

# Clustrix

*“Speed. Scale. Simplicity.”*



---

BROWN

# History

- Founded 2006
- San Francisco based
- ~60 employees
- CEO and President Robin Purohit
  - former HP Vice President



---

BROWN

Clustrix

# Overview

- focuses on seamless scalability
- distributed shared-nothing architecture
- “moves the query to the data, not the data to the query”
- full ACID compatibility
- uses Sierra Database Engine



# “slicing” data

- separate tables into “slices”
  - replicate slices for redundancy
  - number of slices determined by table size
  - placement of the slices and their replication factor is dynamic
  - determination of data to slices can be range or hash based



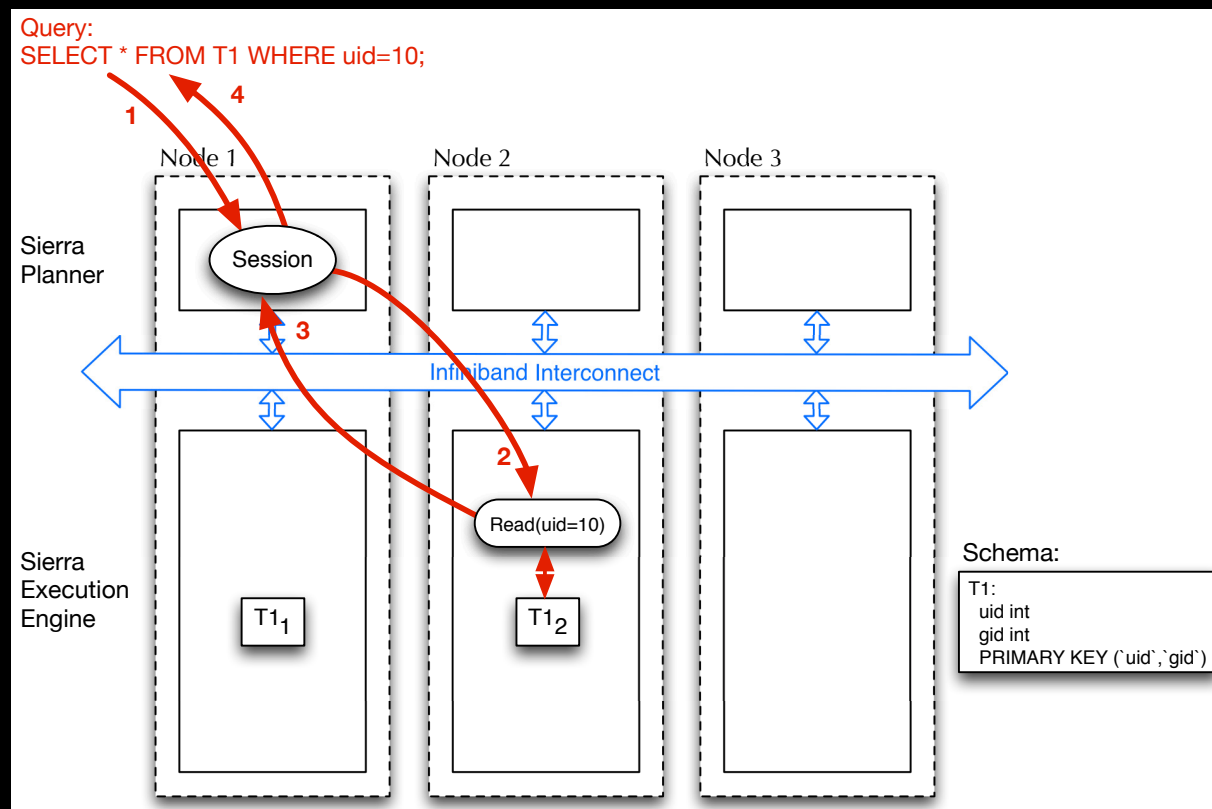
# Query execution

- Incoming query translated into “query fragments” by Sierra Planner
- fragments are sent to the node that contain the data for that fragment
- page-level locks (i.e. no global table or row-level)
  - data consistency handled by MVCC

