ICOODB 2010 - Frankfurt, Deutschland

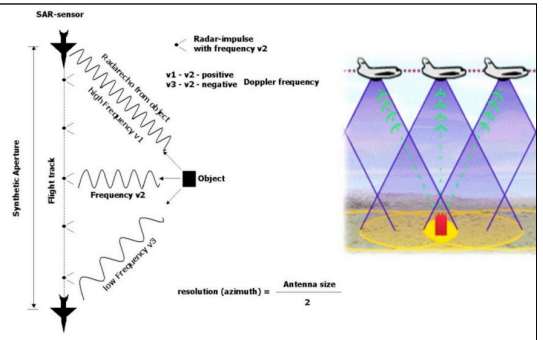

The Synergy Between the Object Database, Graph Database, Cloud Computing and NoSQL Paradigms

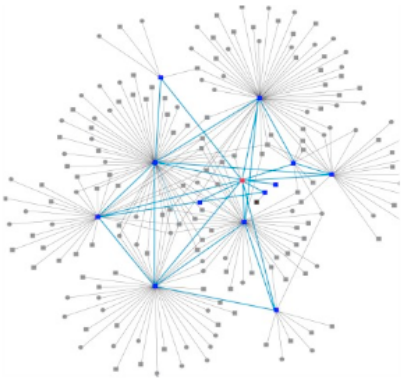

Leon Guzenda - Objectivity, Inc.

Objectivity

AGENDA

- Historical Overview
- Inherent Advantages of ODBMSs
- Technology Evolution
- Leveraging Technologies
- Graph Databases
- Summary

Historical Overview

The ODBMS Players

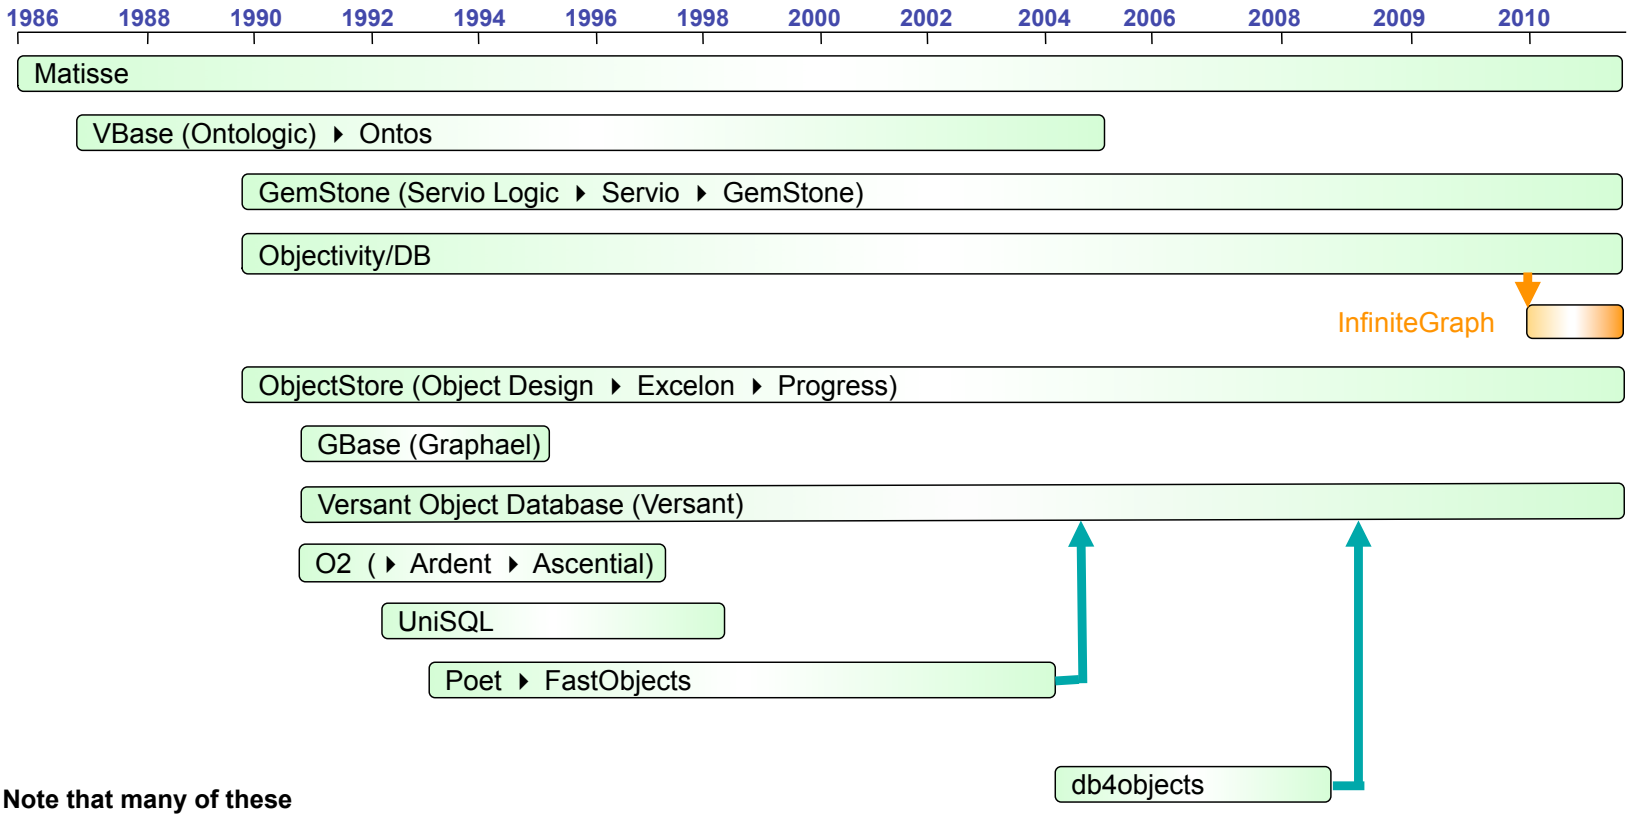

Note that many of these companies were founded earlier, e.g. Objectivity, Inc. was founded in June 1988.

ODBMS Evolution

- 1980s
  - "Performance, Performance, Performance!"
  - Primarily scientific and engineering applications
- 1990s
  - Reliability and Scalability
  - New languages and Operating Systems
  - Large deployments in the scientific domain
- 2000s
  - Ease of use and instrumentation
  - Query languages
  - Performance and scalability
  - Grids and Clouds
  - Embedded systems, government and more...

