Objectivity®

www.Objectivity.com

InfiniteGraph

Objectivity/DB

How we Learned to Stop Worrying and Solve the Distributed Graph Problem

Ibrahim Sallam

Director of Development

What we'll talk about

- Graphs, what are they and why?
- Graph Data Management. Why do we need it?
- Problems in Distributed Graph
- How we solved the problems

Graphs

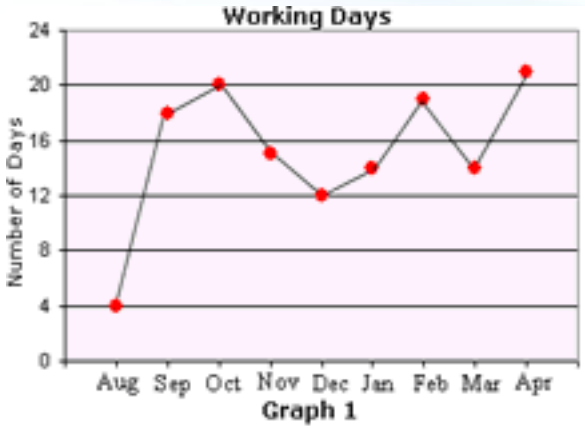

Graph 1

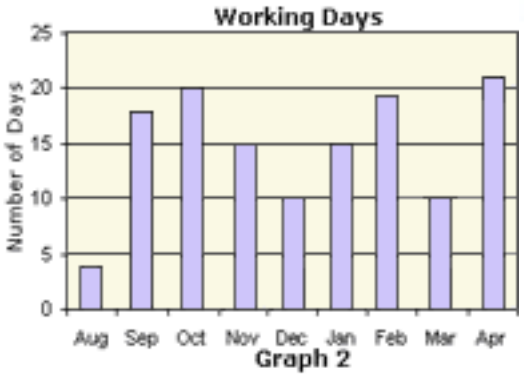

Graph 2

Working Days

Month	Number of School Days
August	●
September	● ● ● ●
October	● ● ● ● ●
November	● ● ● ●
December	● ● ●
January	● ● ● ●
February	● ● ● ● ●
March	● ● ●
April	● ● ● ● ●

● = 4 days

Graph 3

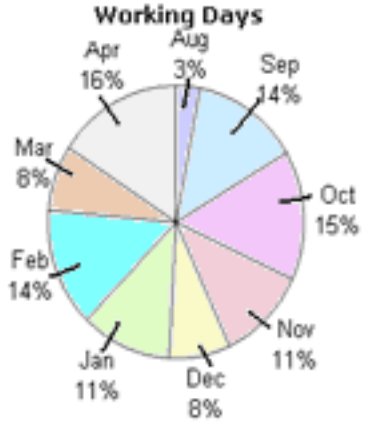

Graph 4

Objectivity

InfiniteGraph
Objectivity/DB

Graphs

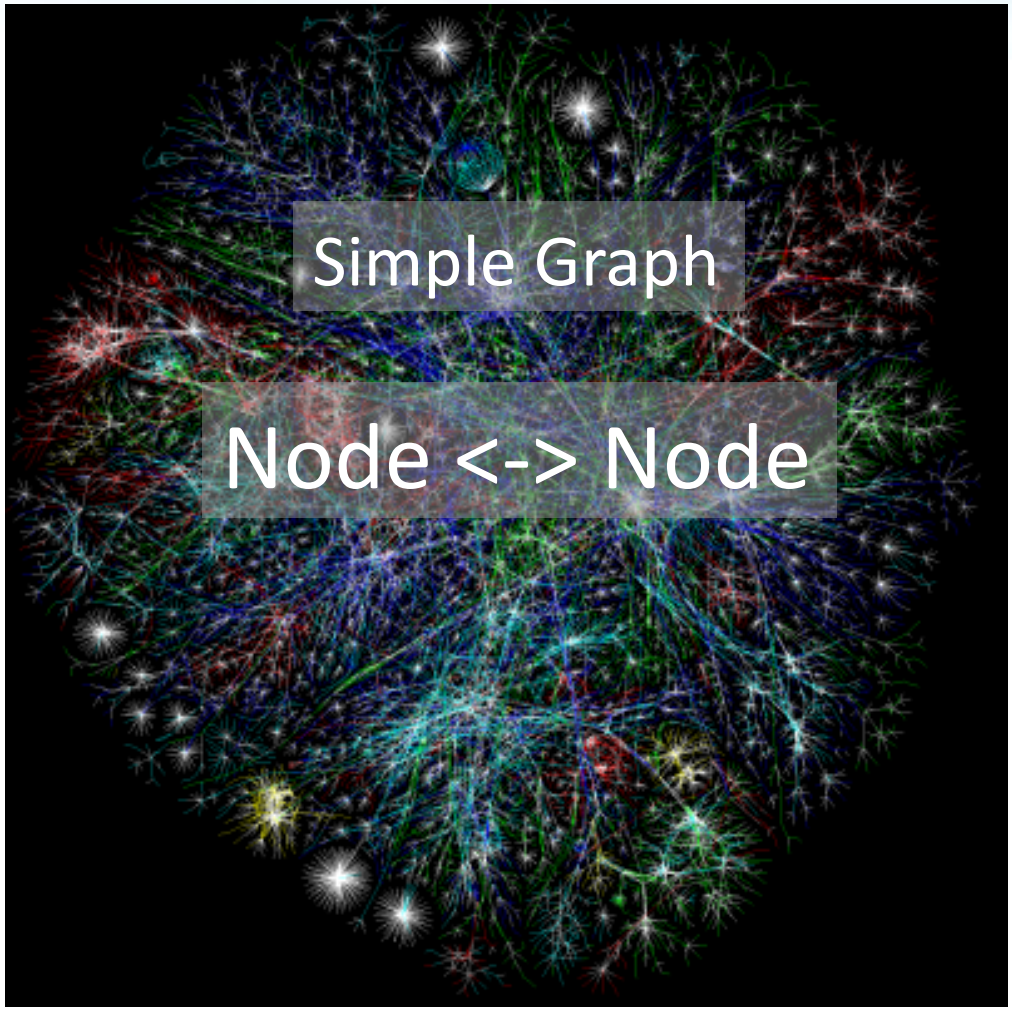

http://opte.org

Objectivity

InfiniteGraph
Objectivity/DB

The Value of Data

- “[finding] Japanese restaurants in New York City liked by people from Japan” – a challenging Facebook query.
- "The value of any piece of information is only known when you can connect it with something else that arrives at a future point in time,” – CIA’s CTO Ira "Gus" Hunt

