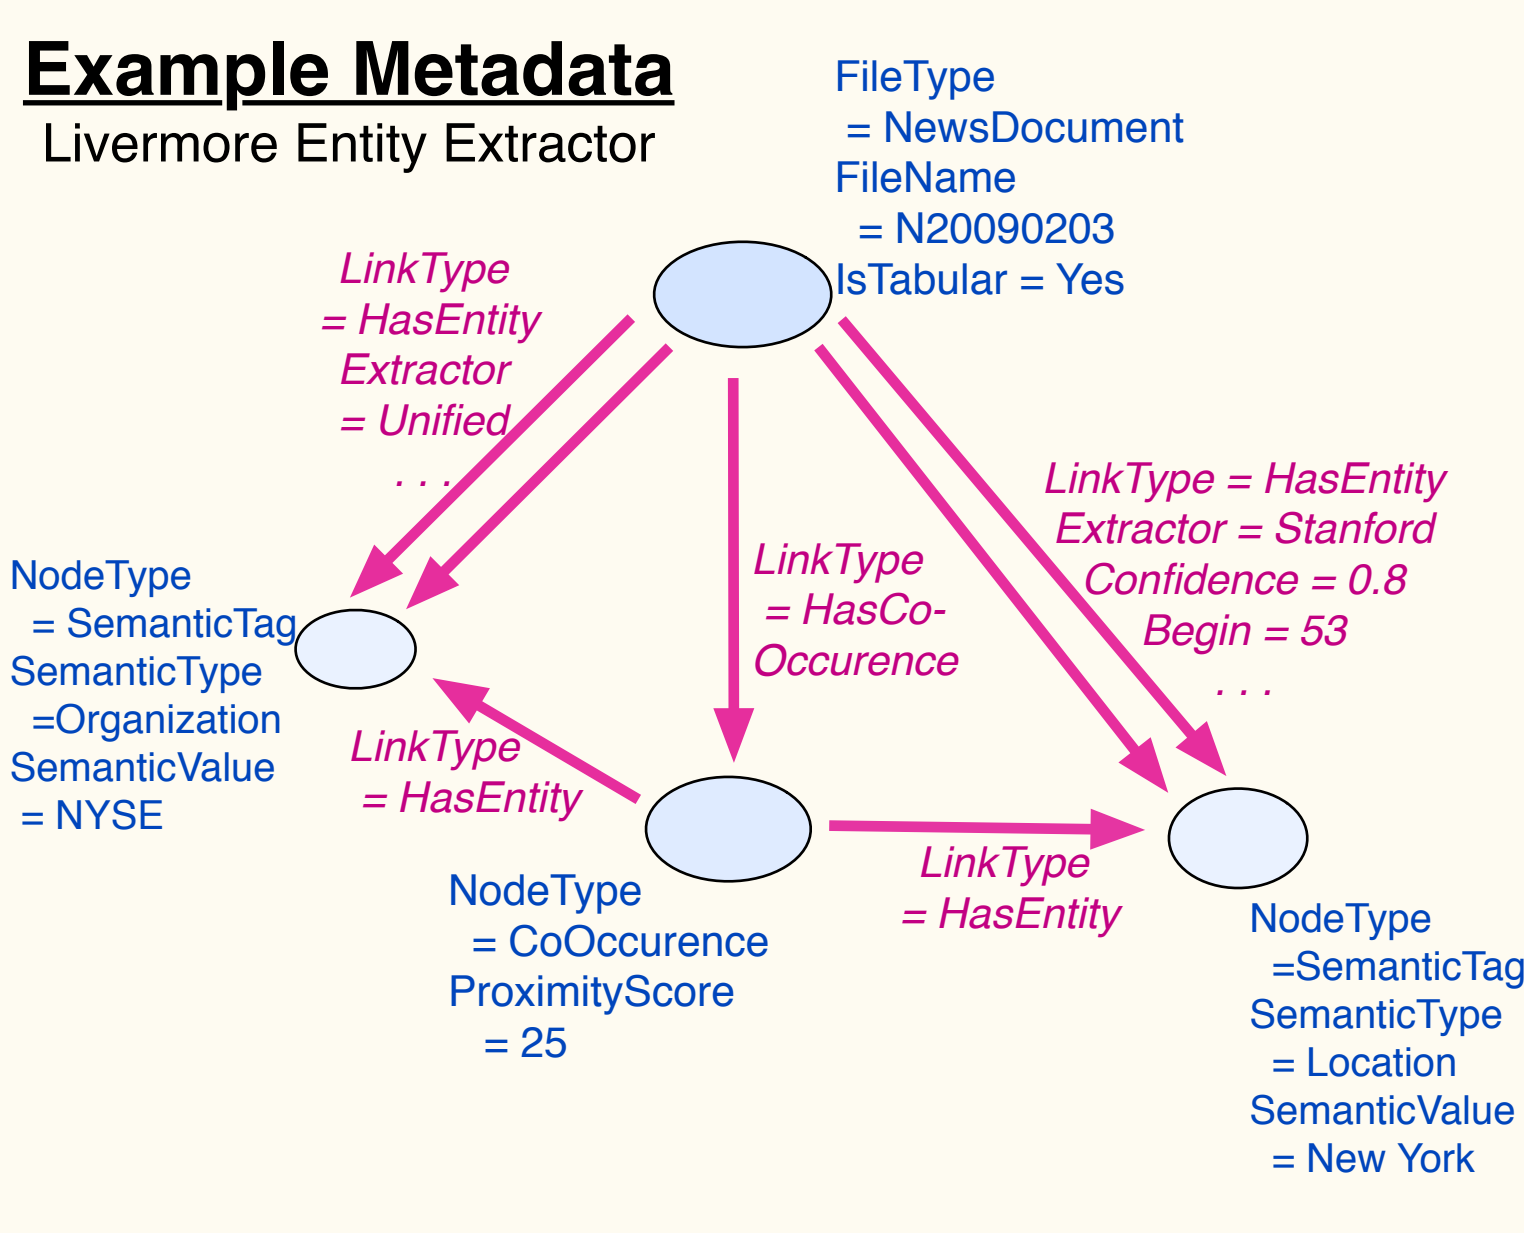


FS+DB Storage, 800,000 Lextrac Documents