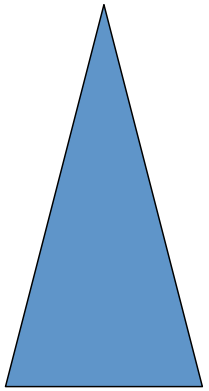

DATA MANIPULATION

Applications tended to generate data and relationships

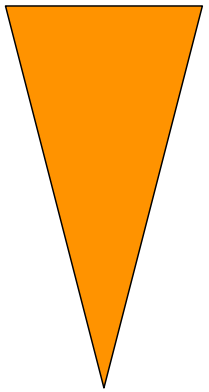

RELATIONSHIP ANALYTICS

Applications ingest and correlate data and relationships

Inherent ODBMS Advantages

Faster Navigation

PROBLEM:
Find all of the Suspects
linked to a chosen Incident

Relational solution:

N * 2 B-Tree lookups

N * 2 logical reads

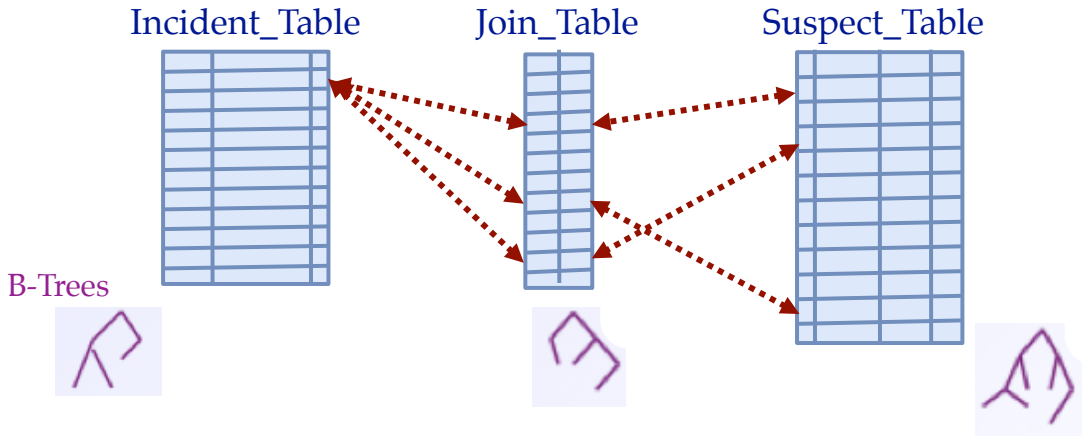

Objectivity/DB solution:

1 B-Tree lookup

1 + N logical reads

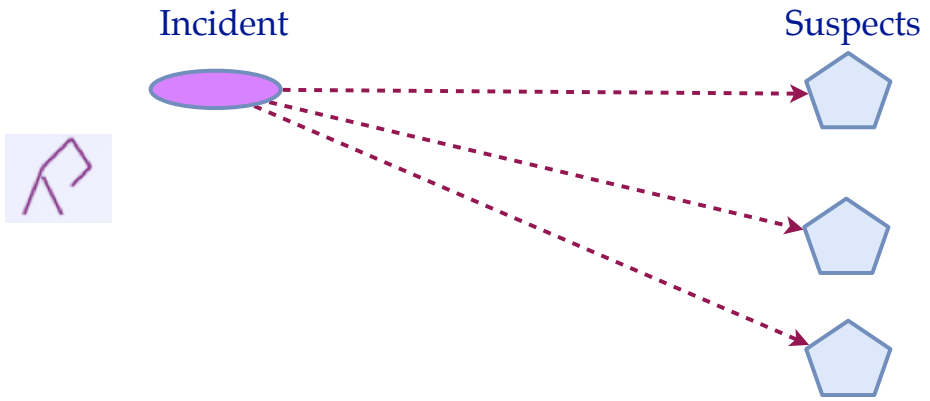

Significantly faster navigation of relationships

Lower Query Latency

Relational

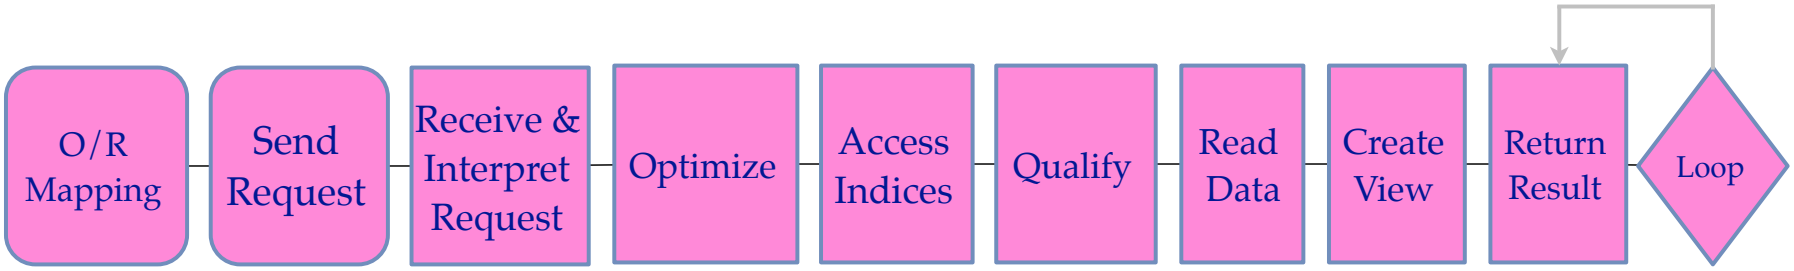

Objectivity/DB

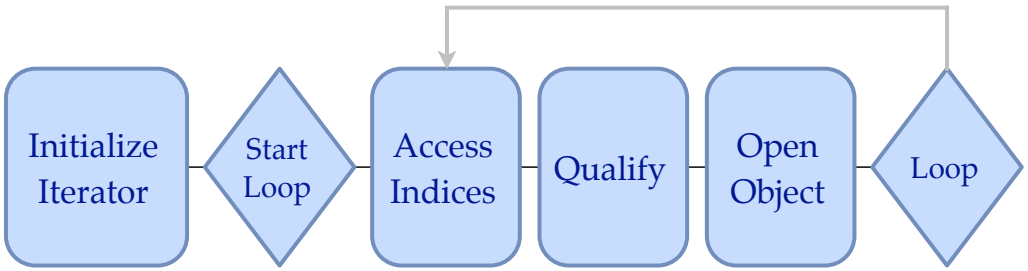

Qualified objects are returned as soon as they are found.

Technology Evolution

Grids and Clouds

- 1996 - CERN started looking for a DBMS for the LHC.
- RD45 team verified that a distributed ODBMS could handle complexity, performance and Petabyte+ scale.
- Lead to grid deployments at SLAC and Brookhaven
- Meanwhile, developers migrated from CORBA to SOA.
- Then, our grid and SOA experience made the migration to cloud environments very easy.

The "NoSQL" Movement

- Some web application developers found RDBMSs too restrictive.
- They were dealing with:
  - Huge parallel ingest streams
  - Applications that scan or navigate rather than query
  - Unconventional transaction models.
- So... they re-invented the wheel!
  - Sharding (Hadoop and Big Table)
  - Key-Value tables (Big Table)
  - Dynamo...

