Excuse me,
but is your data dead?

Toby Lehman

Aerospike

The Relational Landscape

Stonebraker Land

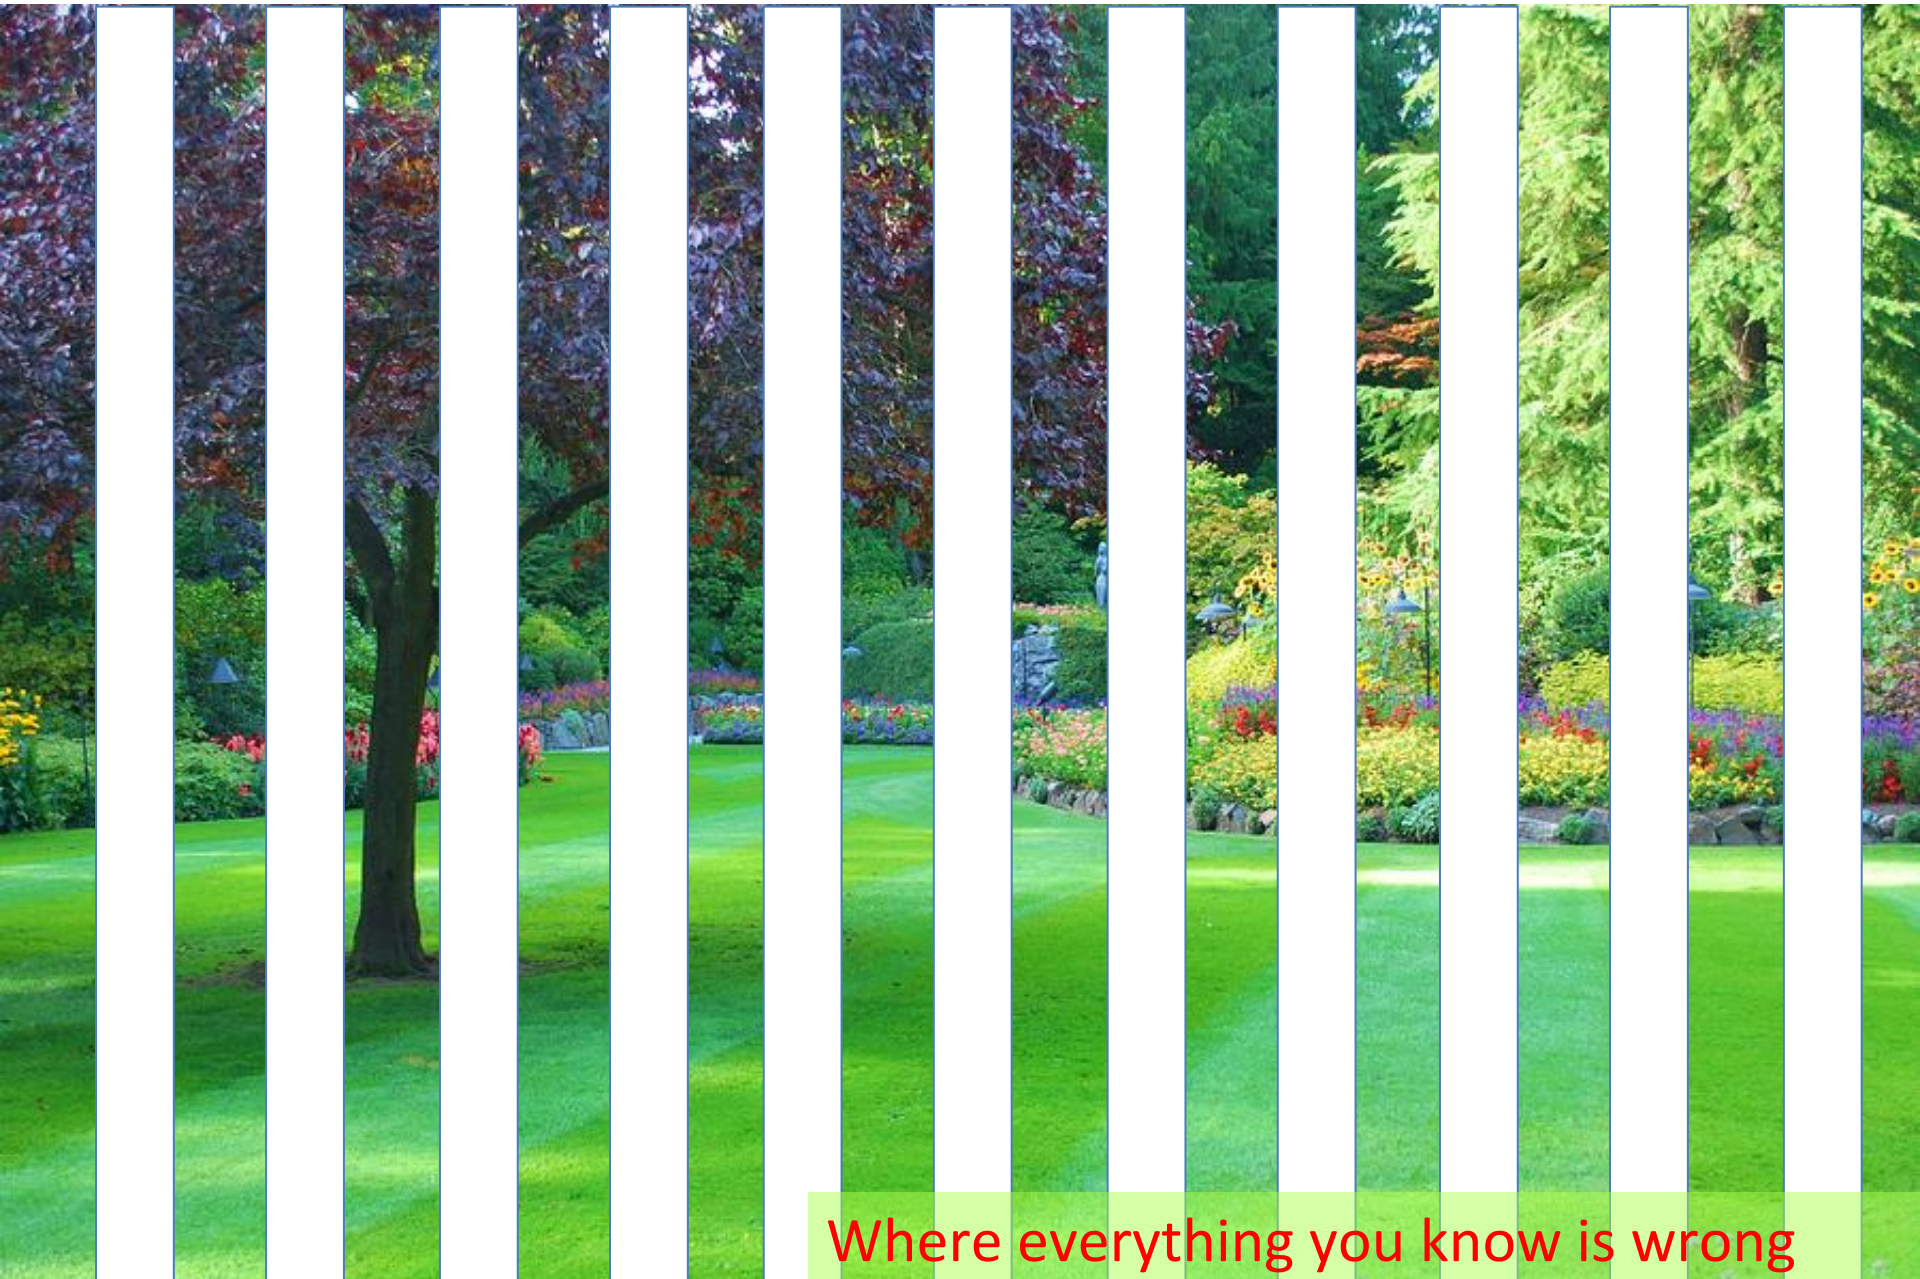

Where everything you know is wrong

The Joy of Data Modeling

- Section
- Course #
- Course Name
- # Seats
- # Enrolled
- Room
- Time
- Days
- Instructor
- Students