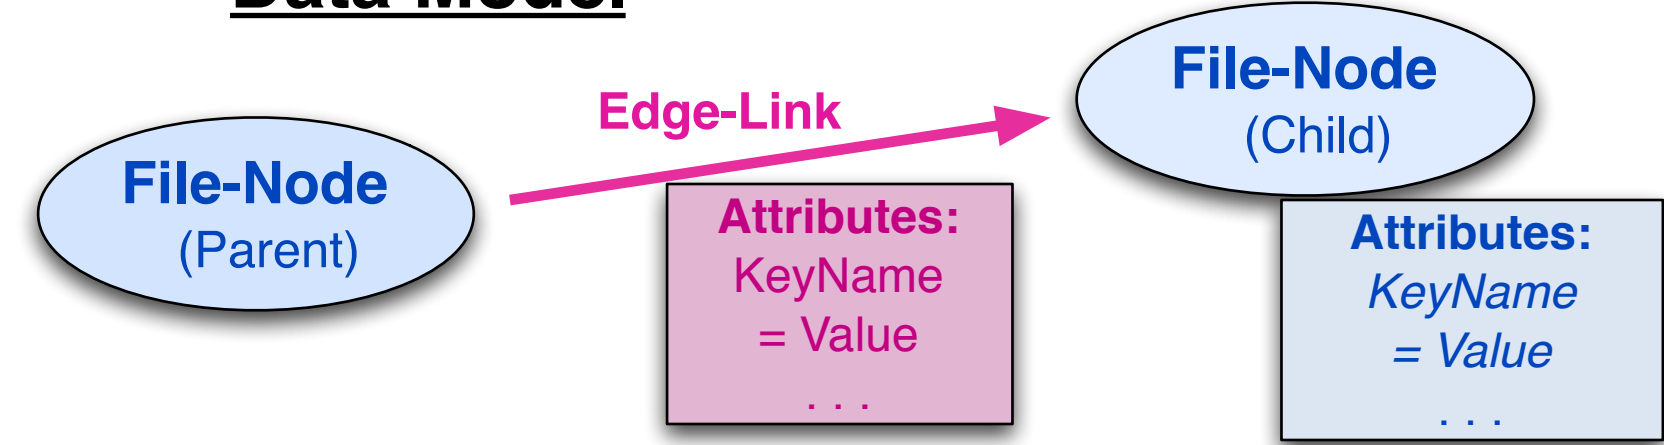
Original Documents	3.8GB
Intermediate Files	14 GB
Database Storage	30 GB