Connecting Data

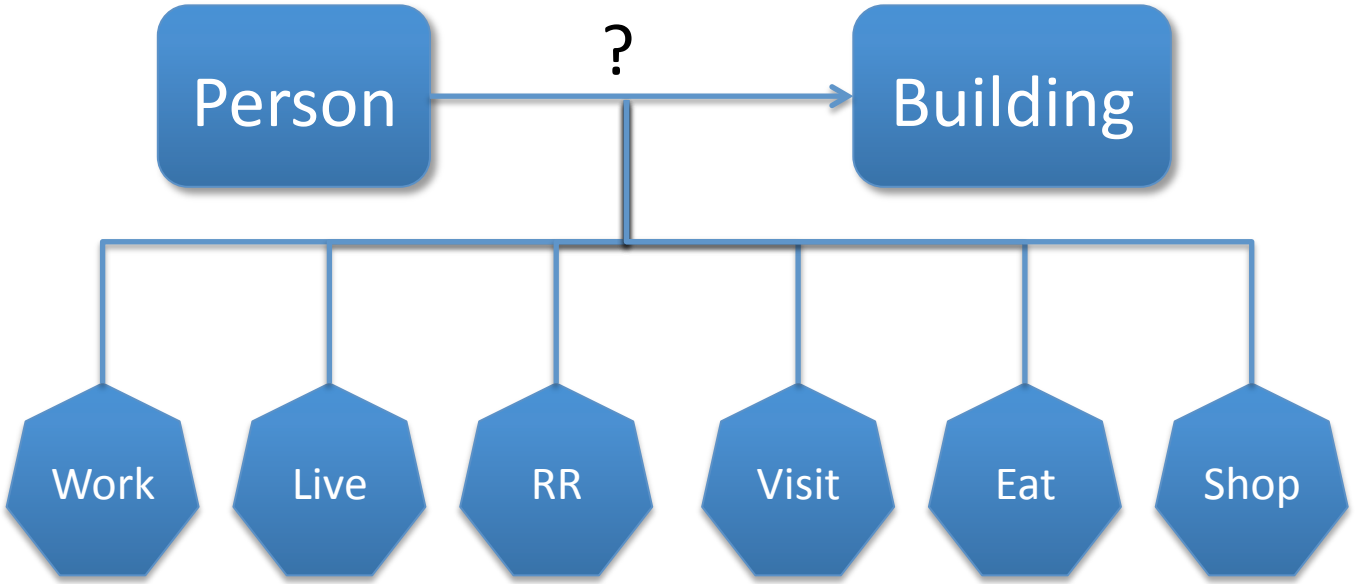

Objectivity

InfiniteGraph
Objectivity/DB

Relationships are everywhere

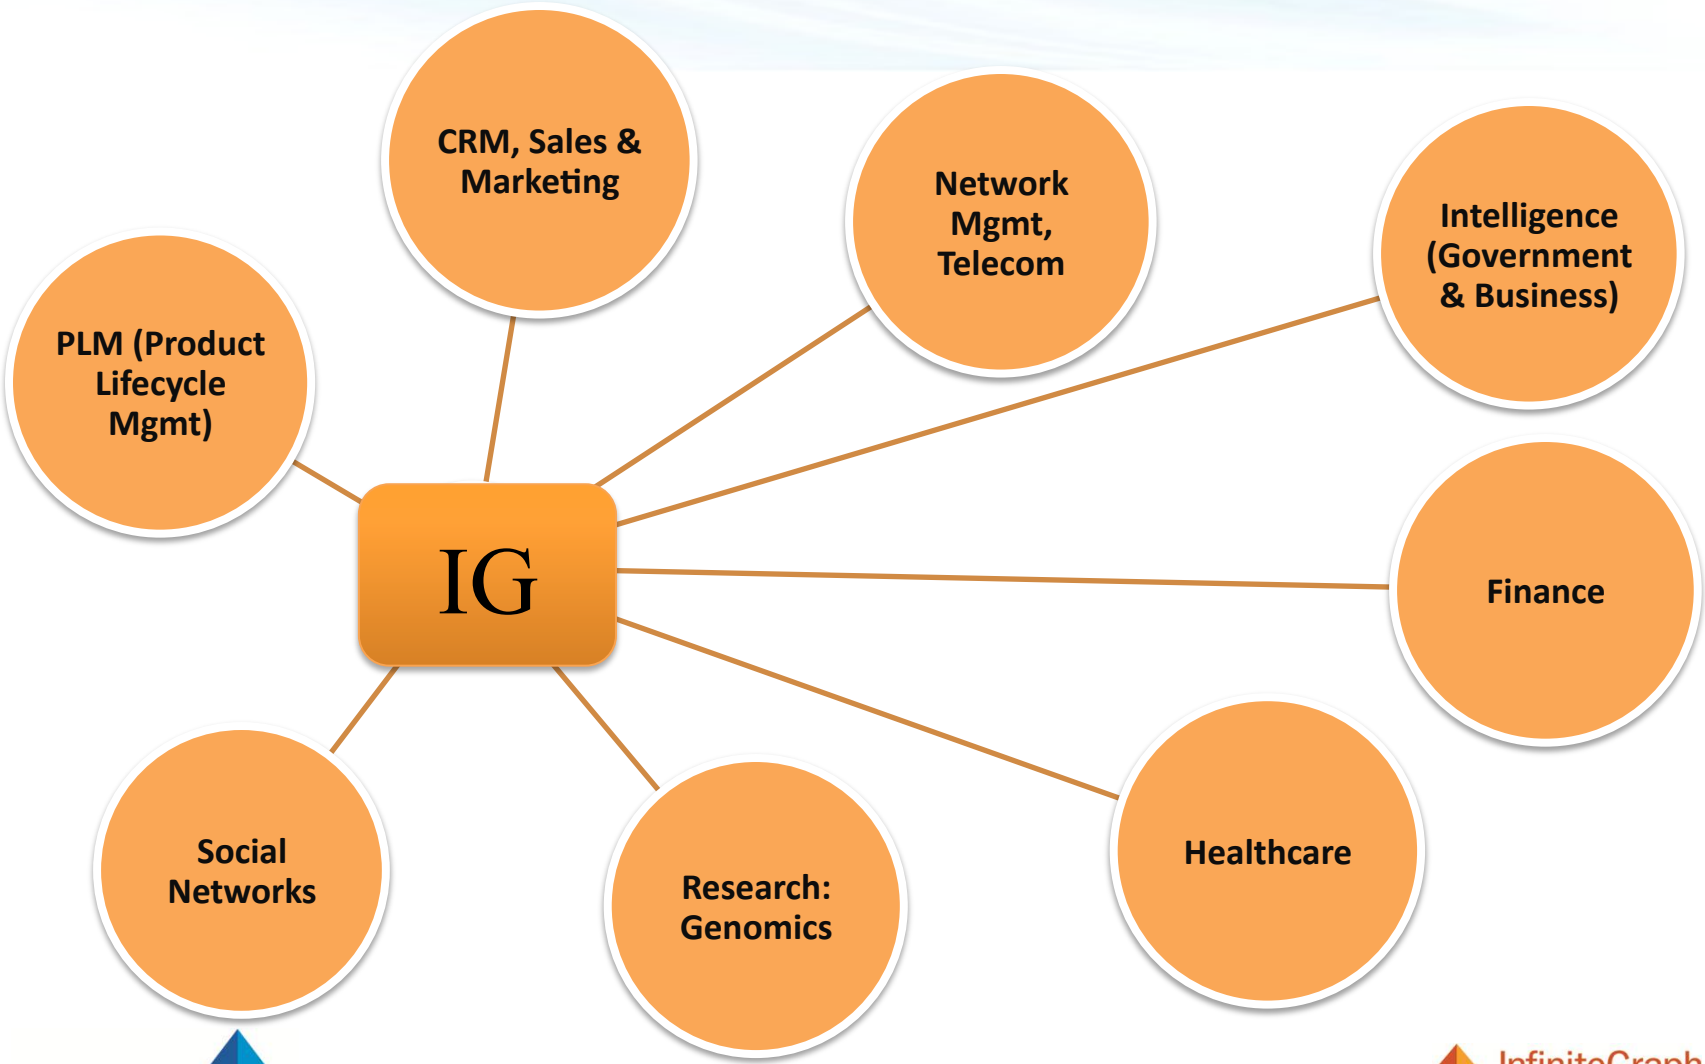

Objectivity

InfiniteGraph
Objectivity/DB

Graph Data

- Data structure for representing complicated relationships between entities.
- Wide applications in…
  - Bioinformatics.
  - Social networks.
  - Network management… etc.

Why a Graph Database ?

Objectivity

InfiniteGraph
Objectivity/DB

The Graph Specialist !

- Everyone specializes
  - Doctors, Lawyers, Bankers, Developers ☺
- Why was data so normalized for so long !
- NoSQL is all about the data specialist
- Specializing in…
  - Distribution / deployment
  - Physical data storage
  - Logical data model
  - Query mechanism

Objectivity

InfiniteGraph
Objectivity/DB

What Problem!

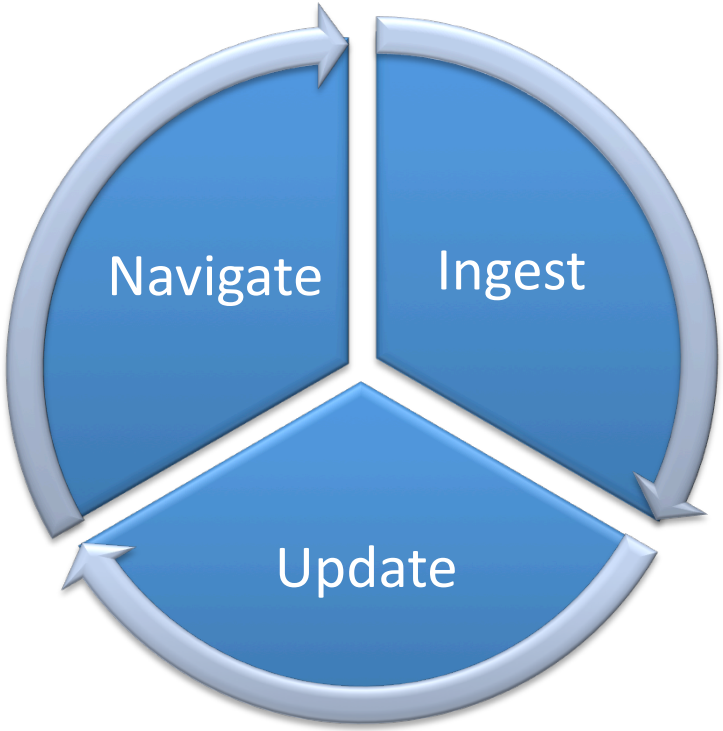

Massive graph data require efficient and intelligent tools to analyze and understand it.