Sharding

Splits large tables into groups of related rows and puts them onto separate servers. They each have a schema.

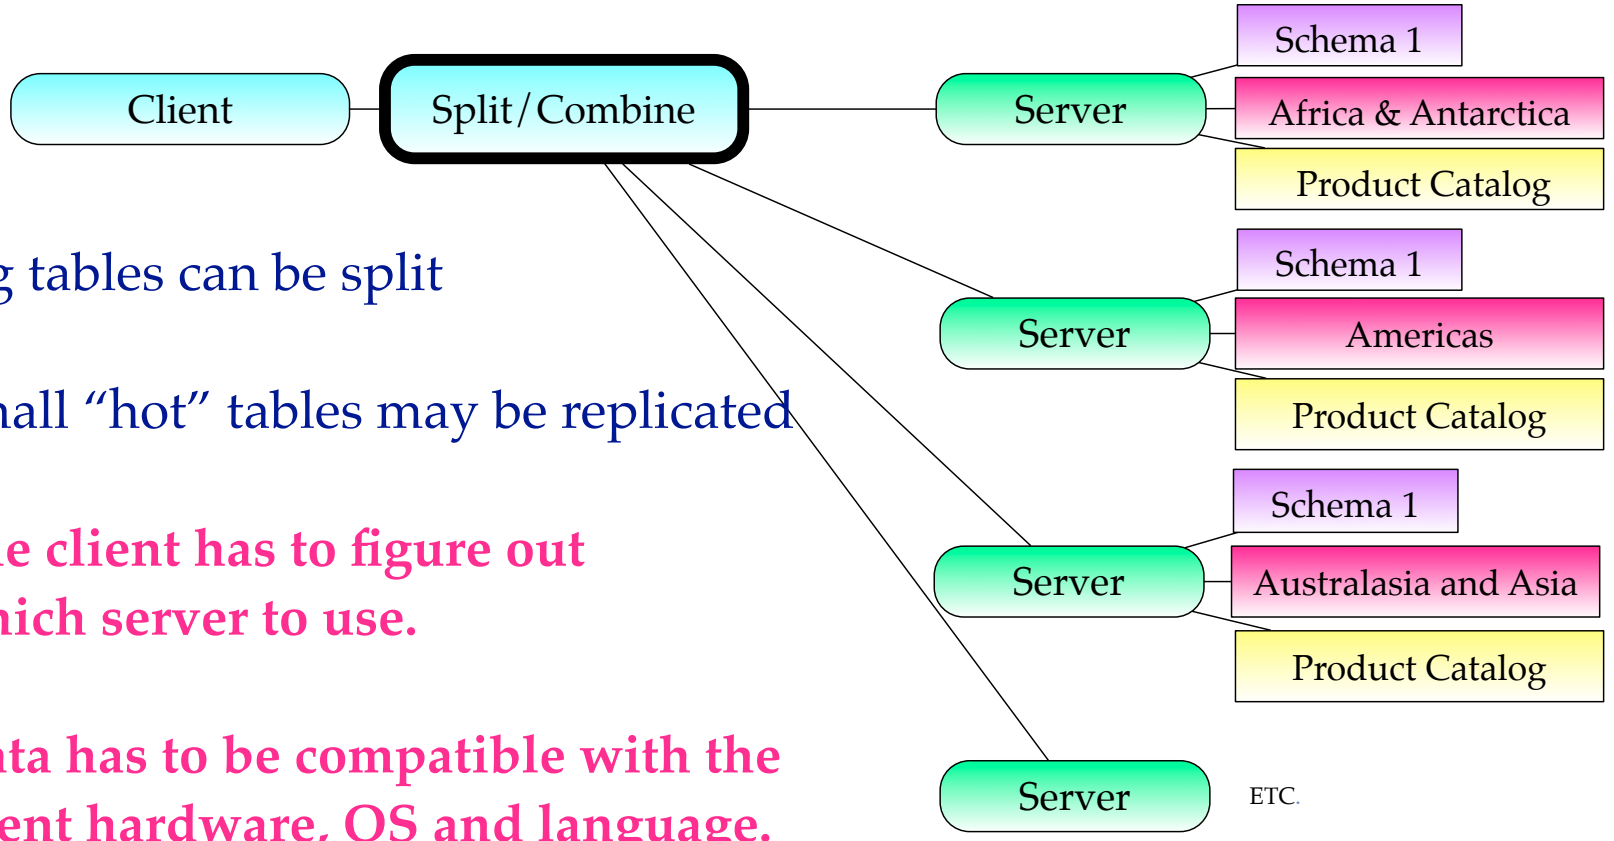

- Big tables can be split
- Small "hot" tables may be replicated
- **The client has to figure out which server to use.**
- **Data has to be compatible with the client hardware, OS and language.**

Hadoop's HDFS

- Hadoop Distributed File System (HDFS) is the primary storage system used by Hadoop applications.
- HDFS creates multiple replicas of **64+ Megabyte data blocks** and distributes them on compute nodes throughout a cluster to enable reliable, extremely rapid computations.
- Each data block is replicated 3 times - twice on the same rack and once on another rack.