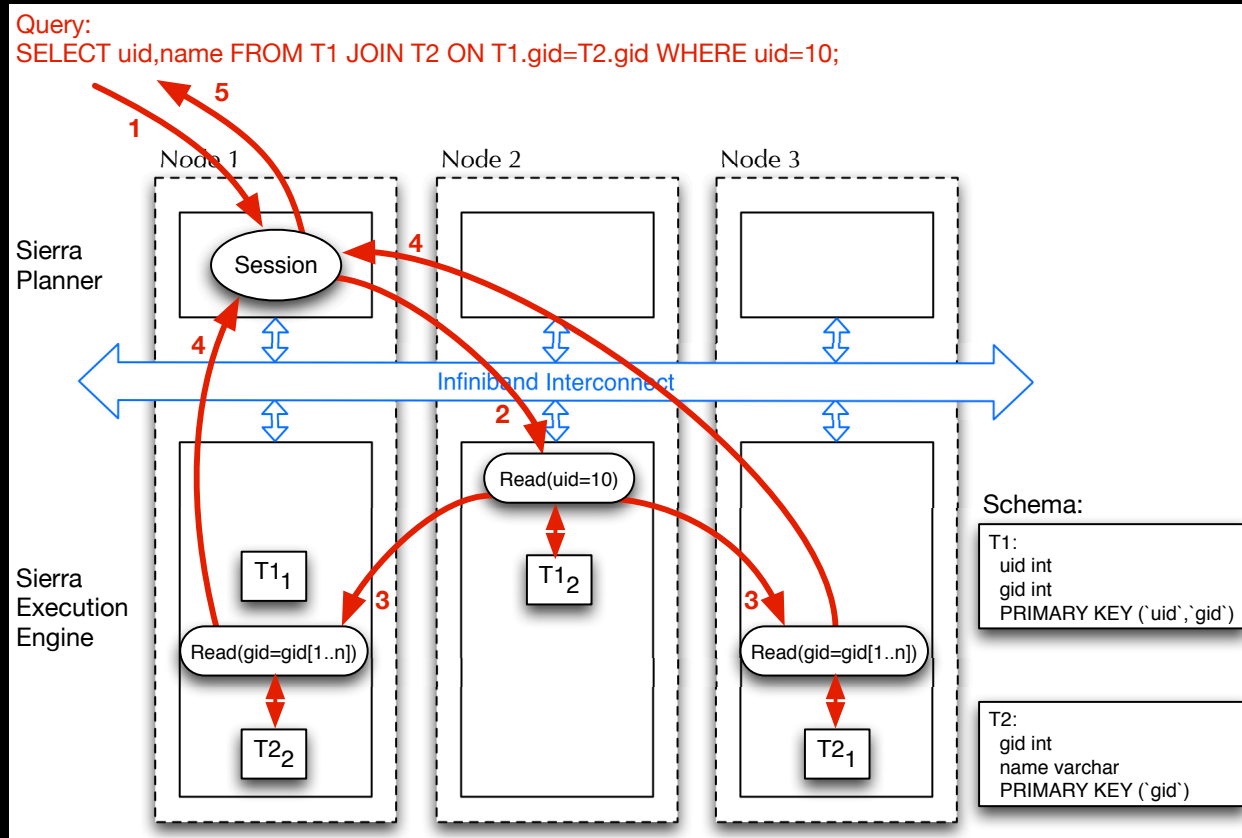


BROWN

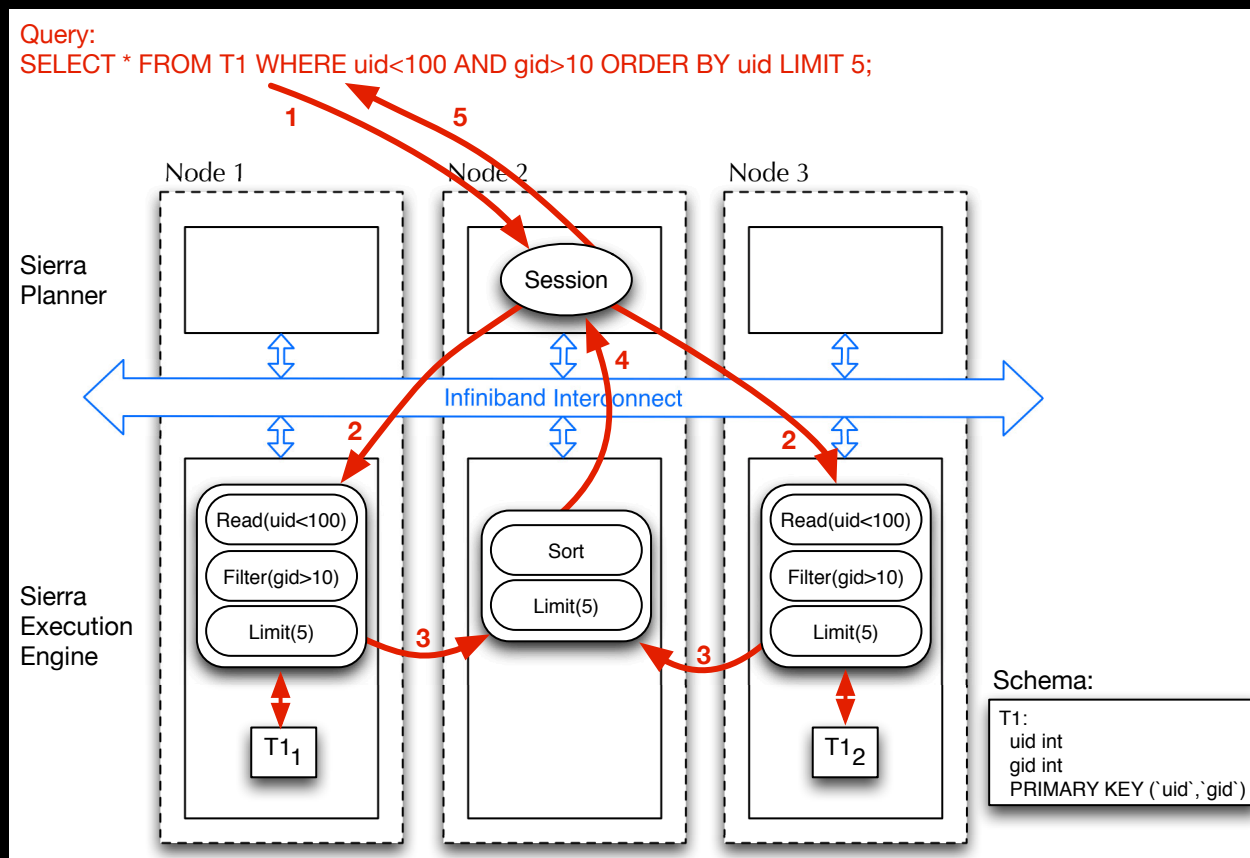
# Example 1 (read)