Objectivity

InfiniteGraph
Objectivity/DB

Scaling Writes

- Big/Fast data demands write performance
- Most NoSQL solutions allow you to scale writes by…
  - Partitioning the data
  - Understanding your consistency requirements
  - Allowing you to defer conflicts

Objectivity

InfiniteGraph
Objectivity/DB

Scaling Graph Writes
ACID Transactions

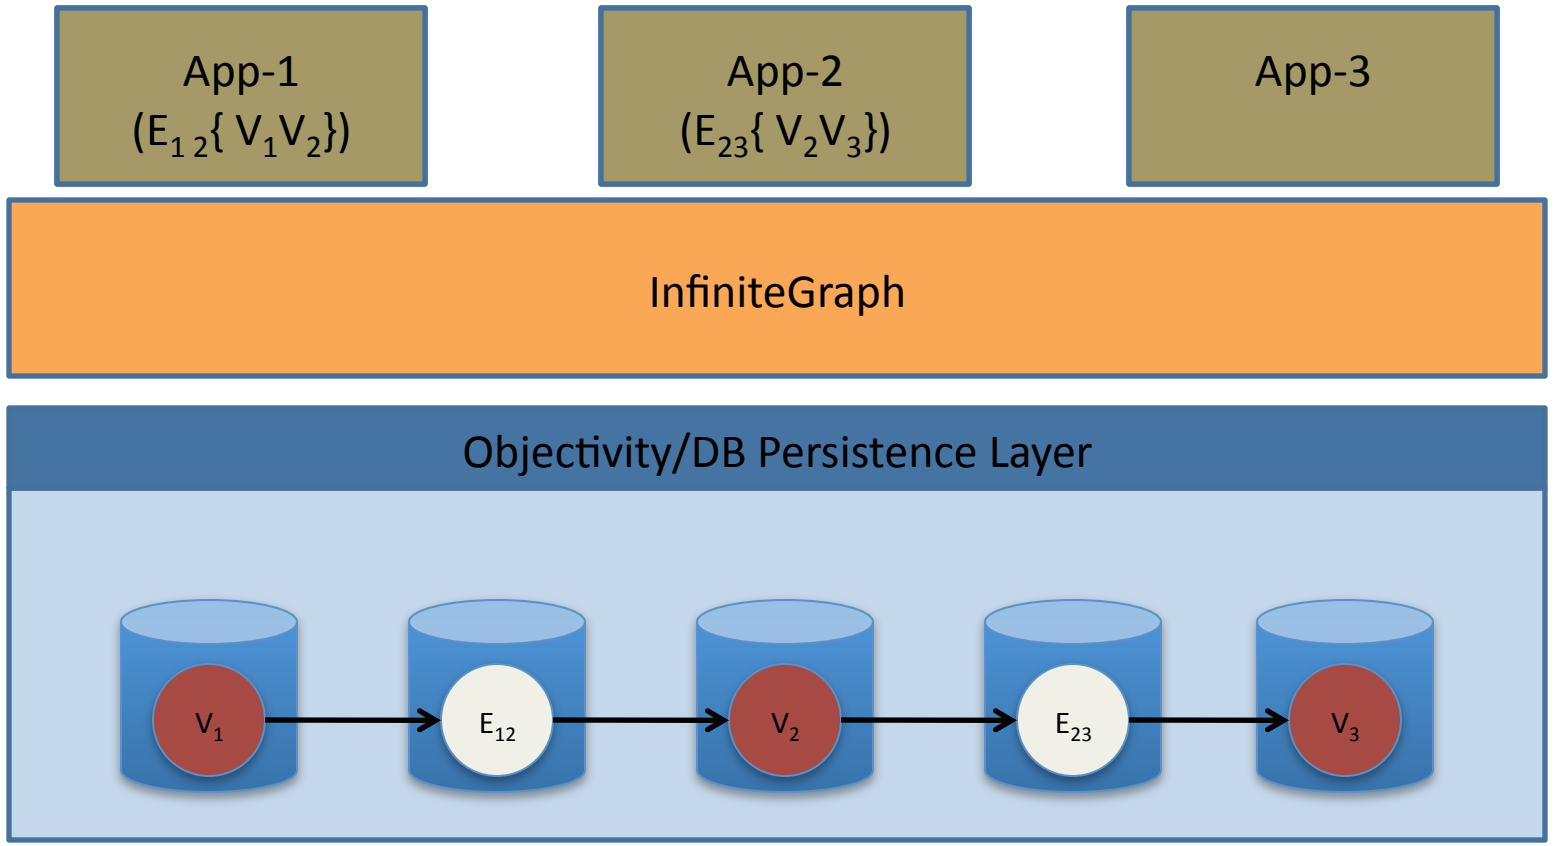

Ingest

Objectivity

InfiniteGraph
Objectivity/DB

High Performance Edge Ingest

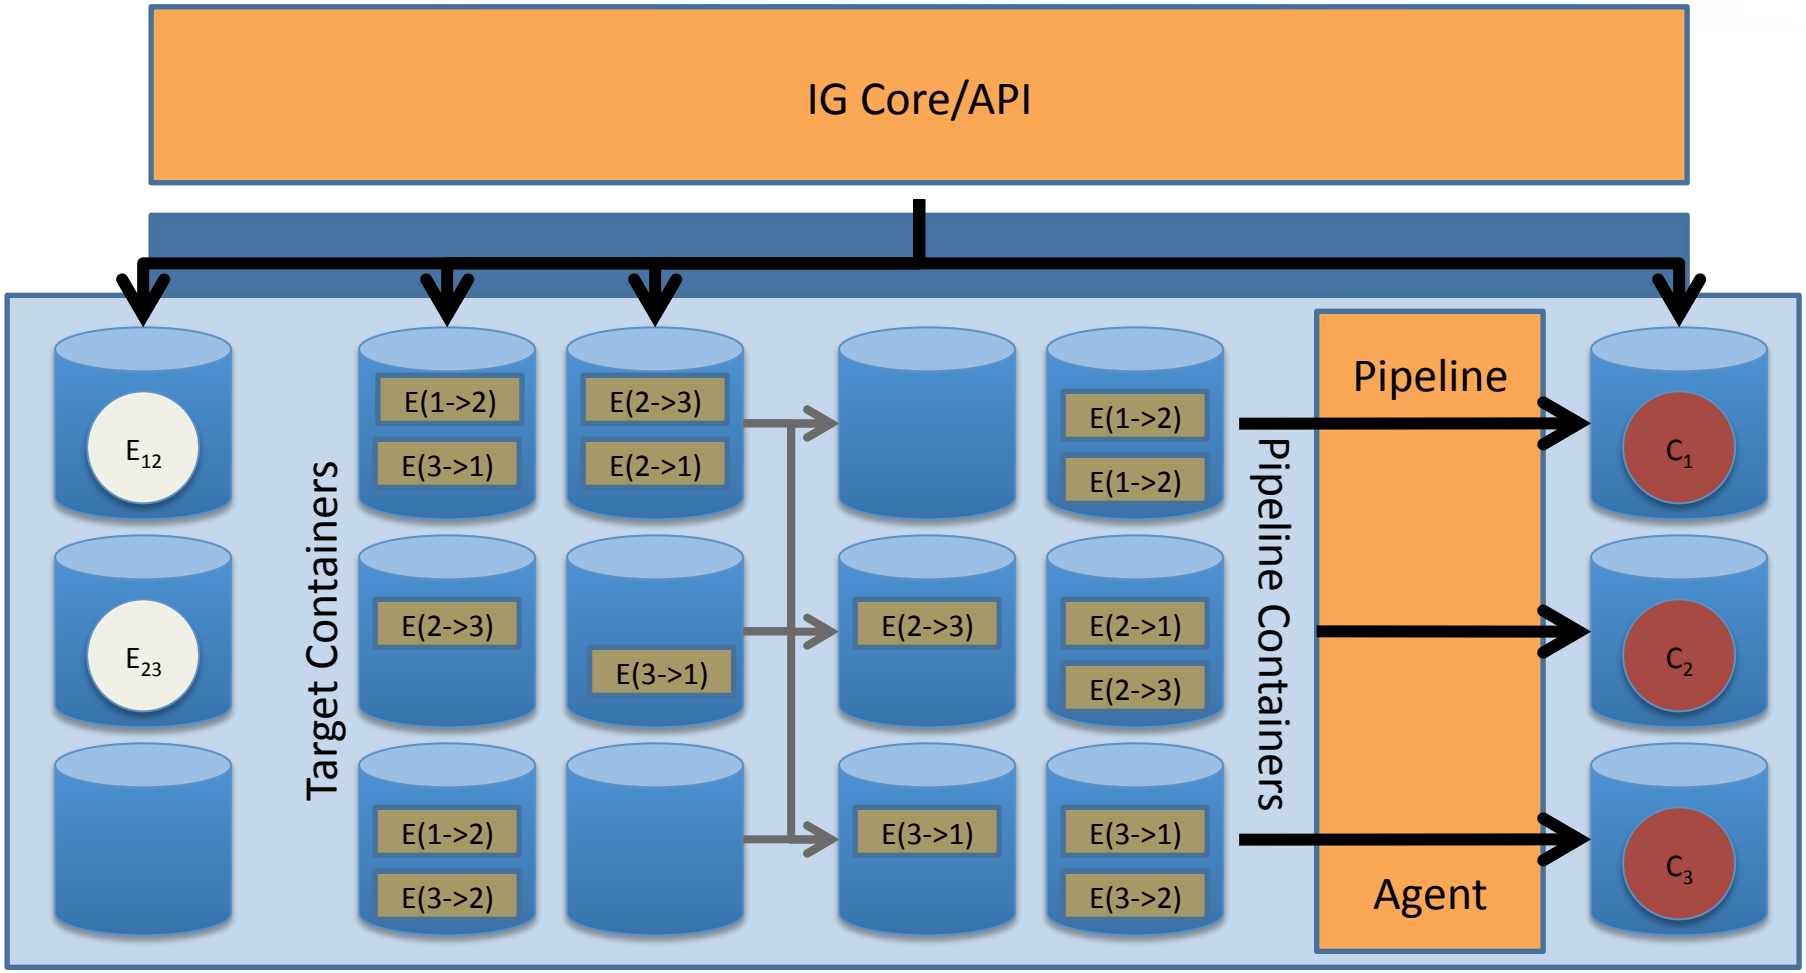

Ingest

Objectivity

InfiniteGraph
Objectivity/DB

Trade offs