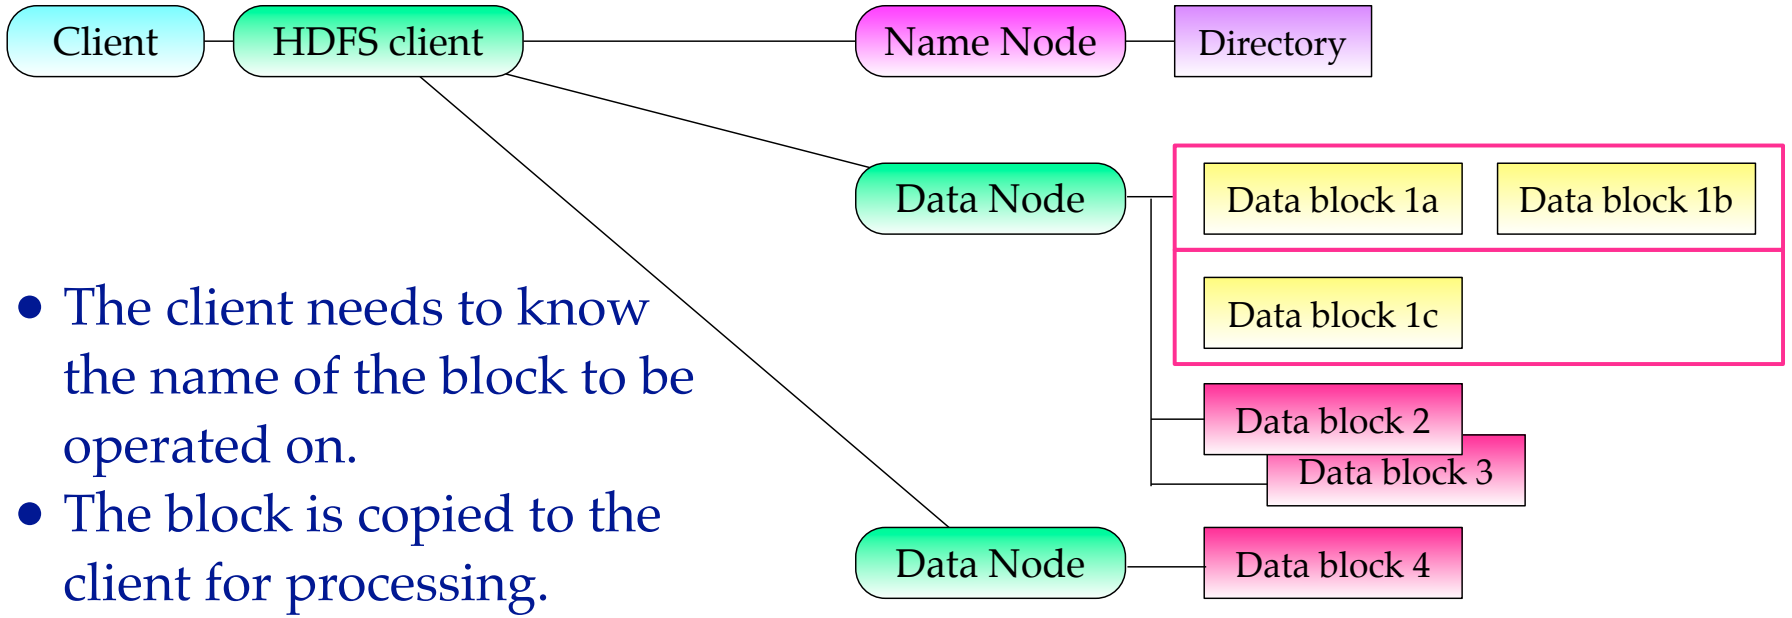

- The client needs to know the name of the block to be operated on.
- The block is copied to the client for processing.

Leveraging Technologies

Objectivity/DB

- Federated object-oriented database platform
- Single Logical View across distributed persistent objects
- Eliminates the OO language to database mapping layer
- Ultra-fast object navigation
- Customizable, distributed query engine
- Dynamic schema changes and low administration overheads
- Interoperability across multiple languages and platforms
- High Availability
- Grid and cloud environment enabled