QMDS Storage requirements, Lextrac application

Document Count (1000s)	20	100	200	400	800
Metadata storage size	1.1GB	5.4GB	10.8GB	21.5GB	43 GB
Total Files* + Directories (1000s)	2720	13,200	26,500	53,331	102,002
Links (1,000,000s)	9.23	46.2	92.3	187.4	368.7



Data Model



Experimental Methodology

Workload studies - Use Livermore Entity Extraction for metadata:

- Ingest performance study** - compare QMDS vs file system + PostgreSQL (FS+DB). Postgres configured with schema specific to Lextrac and indexing on all columns. Look at increasing document counts.

- Query performance studies** - compare QMDS with PostgreSQL using 5 classes of queries.

Query terms selected from the data set based on term frequencies. Sample queries:

Q0 Find all documents containing the "location" "New York."
Q1 Find all documents that contain "New York" and "NYSE" with proximity score of "25".

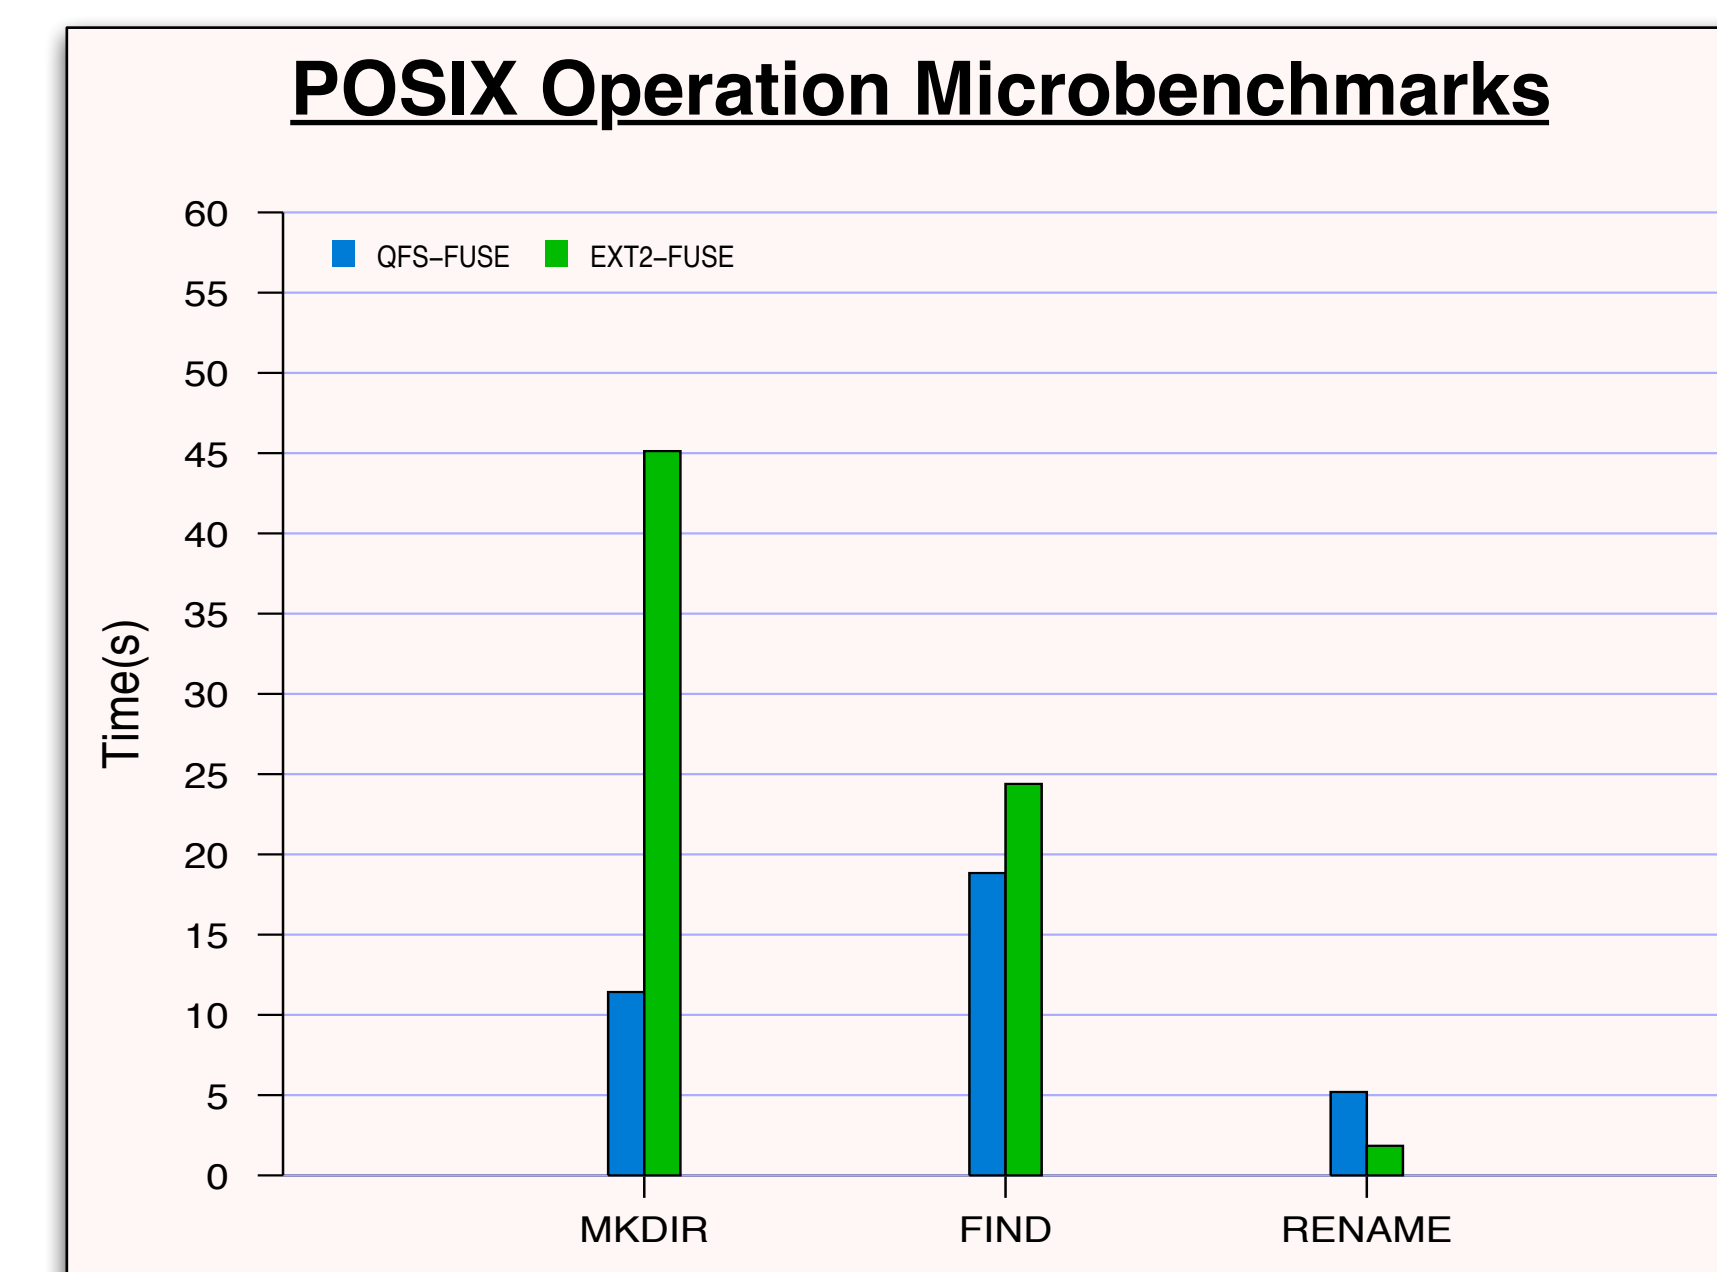
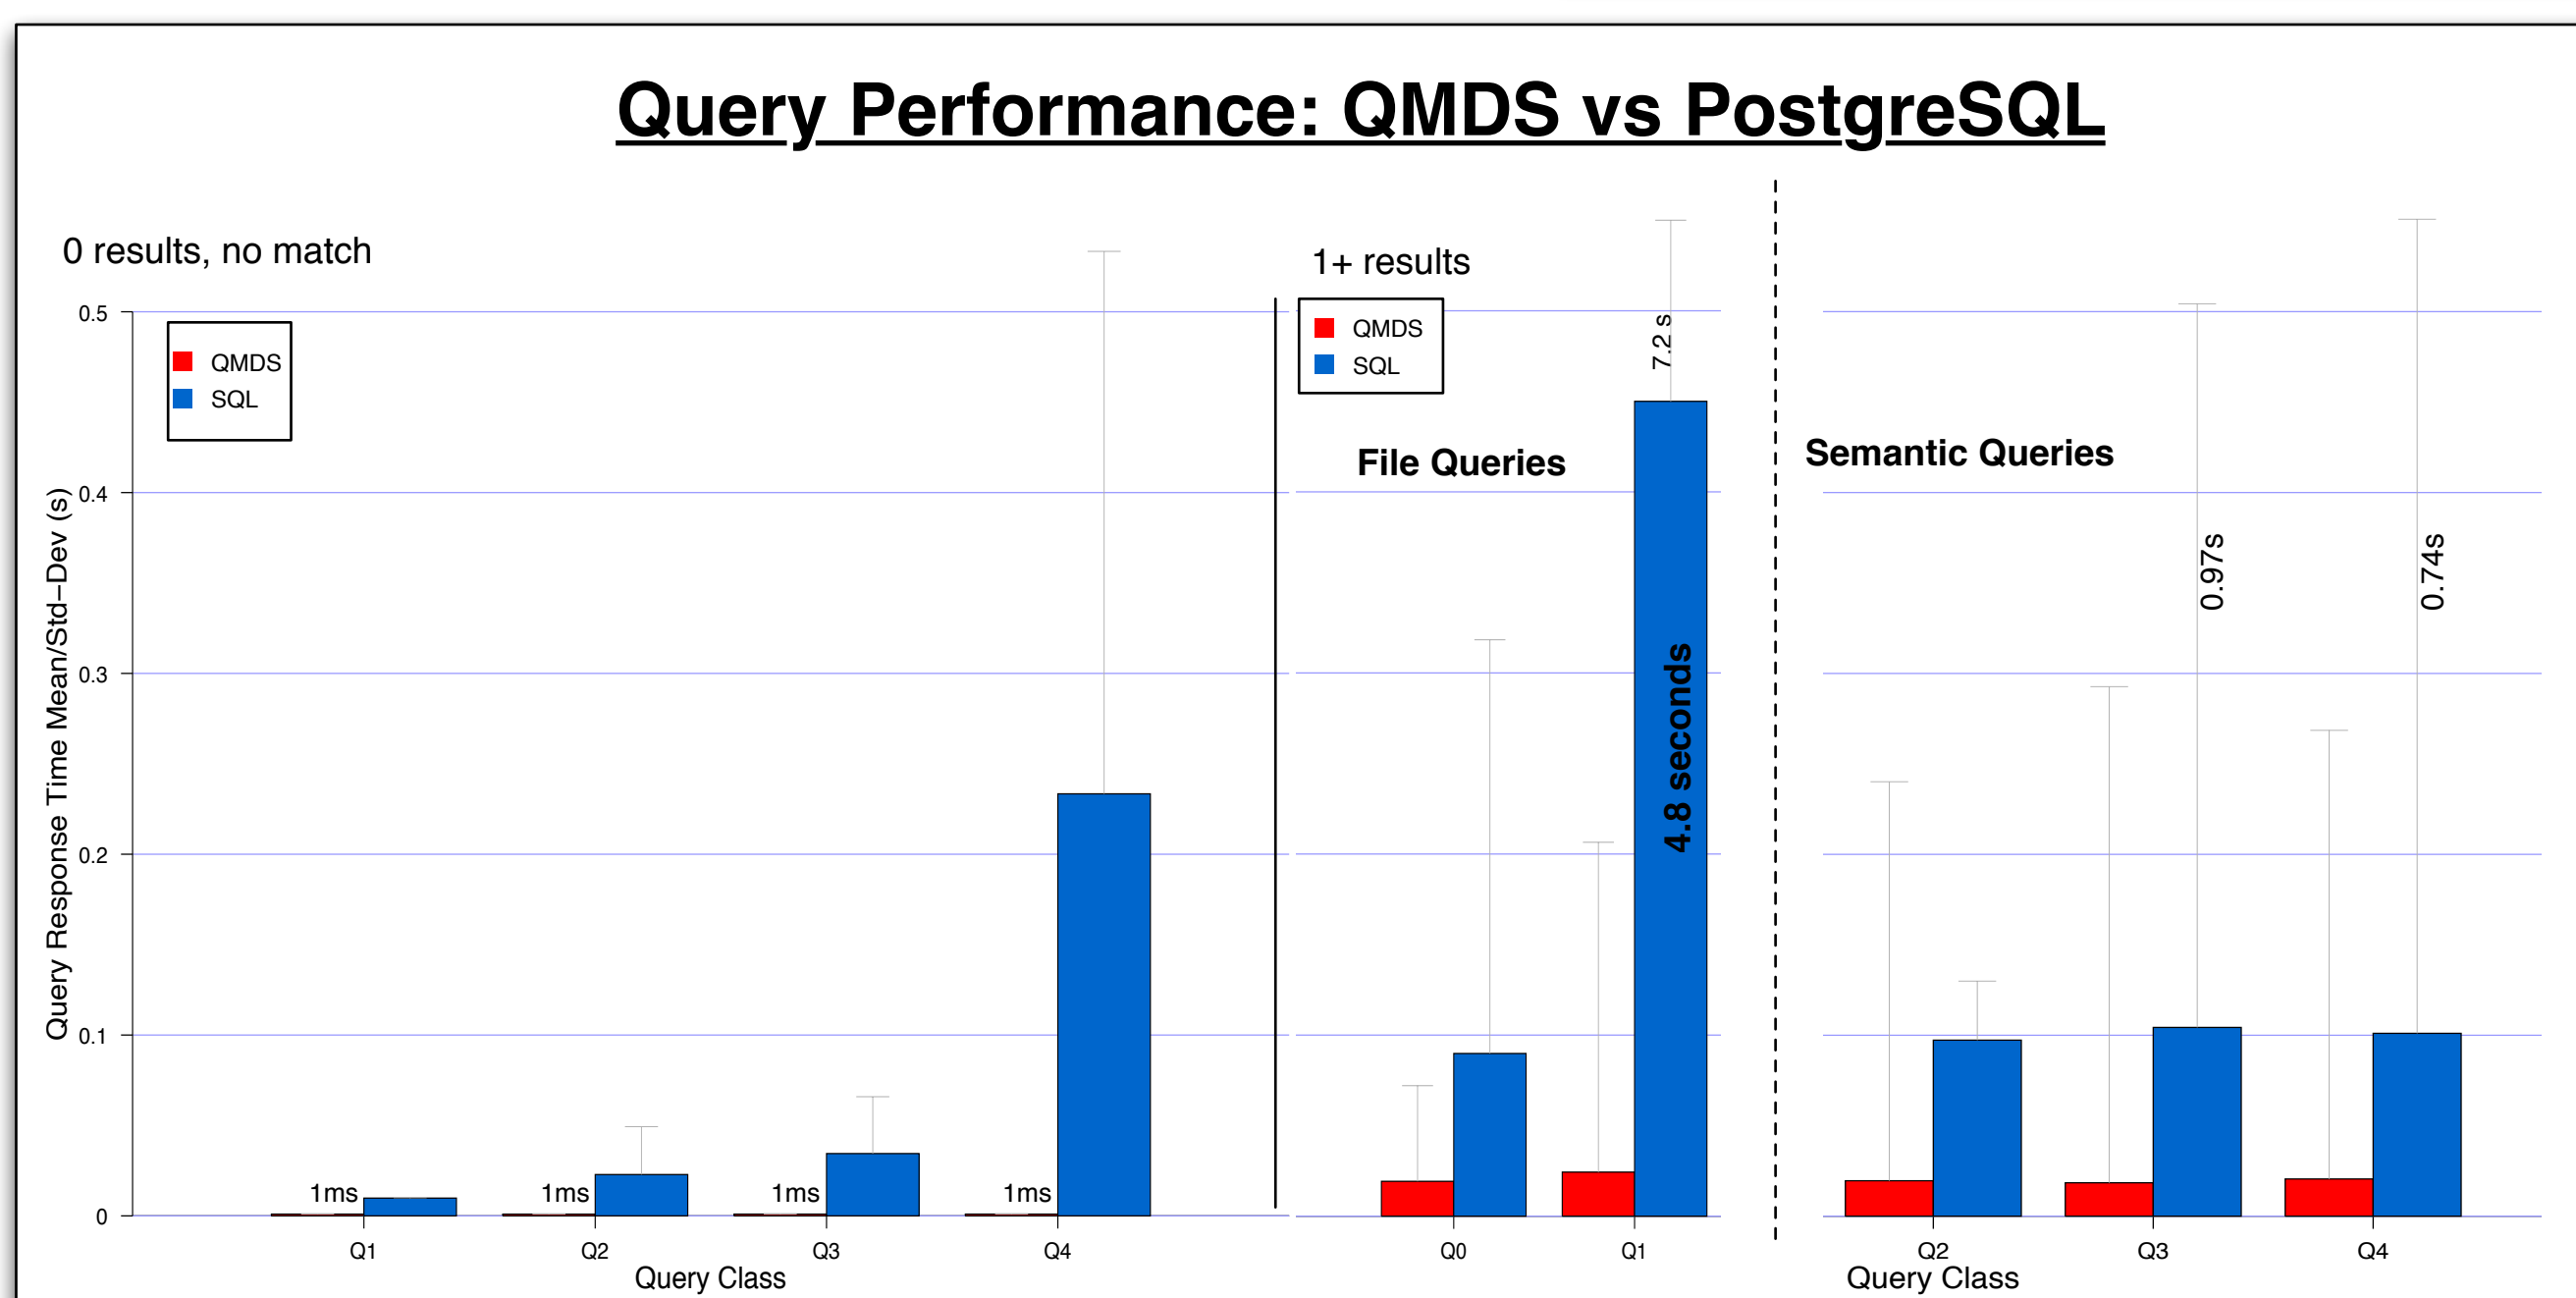
Q2 Find entities co-occurring with "New York" in documents with names in the range "N20090101" - "N20090331" whose proximity score is "25".

Q3 Find the proximity scores in the range 20-30, relating "New York" and "NYSE" in documents with names in the range of "N20090101" - "N20090331."

Q4 Same as Q3 but no range constraint for score.

- POSIX file system operations study** - evaluate overhead of QMDS metadata management over standard file systems. Operations measured:

- Create (MKDIR)
- Stat/openssl/readdir (FIND)
- Move (RENAME)



Results Summary

- QMDS 2.44x faster than FS+DB on Ingest workloads with 800000 documents
- QMDS performs 10-40x faster than PostgreSQL in answering queries that have a non-matching term.
- QMDS is on average 200x faster for Q1, 4-5x faster Q0, Q2-4 (range of 11 documents)
- Range first strategy generally more effective.
 - ➔ Need heuristics to optimize strategy selection.
- QMDS optimized for file, link creation over file moves.