- Excellent for efficient use of page cache
- Able to maintain full database consistency
- Achieves highest ingest rate in distributed environments
- Almost always has highest “perceived” rate
- Trading Off :
  - Eventual consistency in graph (connections)
  - Updates are still atomic, isolated and durable but phased
  - External agent performs graph building

Objectivity

InfiniteGraph
Objectivity/DB

Result…

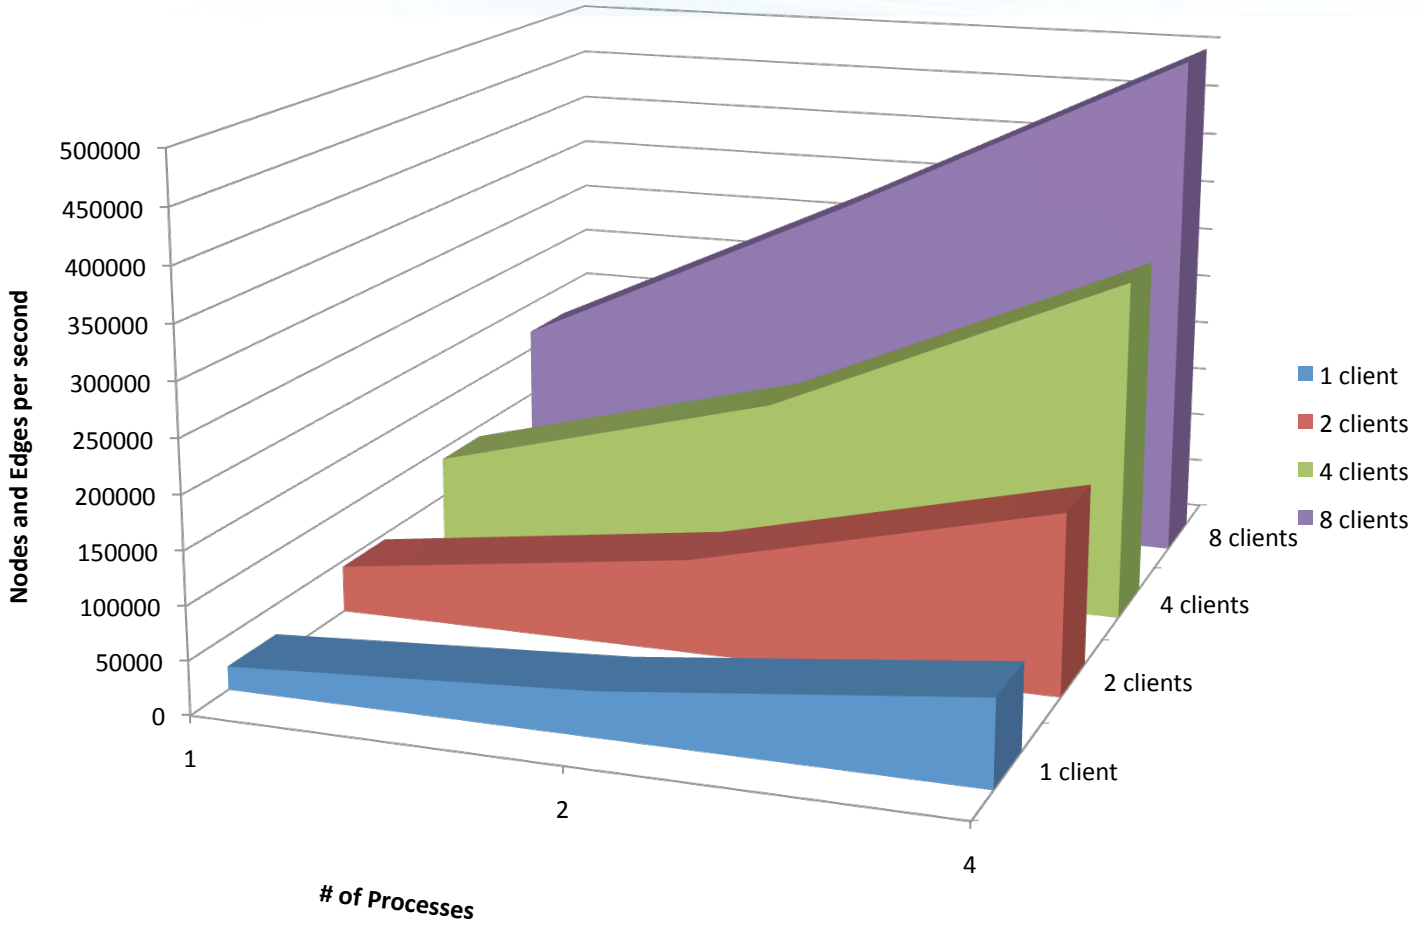

Objectivity

InfiniteGraph
Objectivity/DB

Scaling Reads and Query

Partitioning and Read Replicas… easy right !