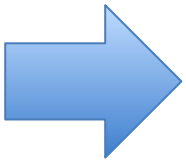

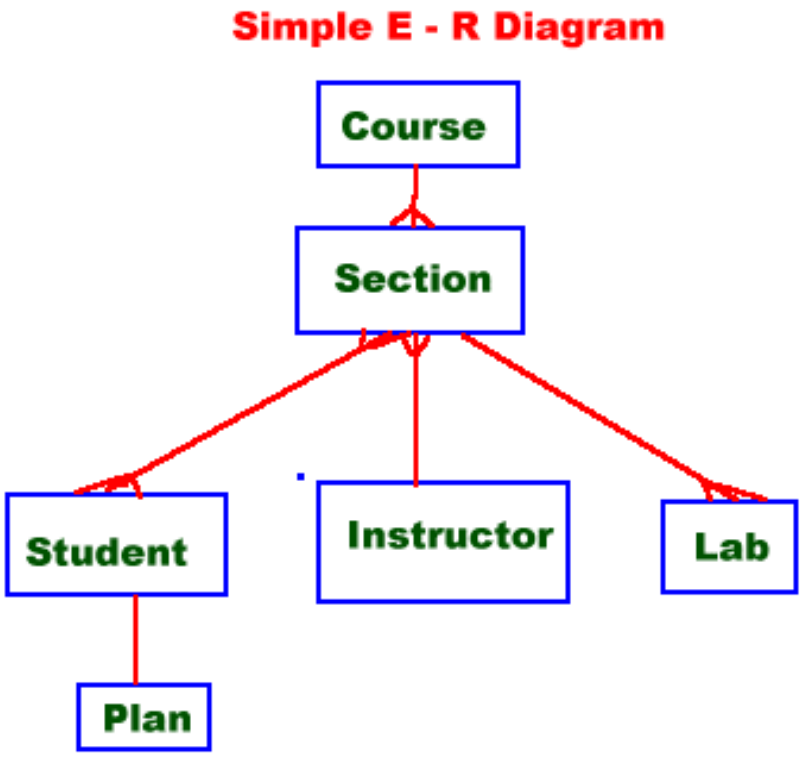

Embrace Normalization

- The Famous Puppy Poster shows the complexities of normalization
- Suffer the penalties of joining the data back again.
- Why Normalize? Because LIVE data may have Update or Delete problems if copied.

