Cache Augmented SQL (CASQL) Systems

Shahram Ghandeharizadeh, Jason Yap
USC Database Lab

Interactive Social Actions

- **Read intensive workloads. Queries that compute the same result repeatedly.**
  - **Same result because the database changes infrequently, e.g., your Facebook account:**
    - **How often do you visit your Facebook profile page?**
      - **Once a second, every minute, every hour, once a week?**
    - **How often do you add and drop friends/change your profile?**

CASQL Architecture

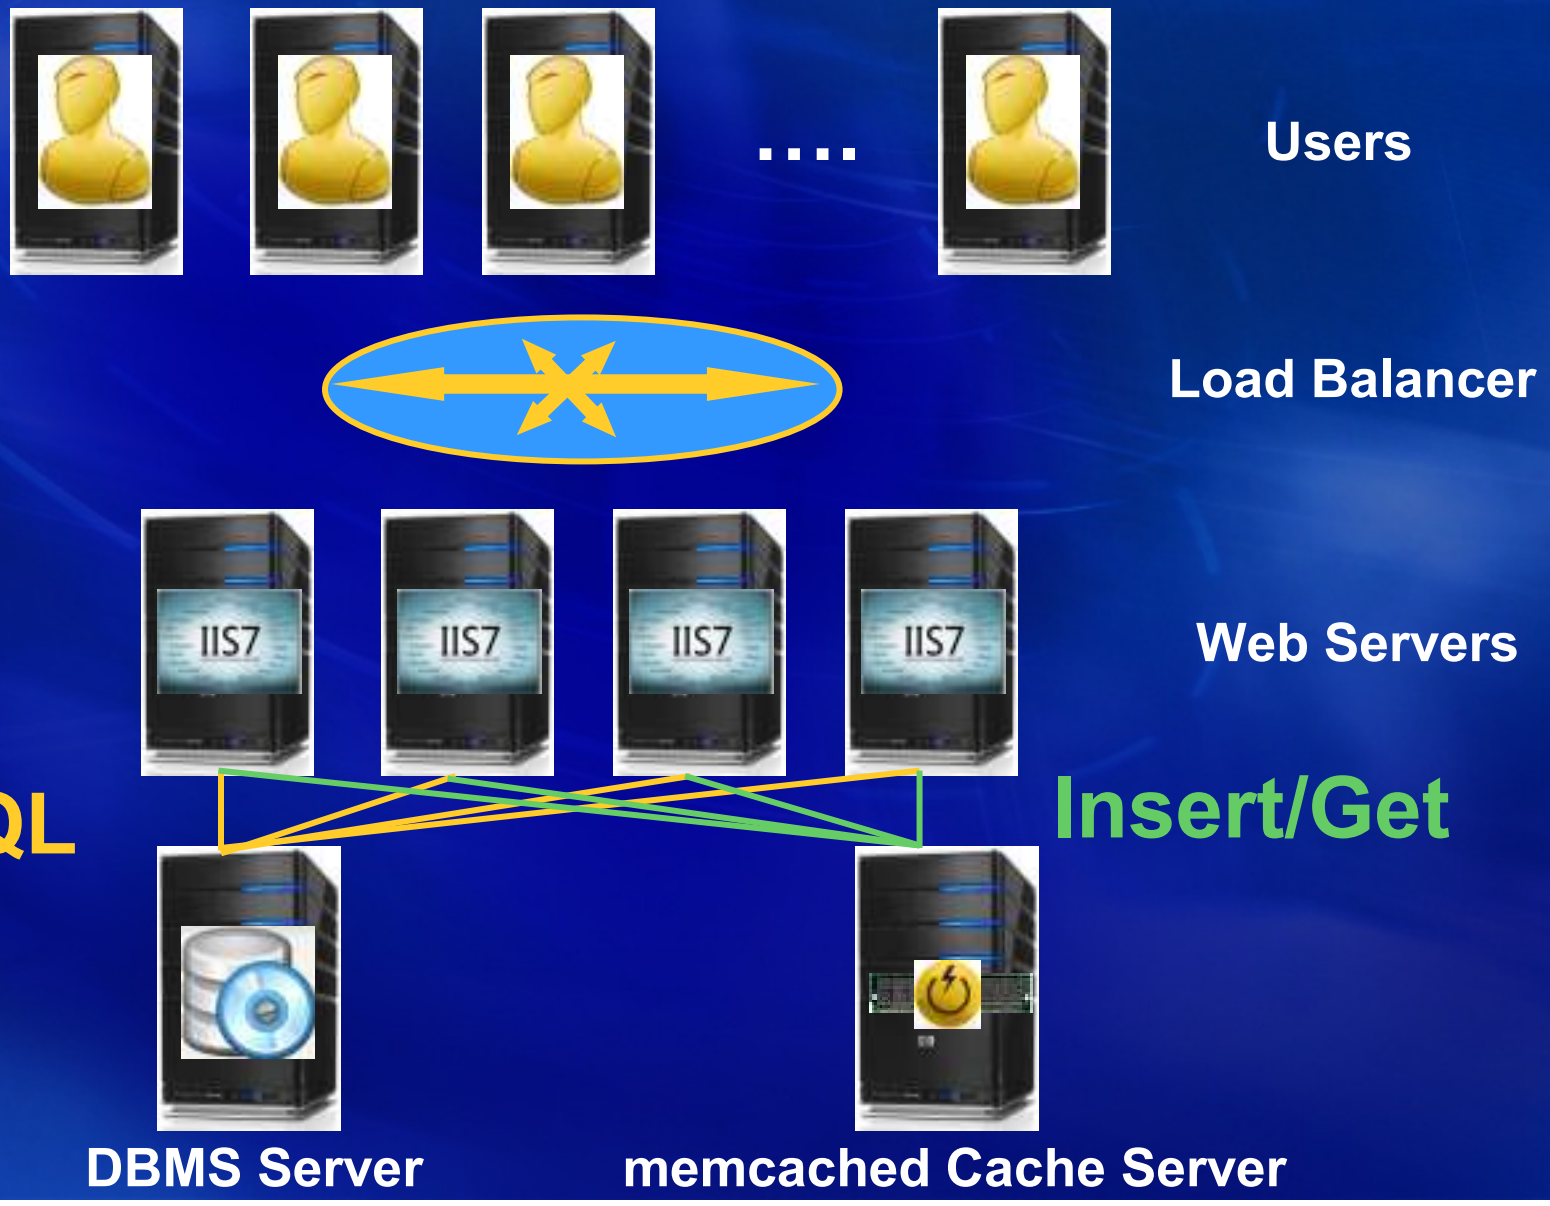

Enhanced Performance

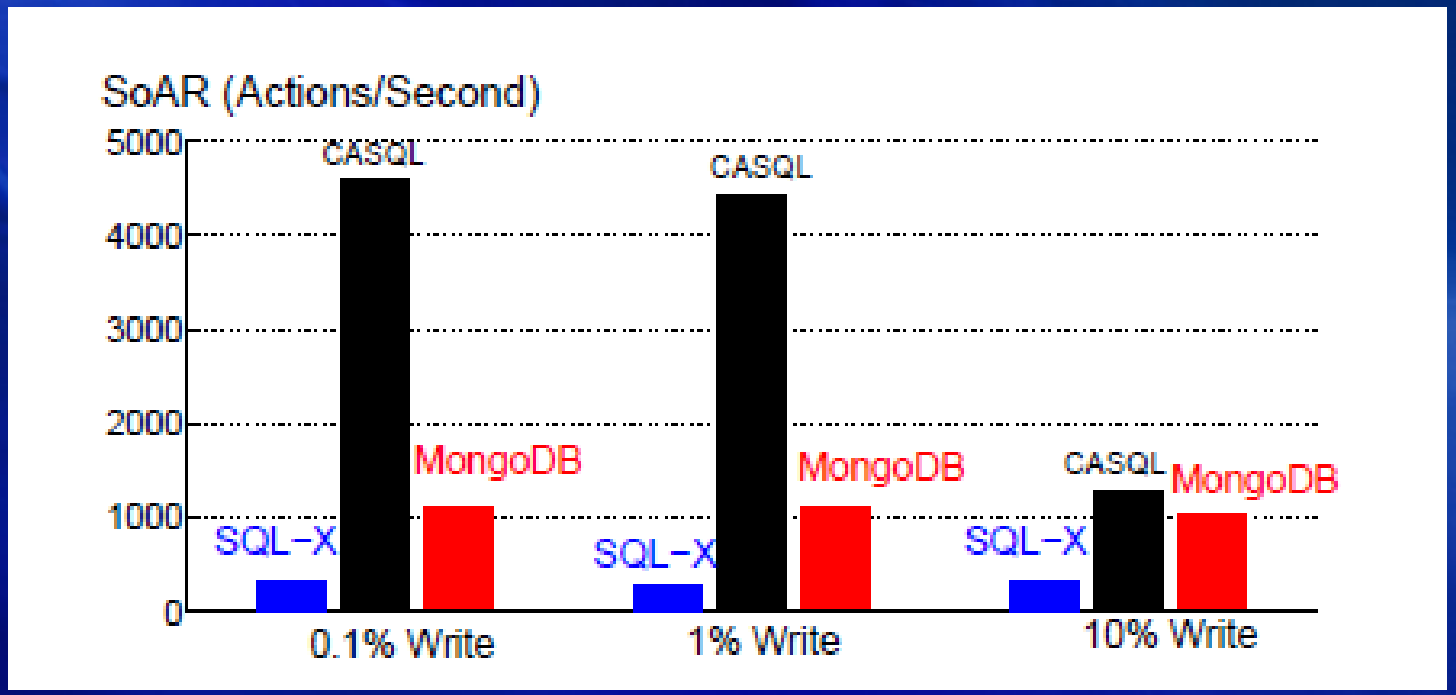

SLA: 95% of actions to observe a response time below 100 msecs

CASQL Today

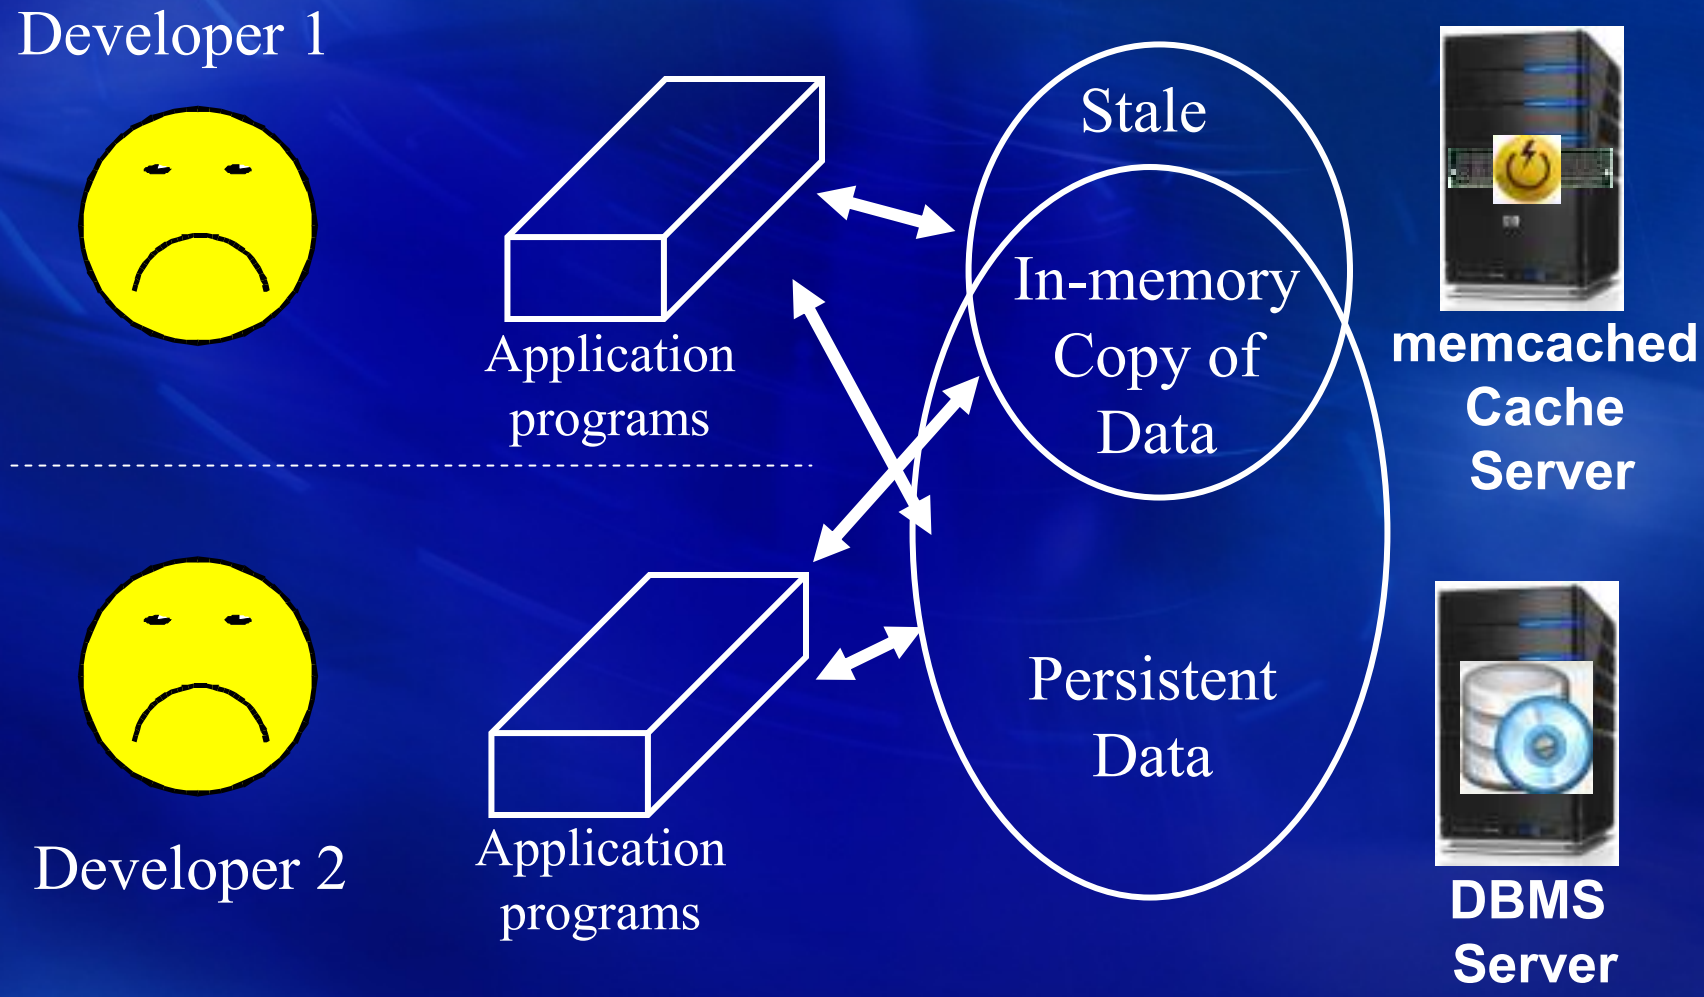
Developer 1
Application programs
Stale
In-memory Copy of Data
memcached Cache Server
Persistent Data
DBMS Server
Developer 2
Application programs

Before DBMS: 1960/70s

Developer 1

Application programs

Developer 2

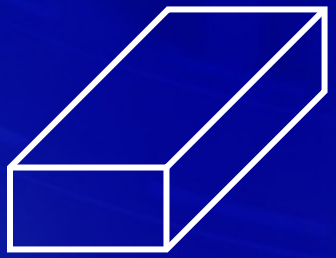

Application programs

CASQL DBMS

Developer 1

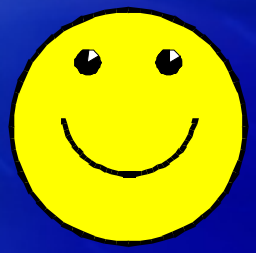

Application programs

Key Value Cache Server

CASQL DBMS

Application programs

Developer 2

DBMS Server

Physical Data Independence.

Stonebraker/Cattel's 10 Rules

Research Challenges

- **Physical data independence:**
  - **Application transparent approach to key-value consistency.**
- **Data availability in the presence of failures.**
- **Scalability.**
- **Elasticity.**

- References: S. Ghandeharizadeh and J. Yap. Cache Augmented Database Management Systems. USC Database Laboratory Tech Report 2012-07.