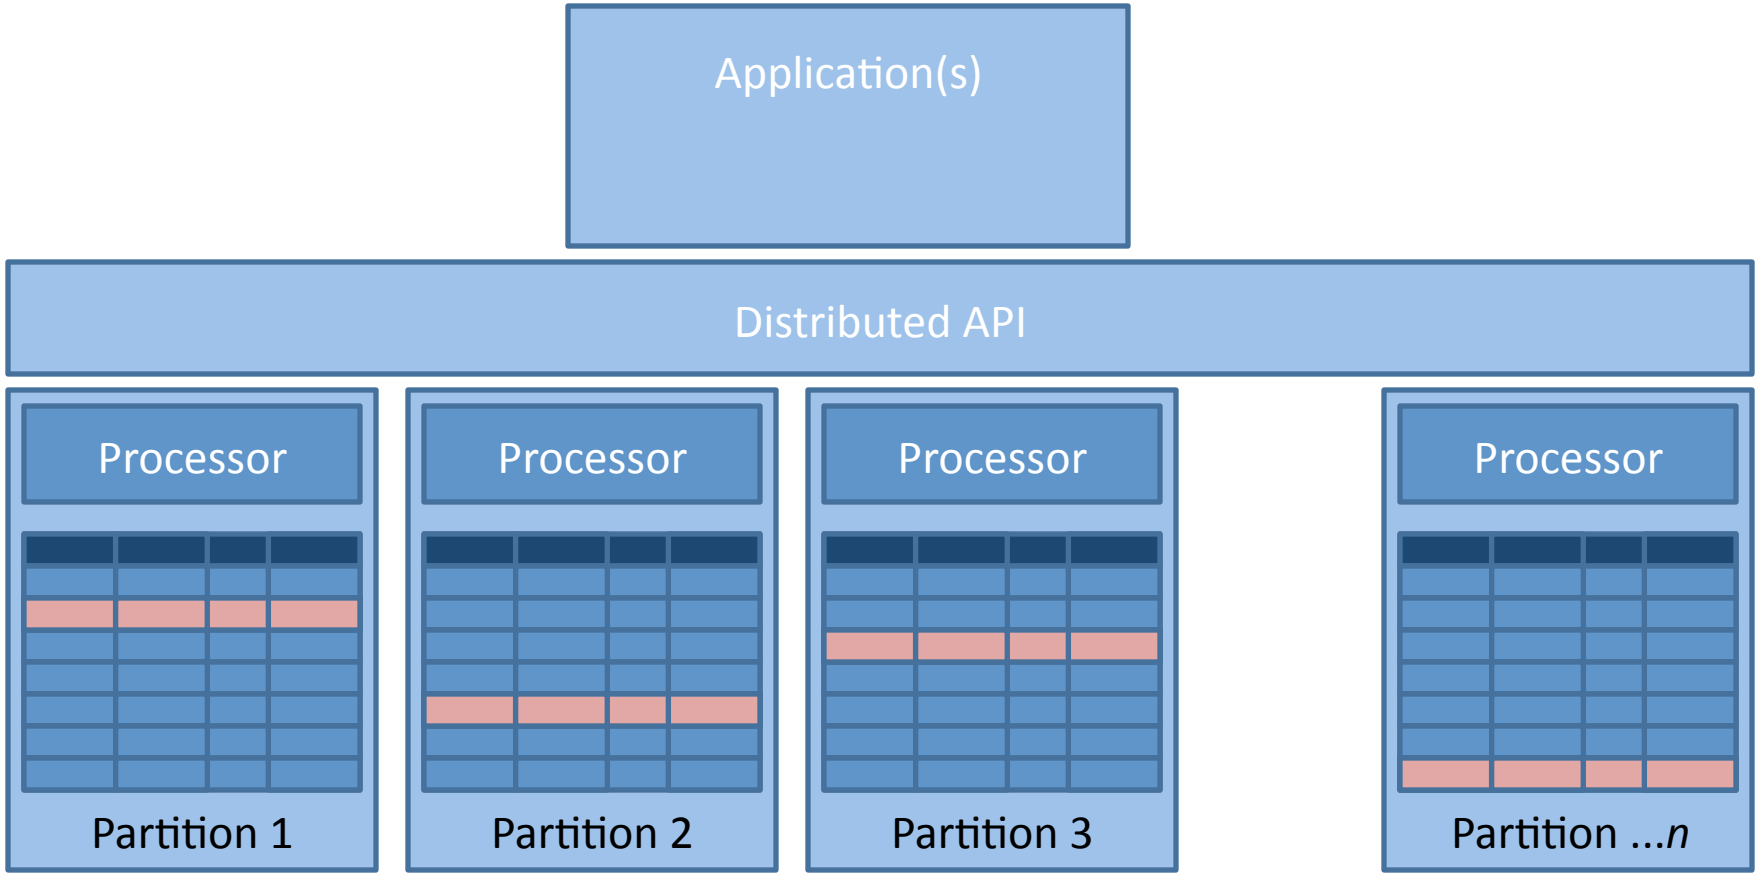

Objectivity

InfiniteGraph
Objectivity/DB

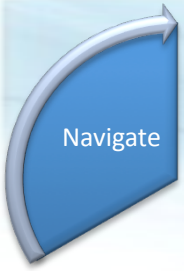

Why are Graphs Different ?

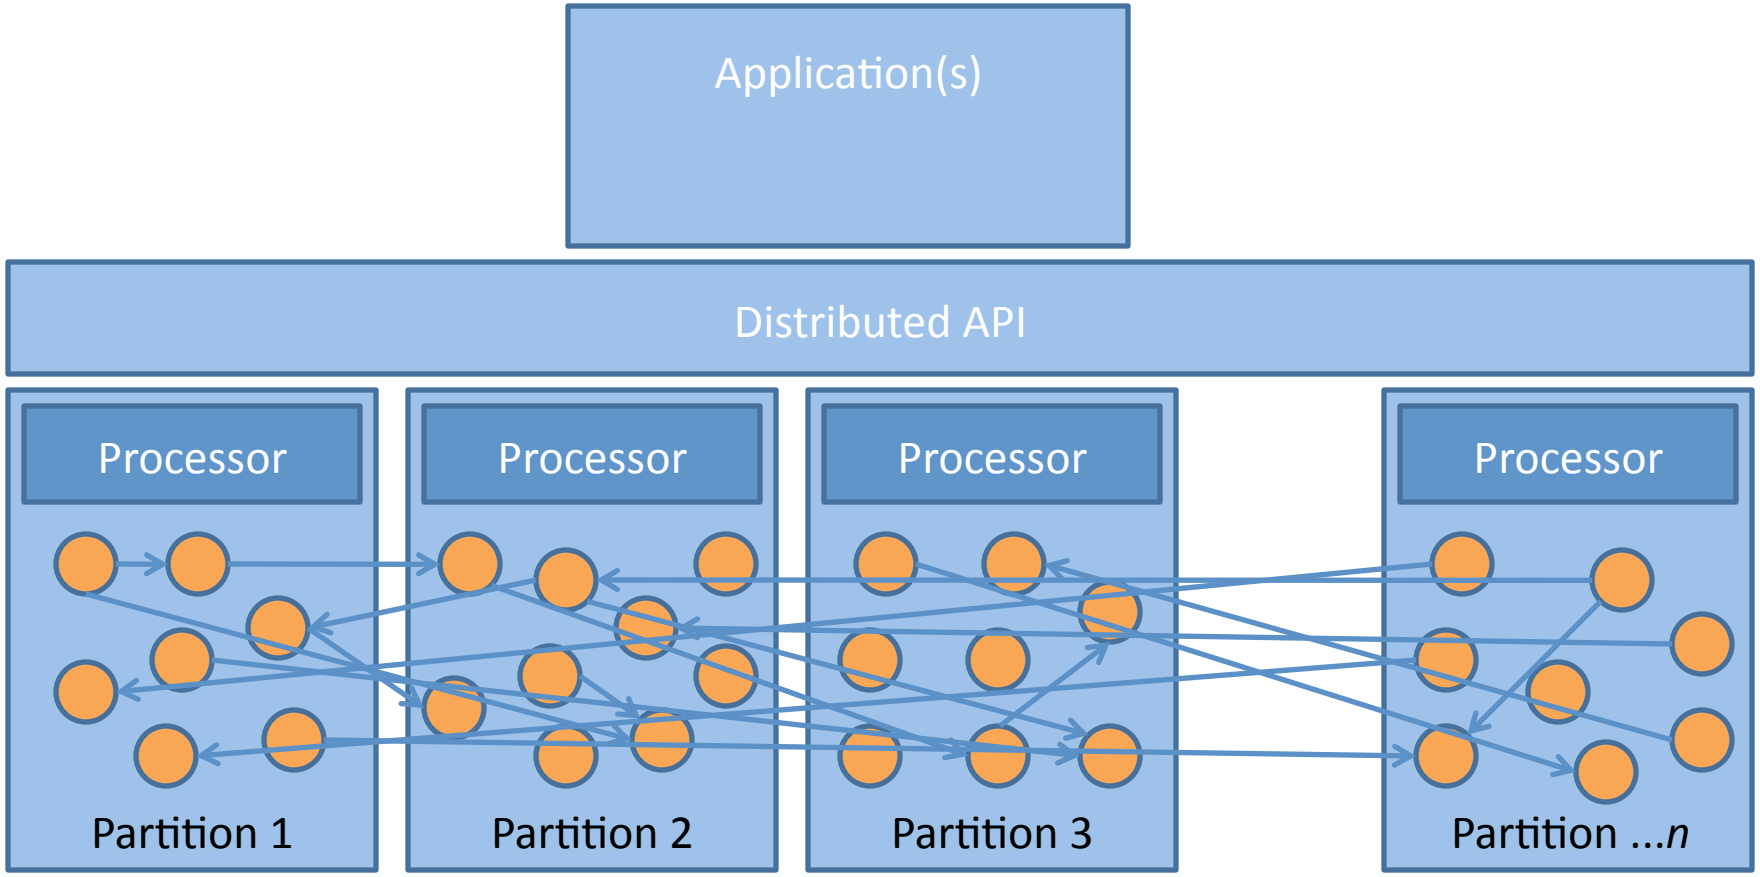

Objectivity

InfiniteGraph
Objectivity/DB

Distributed Navigation

- Detect local hops and perform in memory traversal
- Send the partial path to the distributed processing to continue the navigation.
- Intelligently cache remote data when accessed frequently
    - Route tasks to other hosts when it is optimal