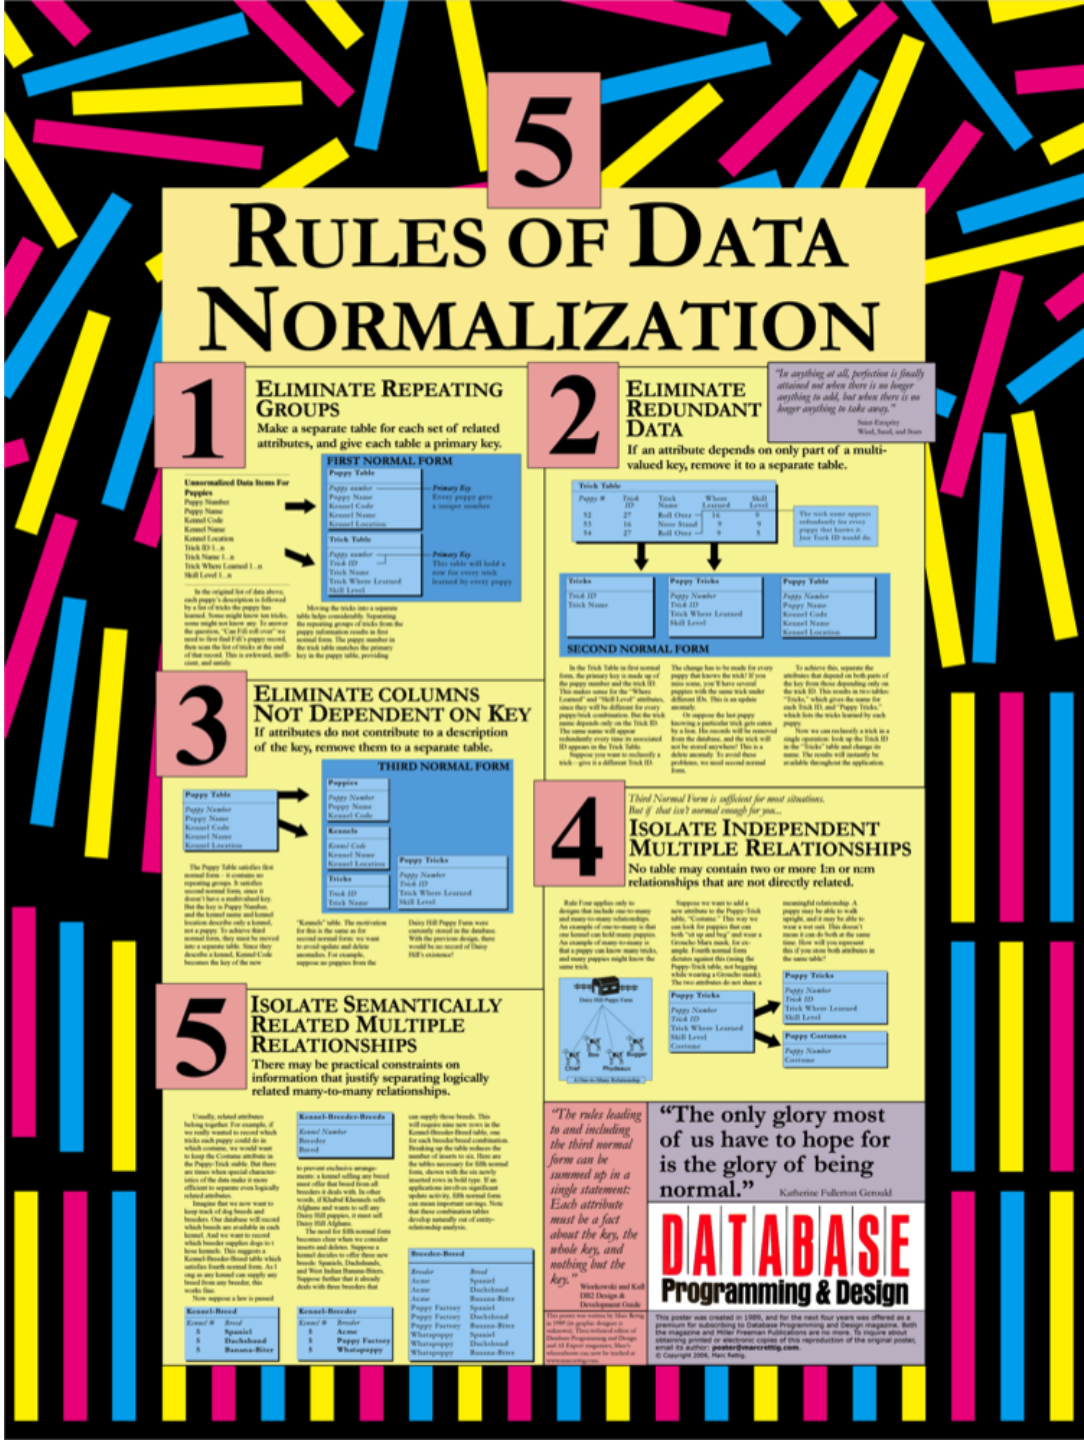

The NoSQL Landscape

But … What about 1:N, N:M Relationships?