Distributed Architecture

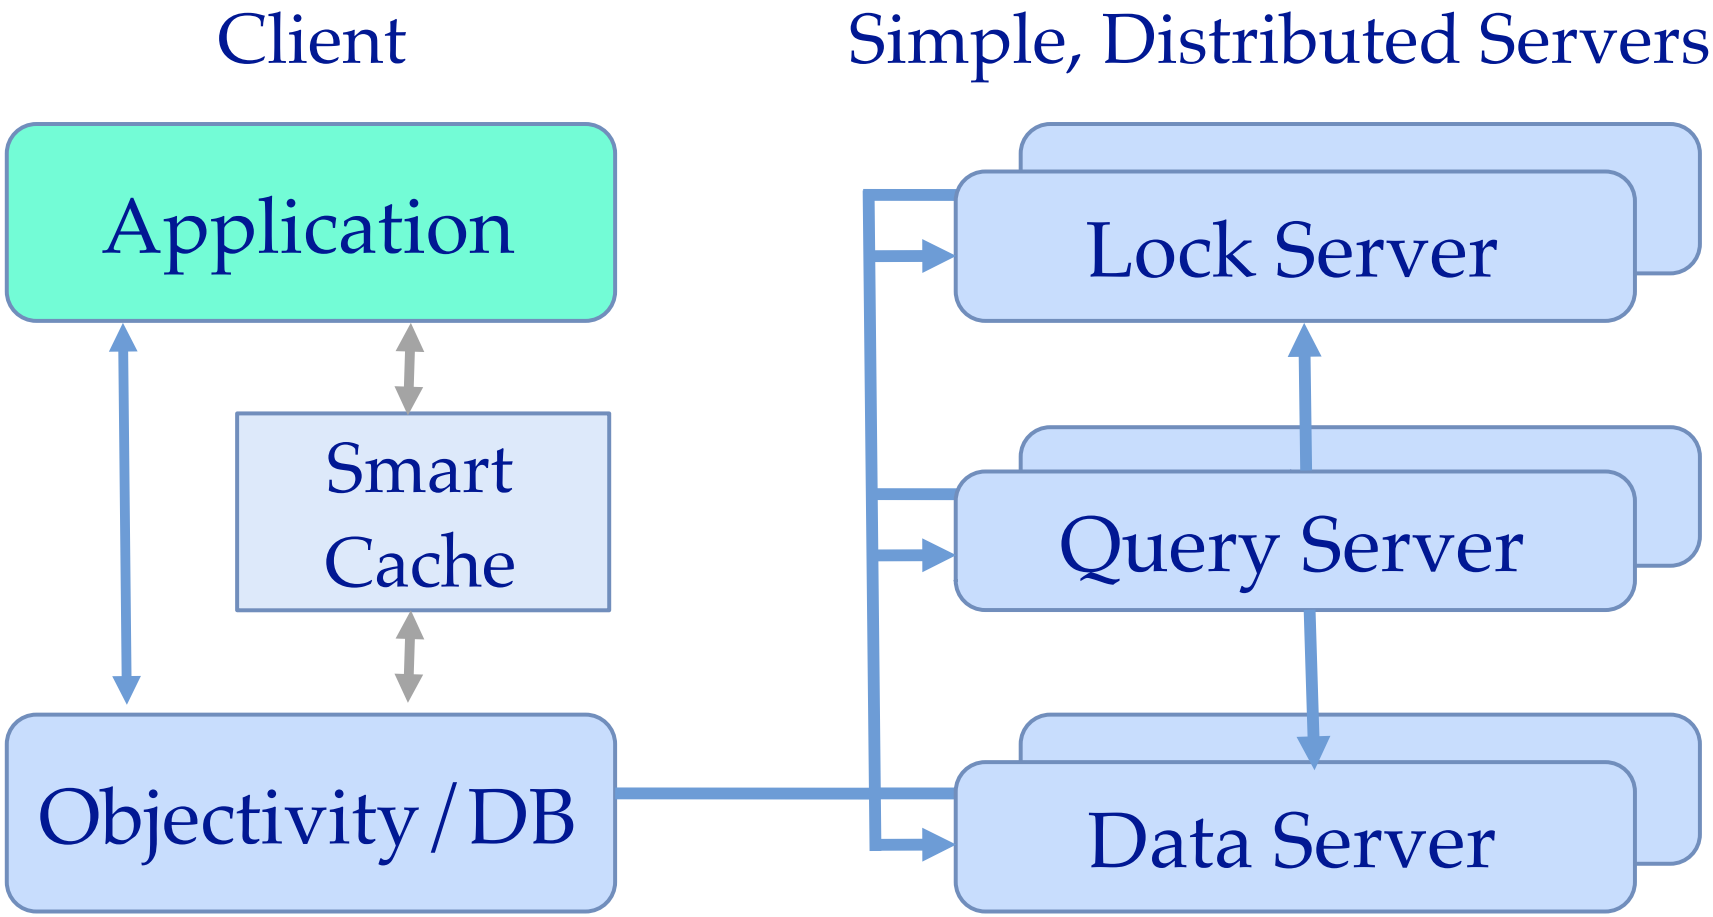

Enhances scalability and availability

Parallel Query Engine

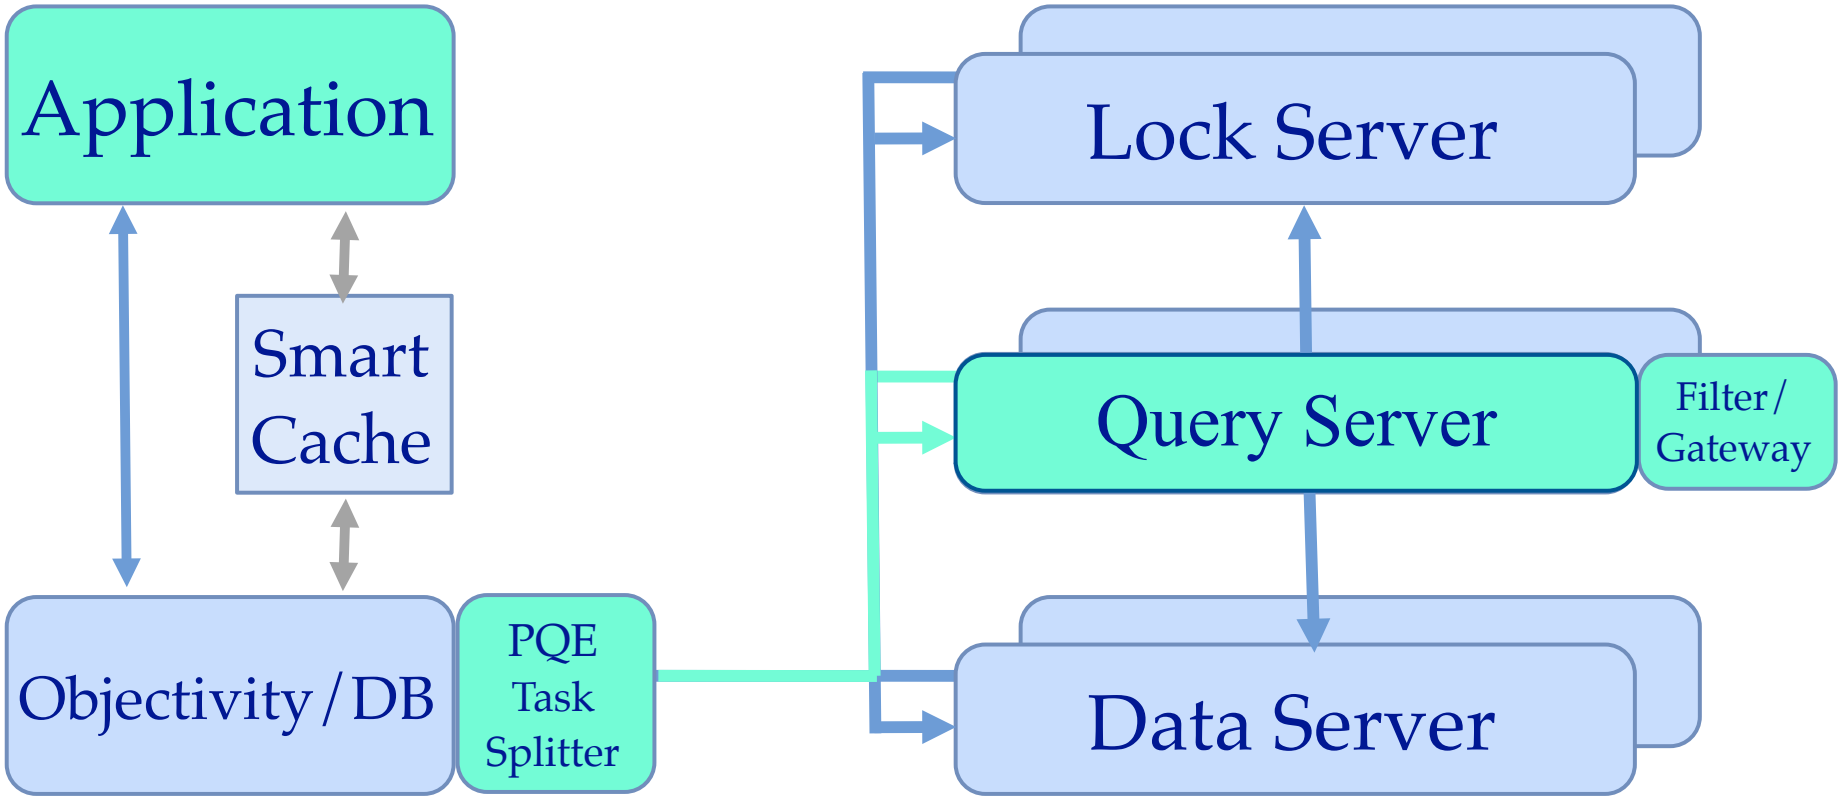

The Task Splitter aims queries at specific databases and containers

Filters can run complex qualification methods. Gateways can access other databases or search engines.

Replaceable components for smarter optimization

Objectivity/DB Advantages

- Fully Distributed with Client-Side Smart Caching
- Highly efficient storage and navigation of relationships
- Flexible Clustering
- Scalable Collections
- Customizable Parallel Query Engine
- Quorum-based Replication and High Availability
- Flexible, Multi-mode Indexing
- Fully Interoperable Across Platforms and Languages

Objectivity/DB 10.1

- User-replaceable Parallel Query Engine search agents
- Page level and partial backups
- Eclipse RCP
- Visual Studio 2010 support
- Mac OS X support