Objectivity

InfiniteGraph
Objectivity/DB

Distributed Navigation Server

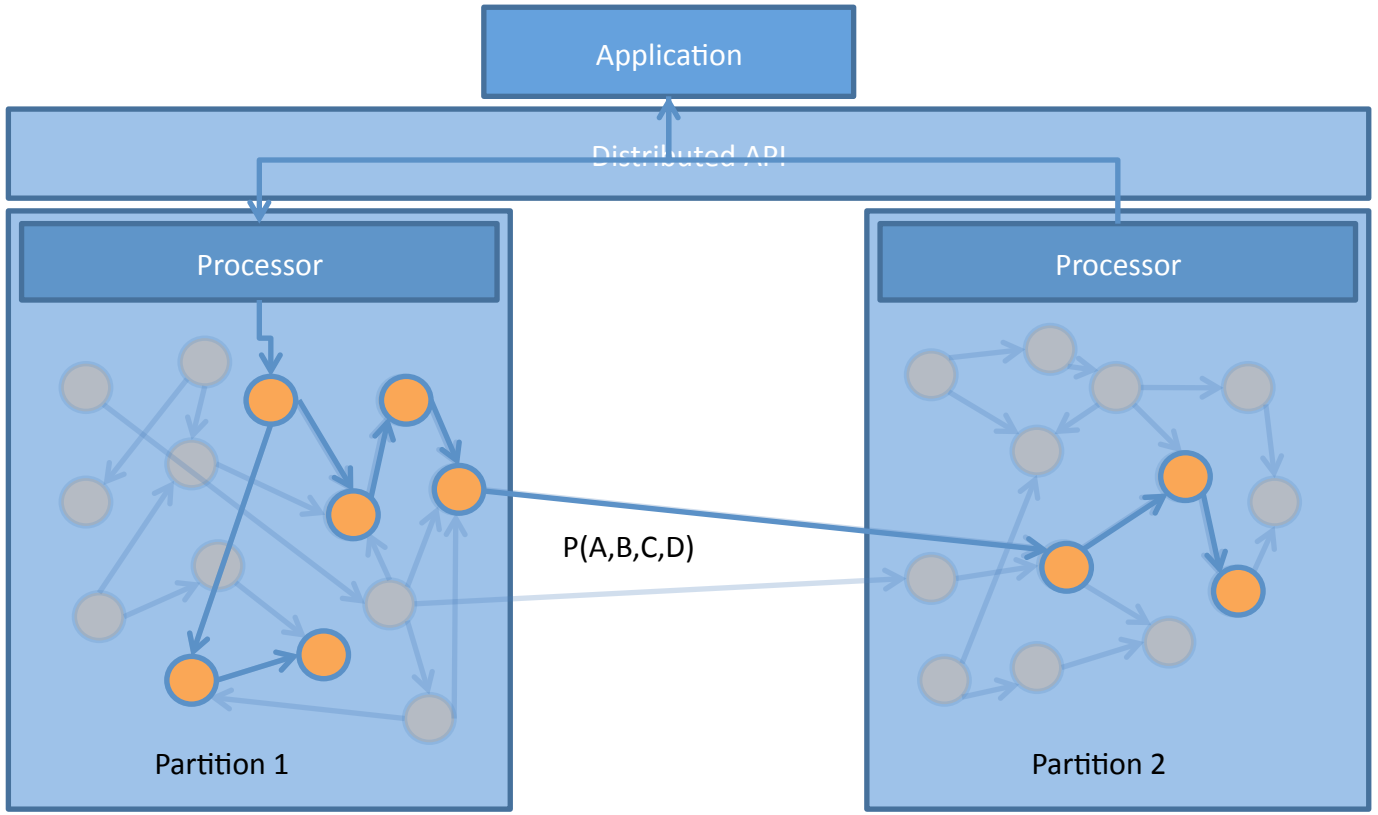

Objectivity

InfiniteGraph
Objectivity/DB

Schema – It's not your enemy !

(at least not all the time…)

- Schema vs Schema-less
  - Database religion
  - InfiniteGraph supports schema, but does not restrict connection types between vertices
  - We take advantage of the schema when pruning.
  - …Working on a hybrid support.

Objectivity

InfiniteGraph
Objectivity/DB

Navigate

GraphViews
Leveraging Schema in the Graph

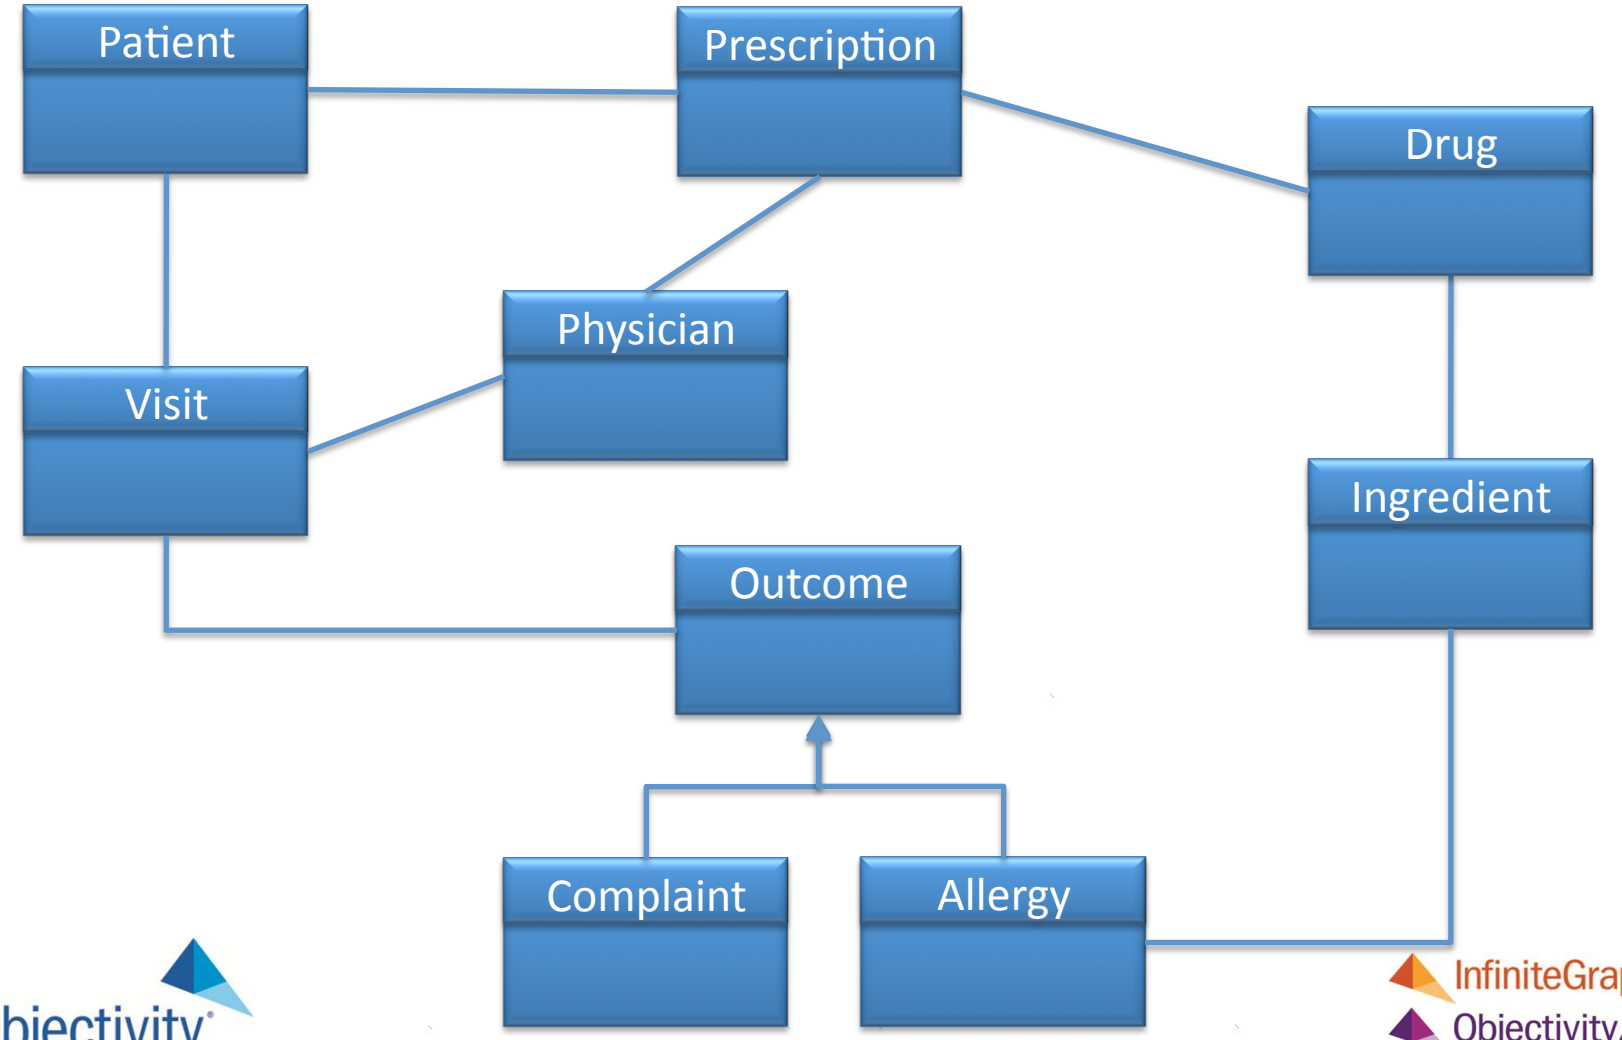

Objectivity

InfiniteGraph
Objectivity/DB

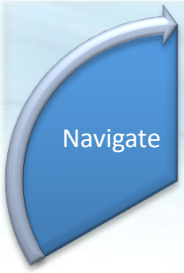

Schema Enables Views

- GraphViews are extremely powerful
- Allow Big Data to appear small !
- Connection **inference** can lead to exponential gains in query performance
- Views are reusable between queries
- Built into the native kernel

Why InfiniteGraph™?