Let's Look More Closely at the Data

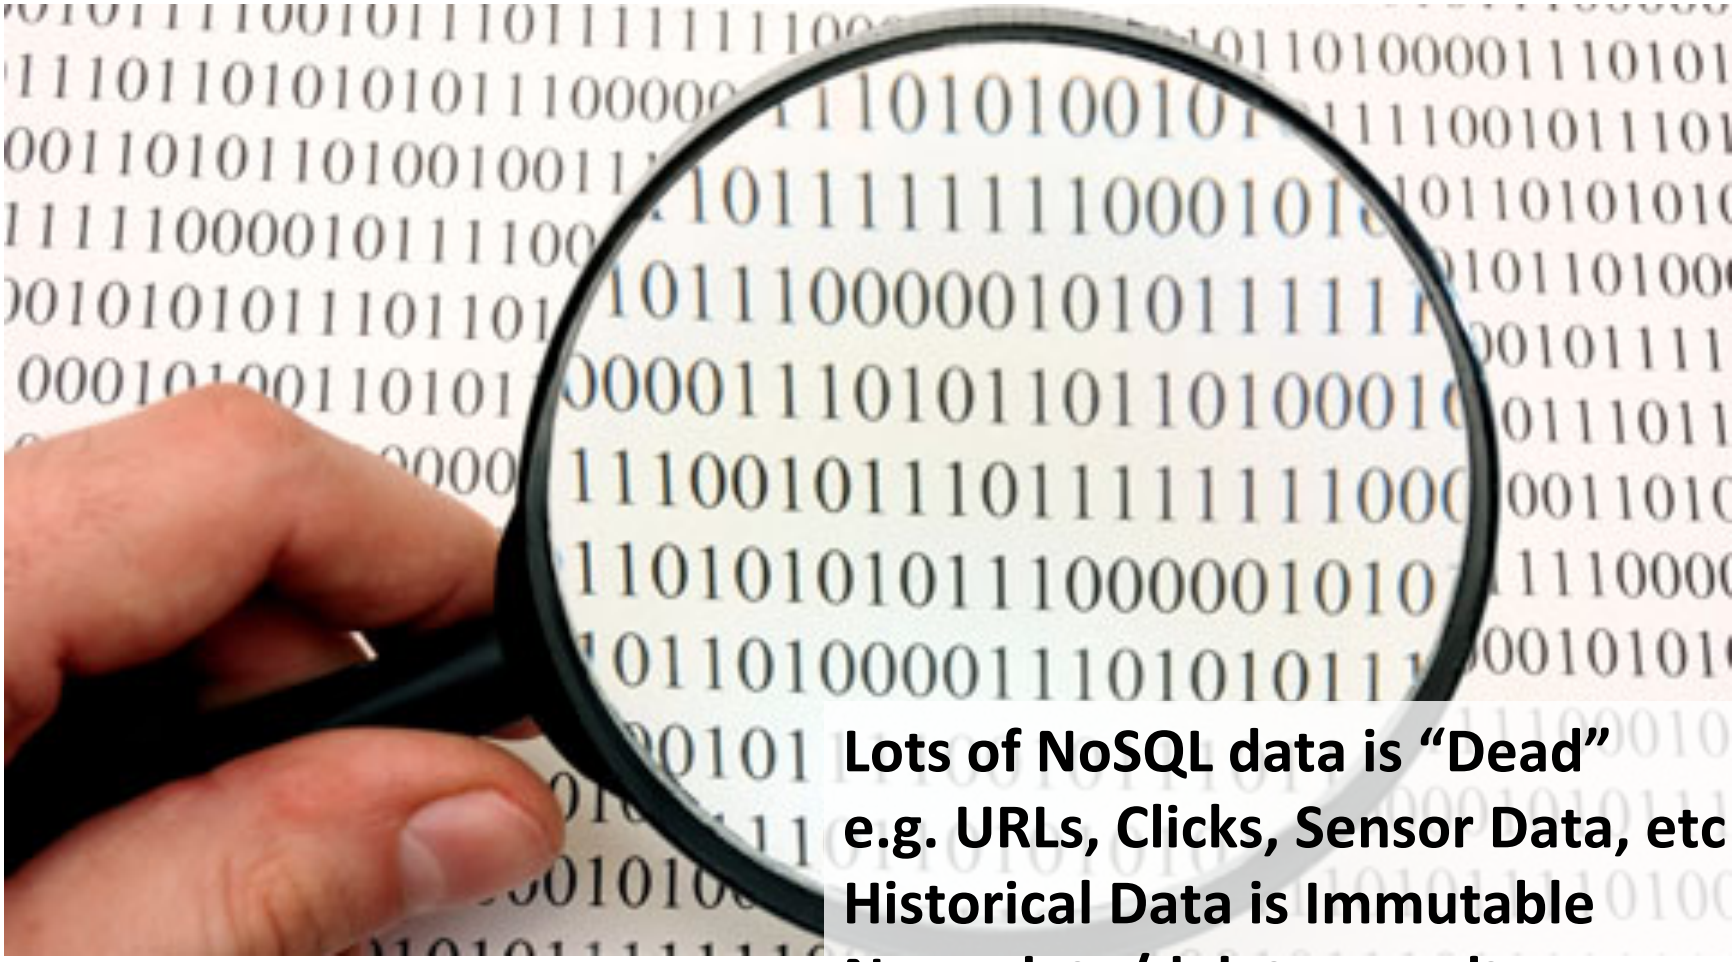

Lots of NoSQL data is "Dead"
e.g. URLs, Clicks, Sensor Data, etc
Historical Data is Immutable
No update/delete penalty
➔ Promotes Collection Types