Graph Databases

The Link Hunter

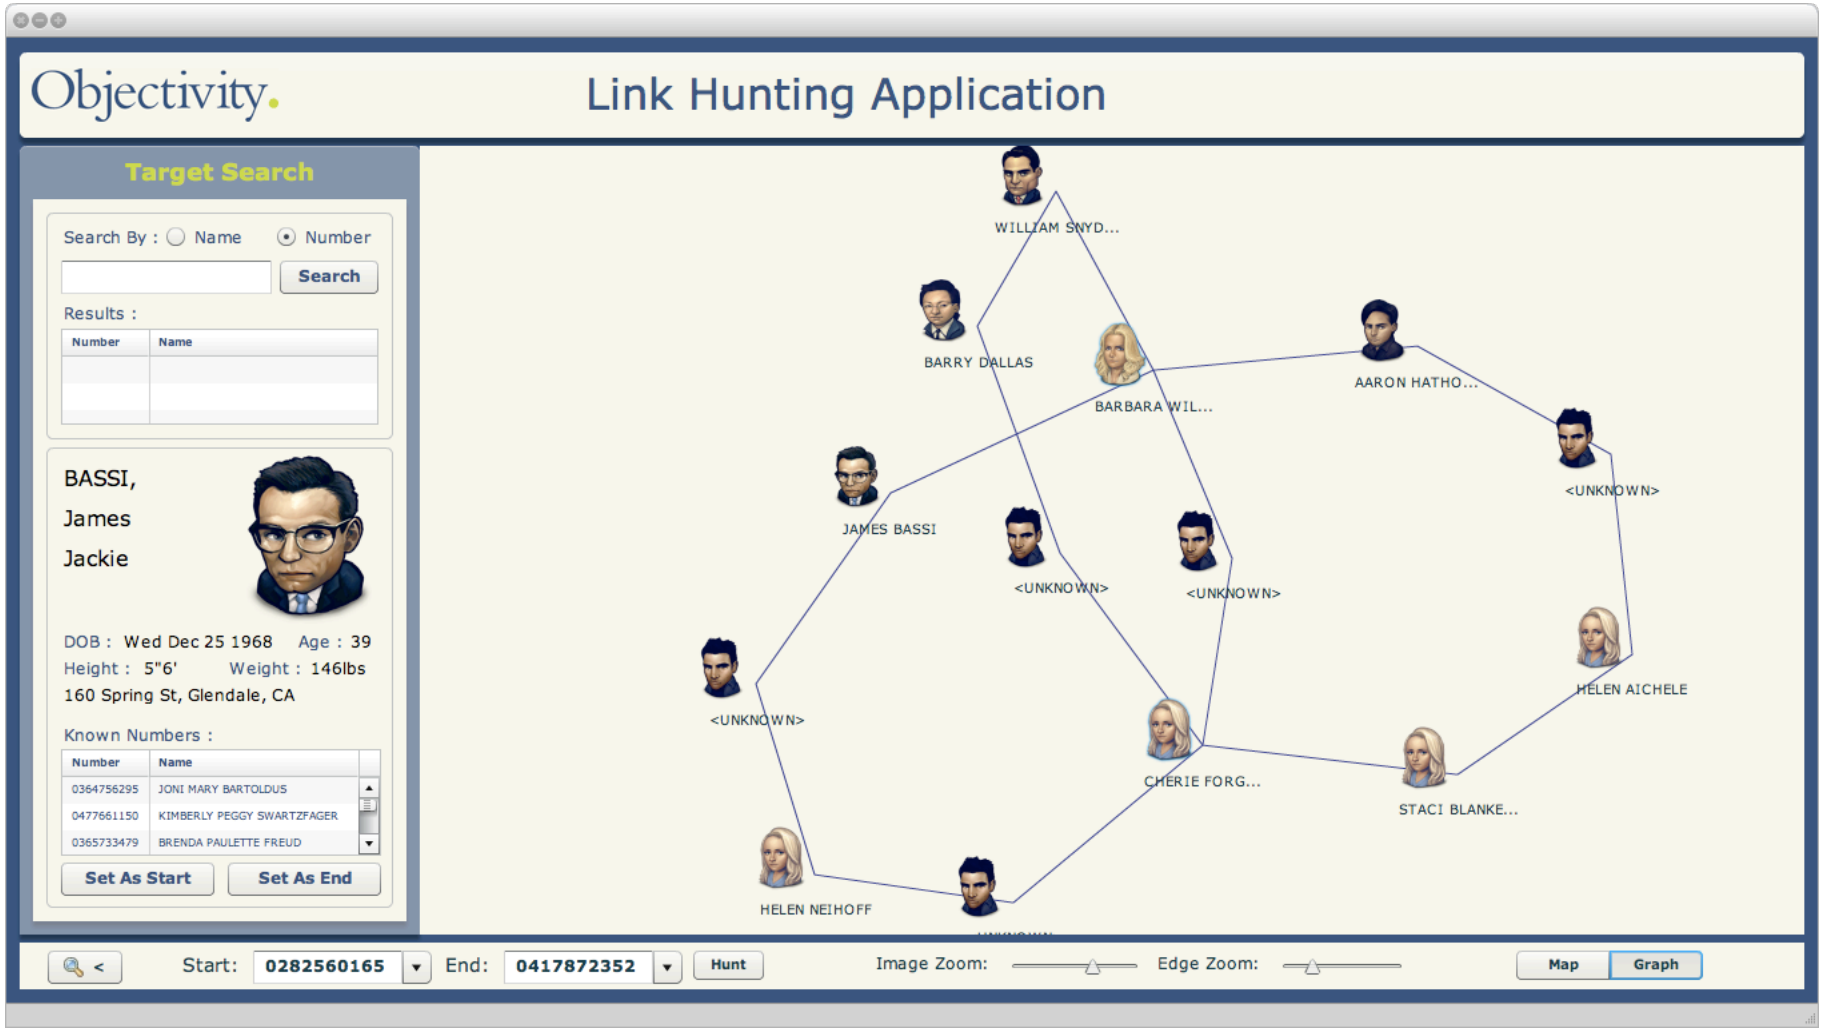

Objectivity

Graph Databases

- Nodes are represented as Vertices
- Relationships are represented as Edges
- Edges may be weighted
- Both are regular objects
  - Properties
  - Methods
  - Inheritance

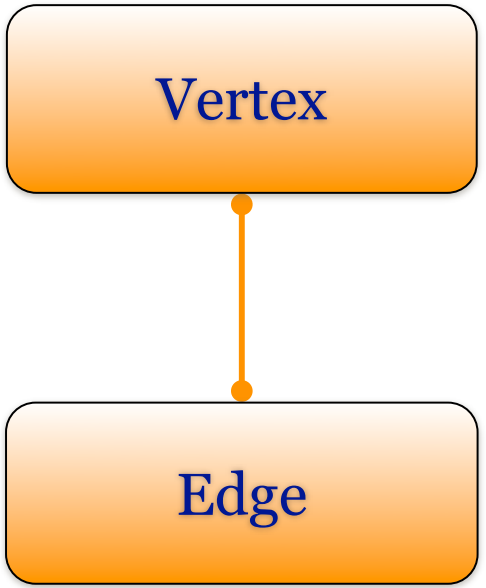

InfiniteGraph...