# Example 2 (join)

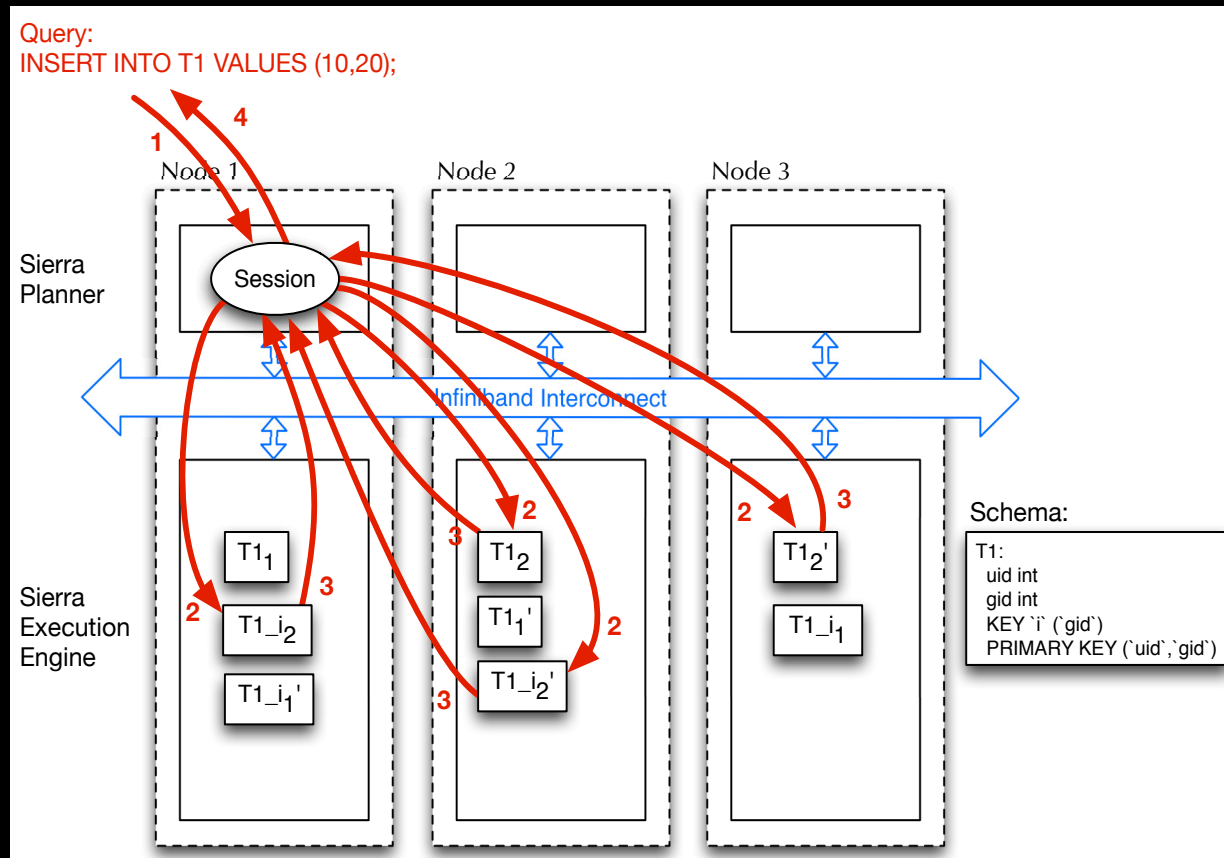


# Example 3 (complicated)





# Example 4 (insert)



BROWN

Clustrix

# Conclusions

- Benefits
  - distributed database with no tuning necessary
  - balanced I/O
    - Clustrix is purchased by buying Clustrix boxes
    - Infiniband interconnect
  - scale by adding a new Clustrix box
    - near-linear by 3<sup>rd</sup> party benchmarks
  - “drop in” replacement for MySQL
- Users
  - The Ladders, AOL, PhotoBox



---

BROWN