The NoSQL Landscape

Object Collections Can Satisfy Relationships

Aerospike Large Data Types

- Aerospike Record

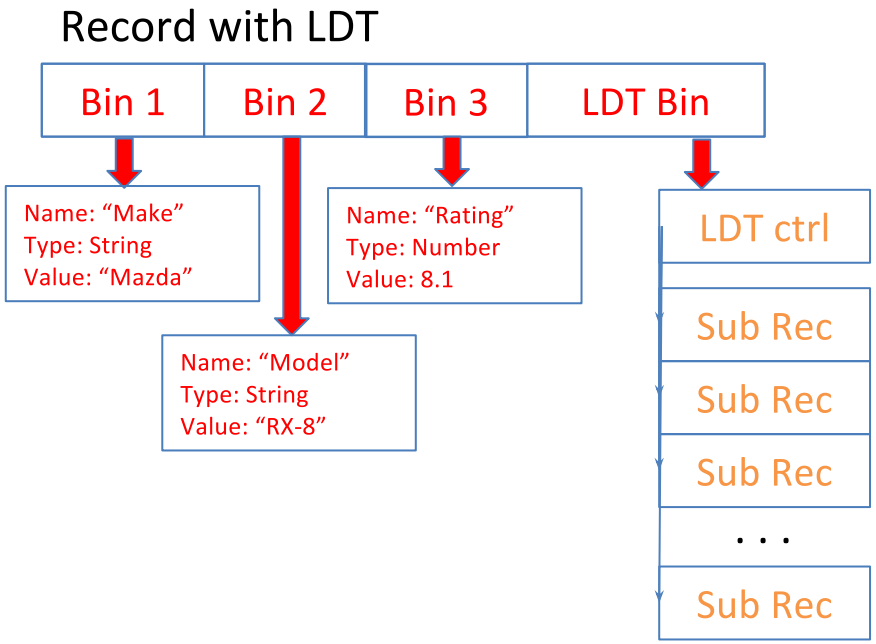

- Set

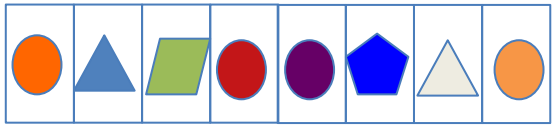

unique collection

- Map

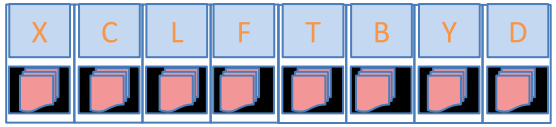

Name/Value Pairs

- List
- Stack

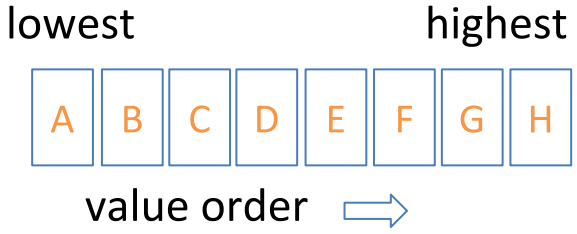

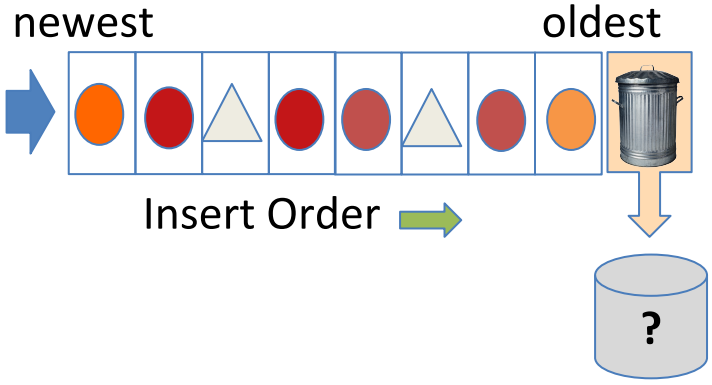

AEROSPIKE

Summary

- In NoSQL Land, much of the data captured is historical (dead data)
- No penalty, other than increased storage, to keep copies
- Often a huge performance win to read a collection as part of a record (much like a pre-computed materialized join)
- Different collection types (stack, list, map, set) match the different application access patterns

- **Questions: toby@aerospike.com**