- Objectivity/DB is a proven foundation
  - Building distributed databases since 1993
  - A complete database management system
    - Concurrency, transactions, cache, schema, query, indexing
- It's a Graph Specialist !
  - Simple but powerful API tailored for data navigation.
  - Easy to configure distribution model

Objectivity

InfiniteGraph
Objectivity/DB

Advanced Configured Placement

- Physically co-locate "closely related" data
- Driven through a declarative placement model
- Dramatically speeds "local" reads

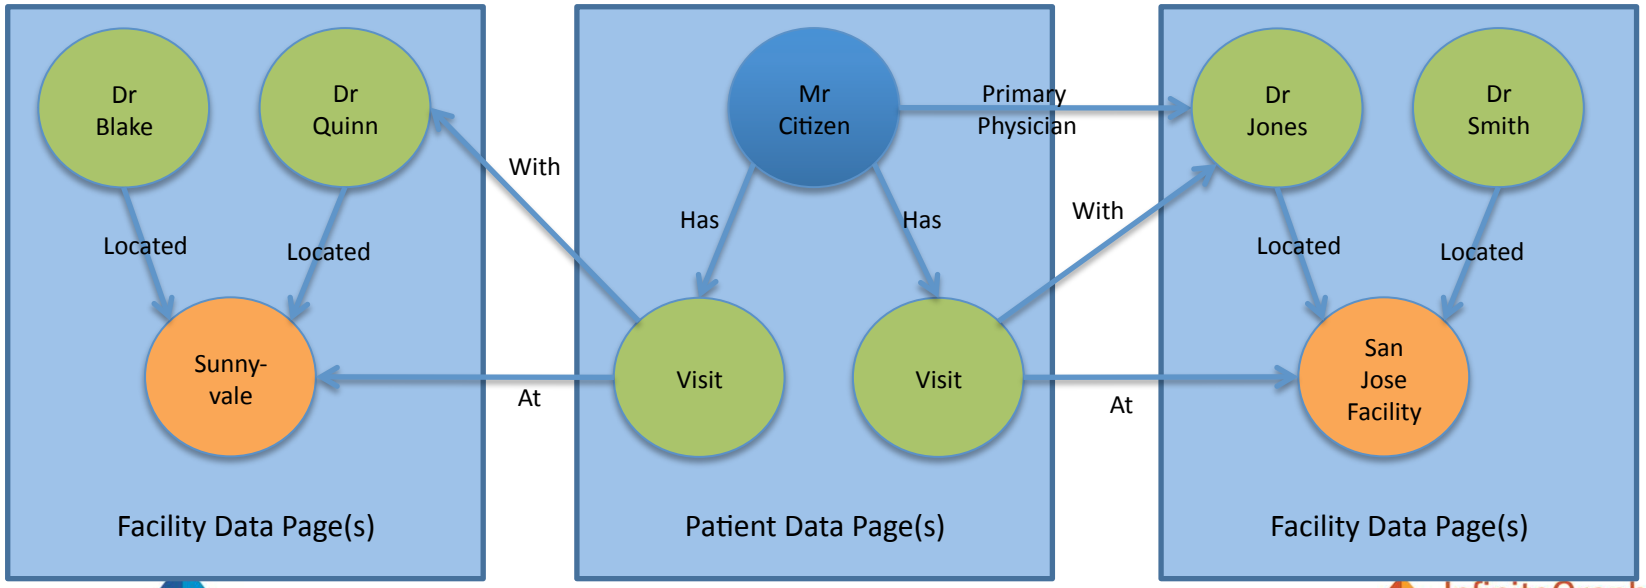

Objectivity

InfiniteGraph
Objectivity/DB

Fully Distributed Data Model

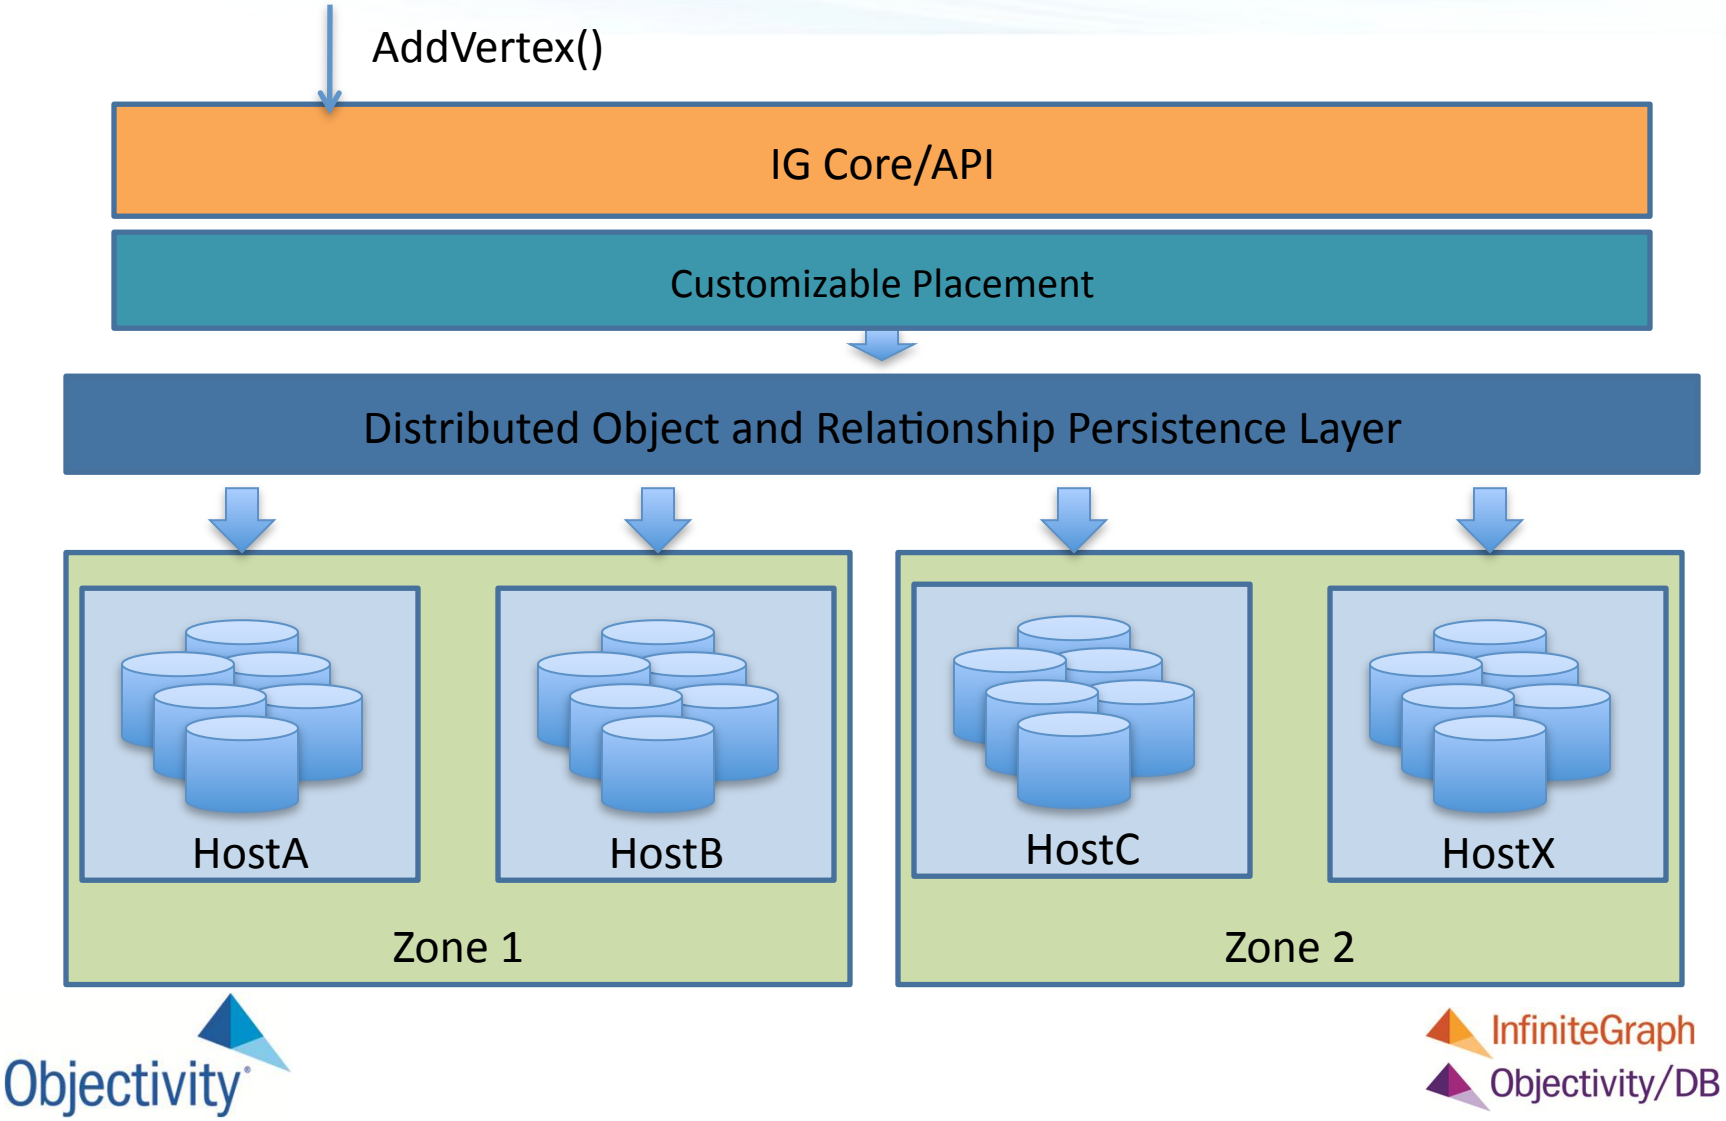

Polyglot NoSQL Architectures

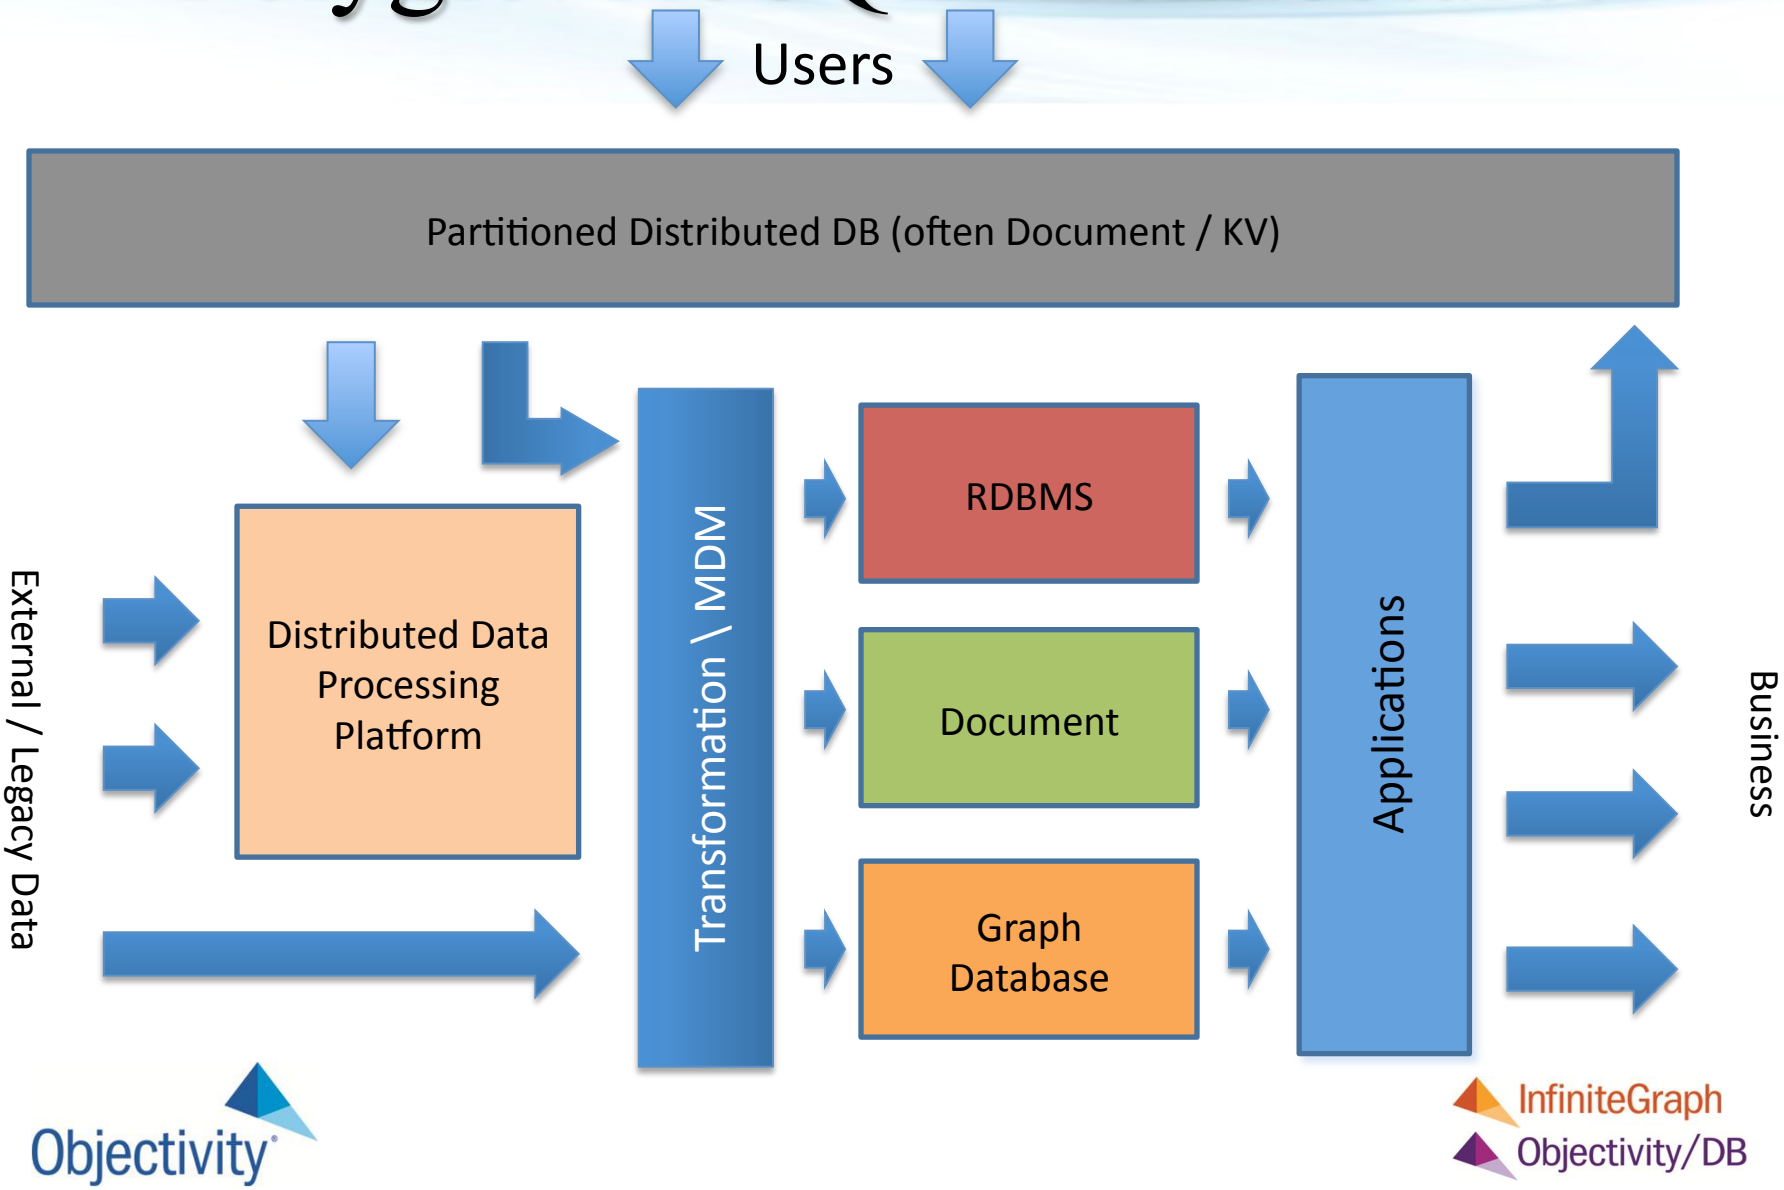

What else!

- Distributed update.

… we are working on it.

Objectivity

InfiniteGraph
Objectivity/DB

Conclusion

Have we solved all the distributed graph problems?

Not quite, but…

We Learned to Stop Worrying.

Objectivity

InfiniteGraph
Objectivity/DB

QUESTIONS?