- Dedicated Graph API
- Easy to use and deploy
- Java now, C# soon...
- Built on Objectivity/DB
- Fully interoperable

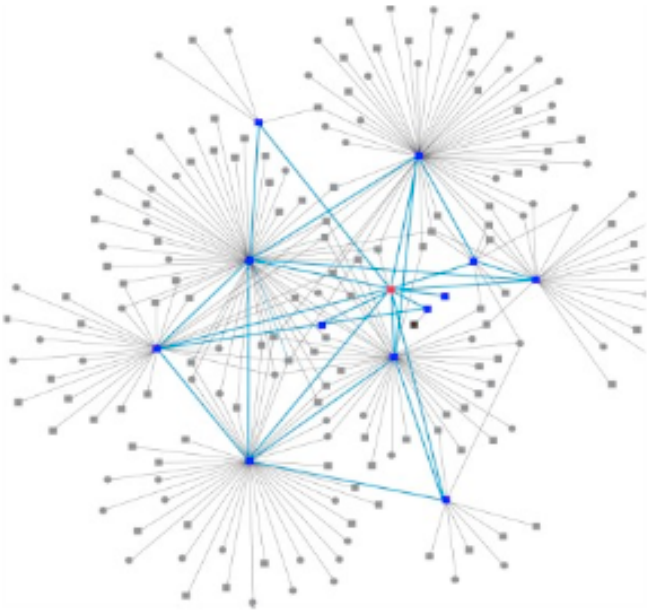

...InfiniteGraph...

- Create and update graph structures
- Traversal with constraints:
  - Return only designated Edge types
  - Return if not an excluded Edge type
  - Direction matches intent
  - Properties do not match a provided predicate
  - Maximum path depth has been reached...
- Optional Event notification
  - Target vertex was found
  - Path is being abandoned

...InfiniteGraph

- Search and Query Processing
  - Supports Objectivity/DB and Lucene indexing
  - Key and range queries
  - Full text searching
  - Regular expression search of string-based keys
- Path Finding
  - Start and end vertices
  - Maximum depth
  - Vertex type inclusion and exclusion lists
  - Edge type list

InfiniteGraph Licensing

- 60-day free trial
- Free **GoGrid** cloud development and deployment environment for qualified startups and non-profits
  - Licenses must be procured after a pre-agreed annual revenue level is reached
- Open Source framework licenses
- Standard commercial licenses

Case Studies

Objectivity/DB & Hadoop...

HDFS Weaknesses

- HDFS cannot be directly mounted by an existing operating system
  - Getting data into and out of the HDFS file system, an action that often needs to be performed before and after executing a job, can be inconvenient.
  - A filesystem in userspace has been developed to address this problem, at least for Linux and some other Unix systems.
  - **It moves 64+ MB blocks and doesn't support POSIX file operations**
- Replicating data three times is costly.
  - However, there is a version that uses a parity block, decreasing the physical storage requirements from 3x to around 2.2x.
- **The Name Node is a single point of failure**
  - **If the name node goes down, the filesystem is offline.** When it comes back up the name node must replay all outstanding operations. This replay process can take over half an hour for a big cluster.

...Objectivity/DB & Hadoop...

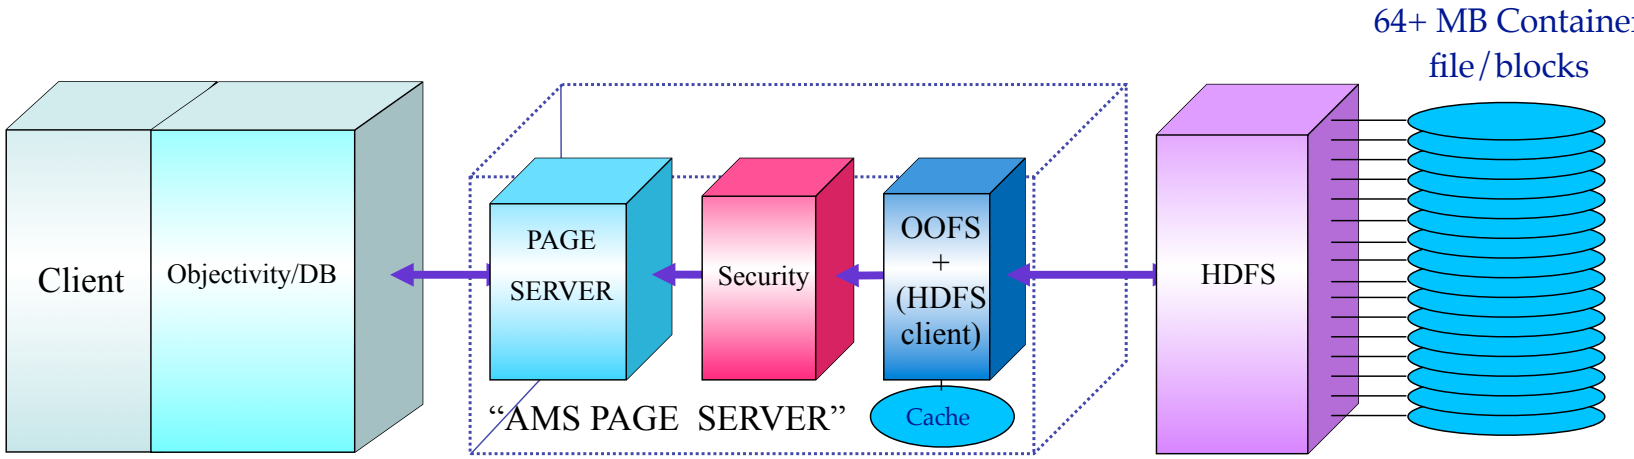

- **Problem:** HDFS works best when transferring 64+ MB blocks of data. Objectivity/DB works best with 16-64KB blocks.
- **Solution:** Implement a memory cache for the HDFS blocks.

...Objectivity/DB & Hadoop

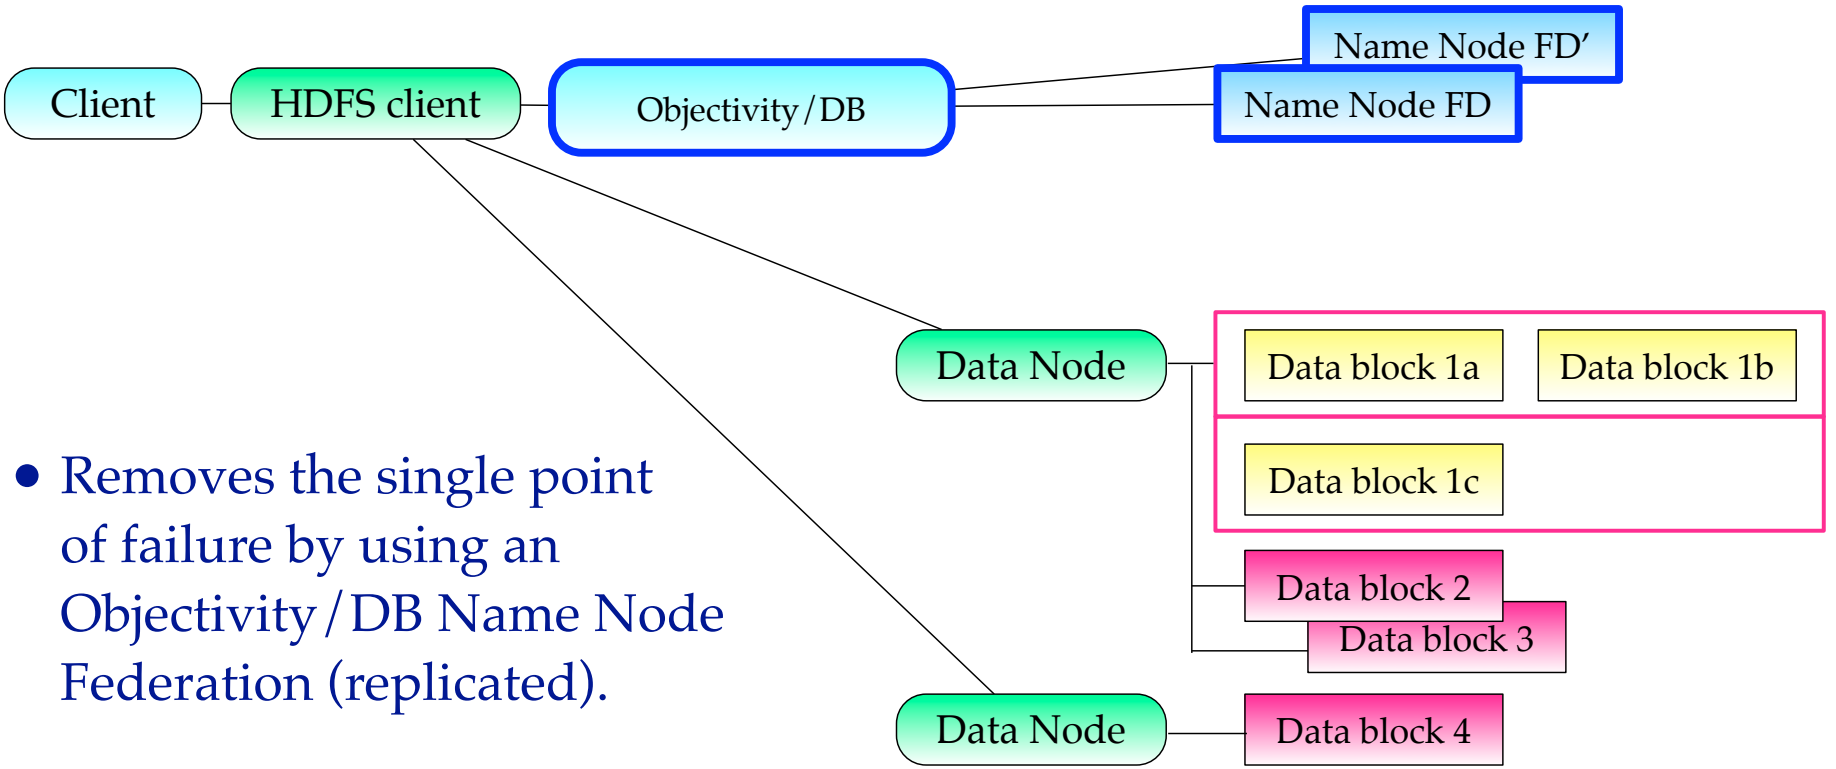

- Removes the single point of failure by using an Objectivity/DB Name Node Federation (replicated).
- Could provide other services too, such as content tagging and connectivity.

InfiniteGraph and Cassandra

- Apache Cassandra provides a structured key-value store with eventual consistency.
- It is a resilient, distributed DBMS.
- The prototype uses social network data.
- It extracts data from Cassandra, then finds the shortest paths between people.

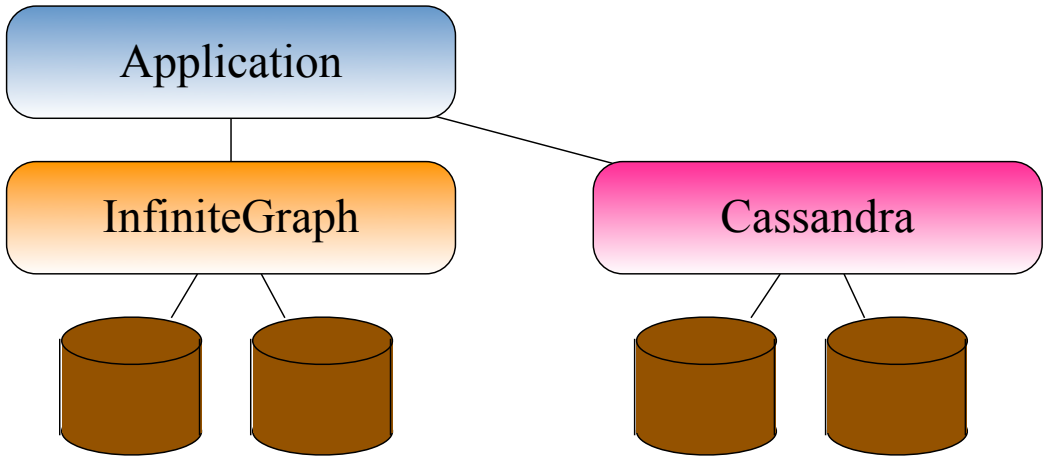

Summary

NoSQL?

- All of these features could have been obtained from "Commercial Off The Shelf" ODBMSs:

  - Unique object/document IDs
  - Sharding
  - Shared-Nothing
  - Fully distributed
  - No lock and novel transaction modes
  - Iterators and fast, predictable traversals
  - Fast scans
  - Optimization for random access
  - In Memory Database configurability
  - Flexible object clustering
  - Effectivity (data in a relationship)
  - Geospatial and multi-dimensional indexing
  - Hash table (key-value) lookups
  - Hyperspace = = single logical view of a federation
  - Multi-way replication
  - High Availability
  - Text searching

Summary

- We need to re-examine the reasons that the NoSQL movement didn't just use ODBMSs.
- The NoSQL movement helps strengthen our argument that RDBMSs aren't always the best choice.
- Graph DBMSs can supplement RDBMS, NoSQL and ODBMS technologies.
- The best Graph DBMSs are built on ODBMSs.
- ODBMSs are here to stay.

Questions?

- objectivity.com - White papers, downloads etc.
- EMAIL: info @ objectivity.com

- infinitegraph.com - WPs, downloads etc.
- EMAIL: info @ infinitegraph.com

- Presenter: leon @ objectivity.com

If they give you paper with lines on, write sideways!