University of Waterloo

David R. Cheriton School of Computer Science

Cloud Data Management Course

Application Specific Data Replication for Edge Services

Presented By

Fathiyeh Faghih

Feb. 2010

University of Waterloo

Table of Content

- Architecture
- Motivation
- TPC-W Objects
- System Evaluation
- Conclusion
- Issues and Observations

University of Waterloo

Table of Content

- **Architecture**
- Motivation
- TPC-W Objects
- System Evaluation
- Conclusion
- Issues and Observations

University of Waterloo

Edge Services Architecture

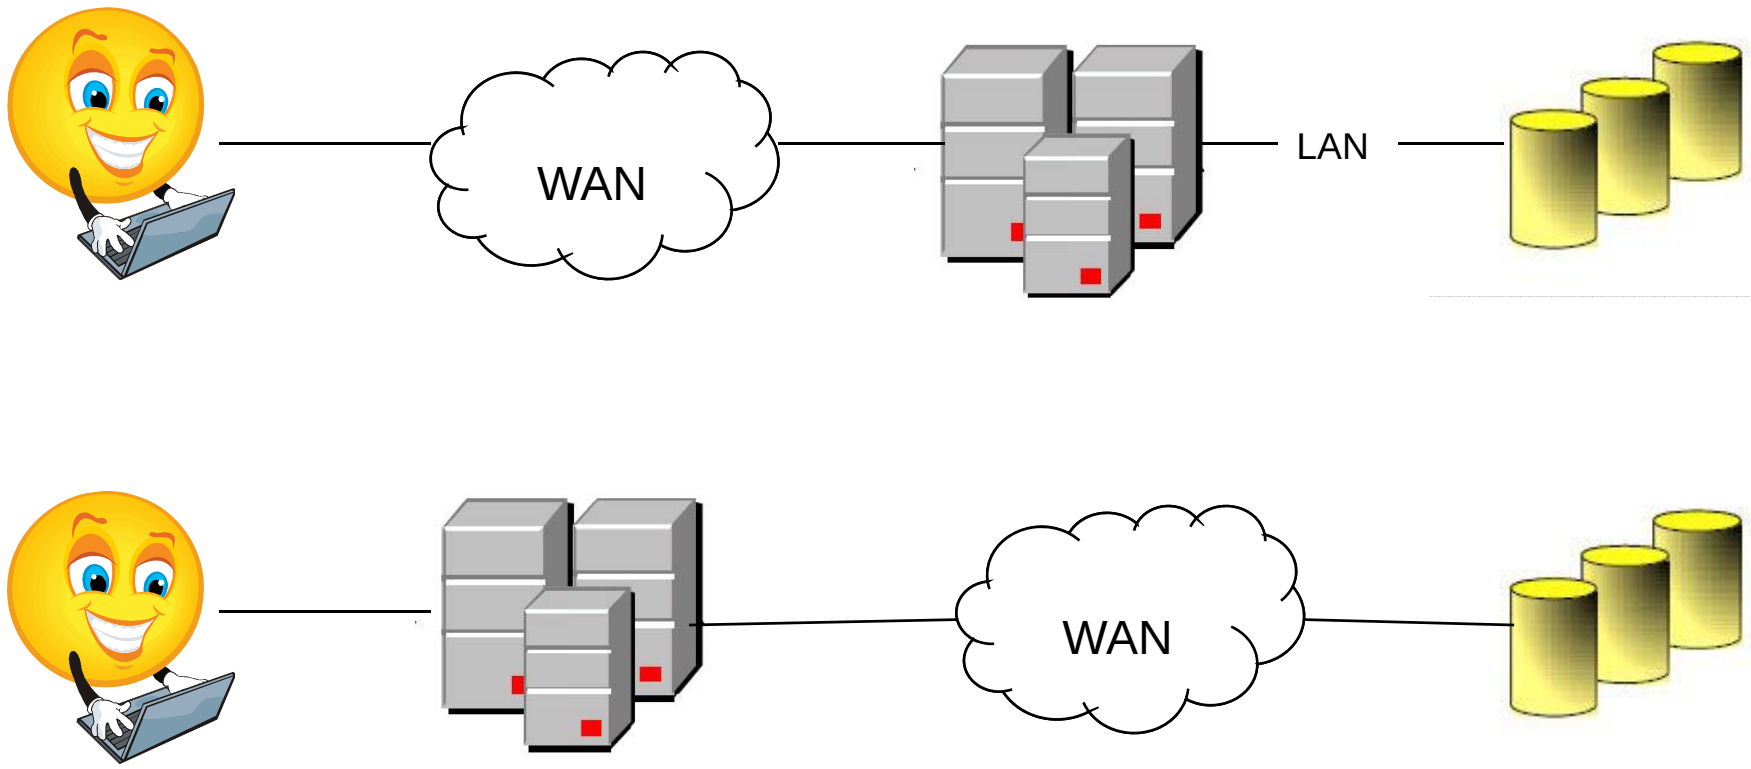

The Distributed Object Architecture

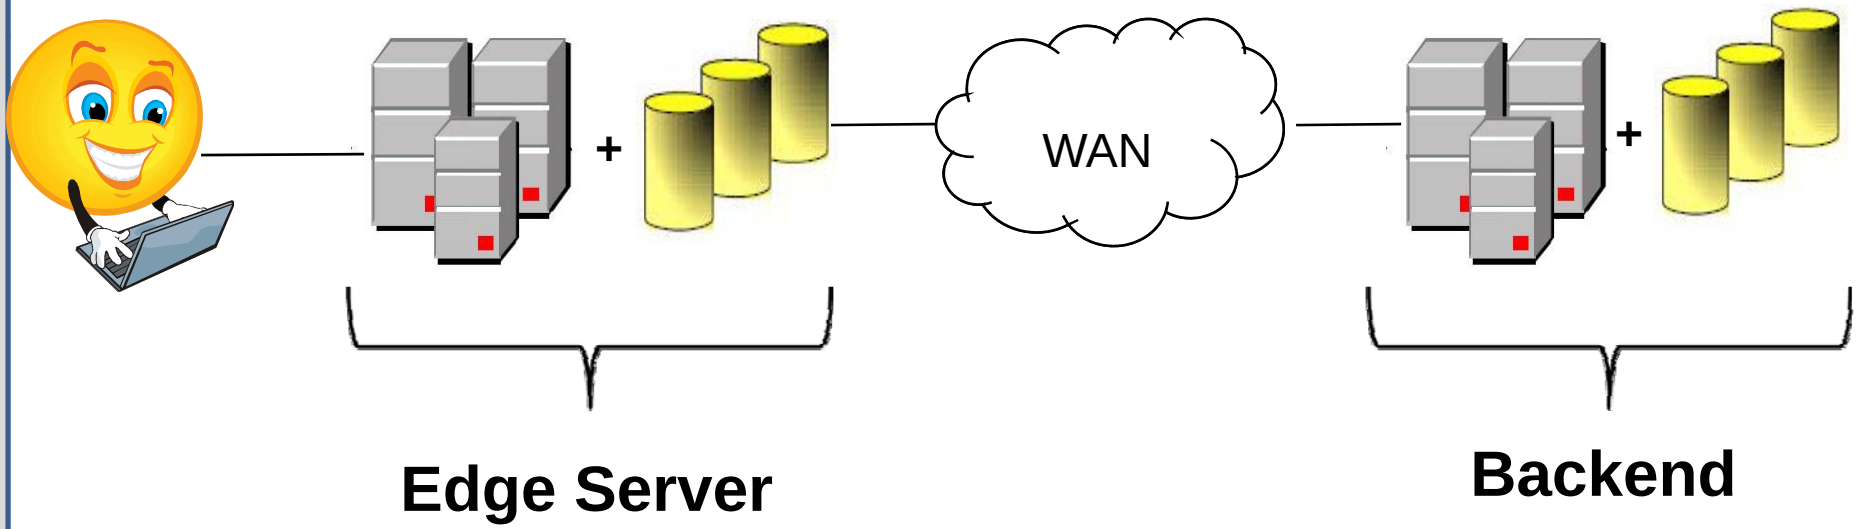

University of Waterloo

The Distributed Object Architecture

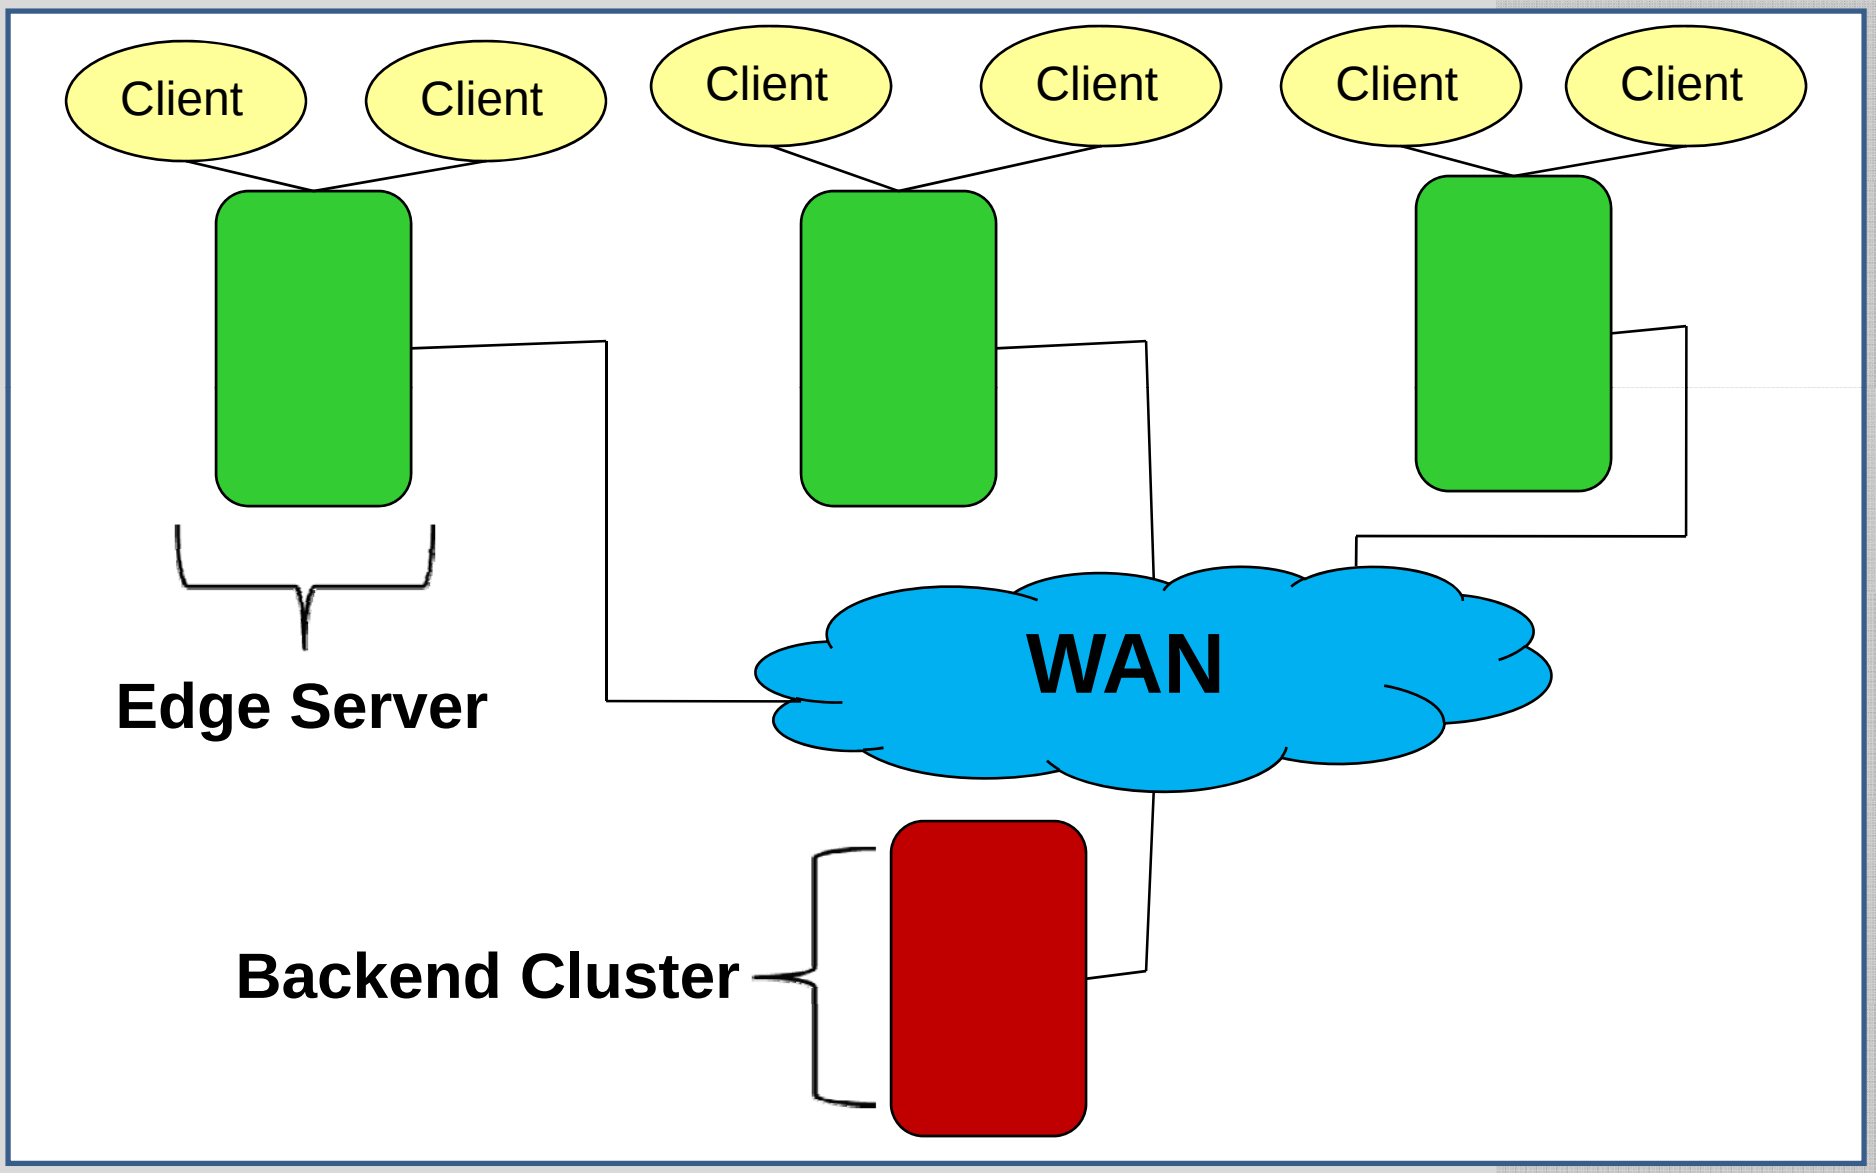

University of Waterloo

The Distributed Object Architecture

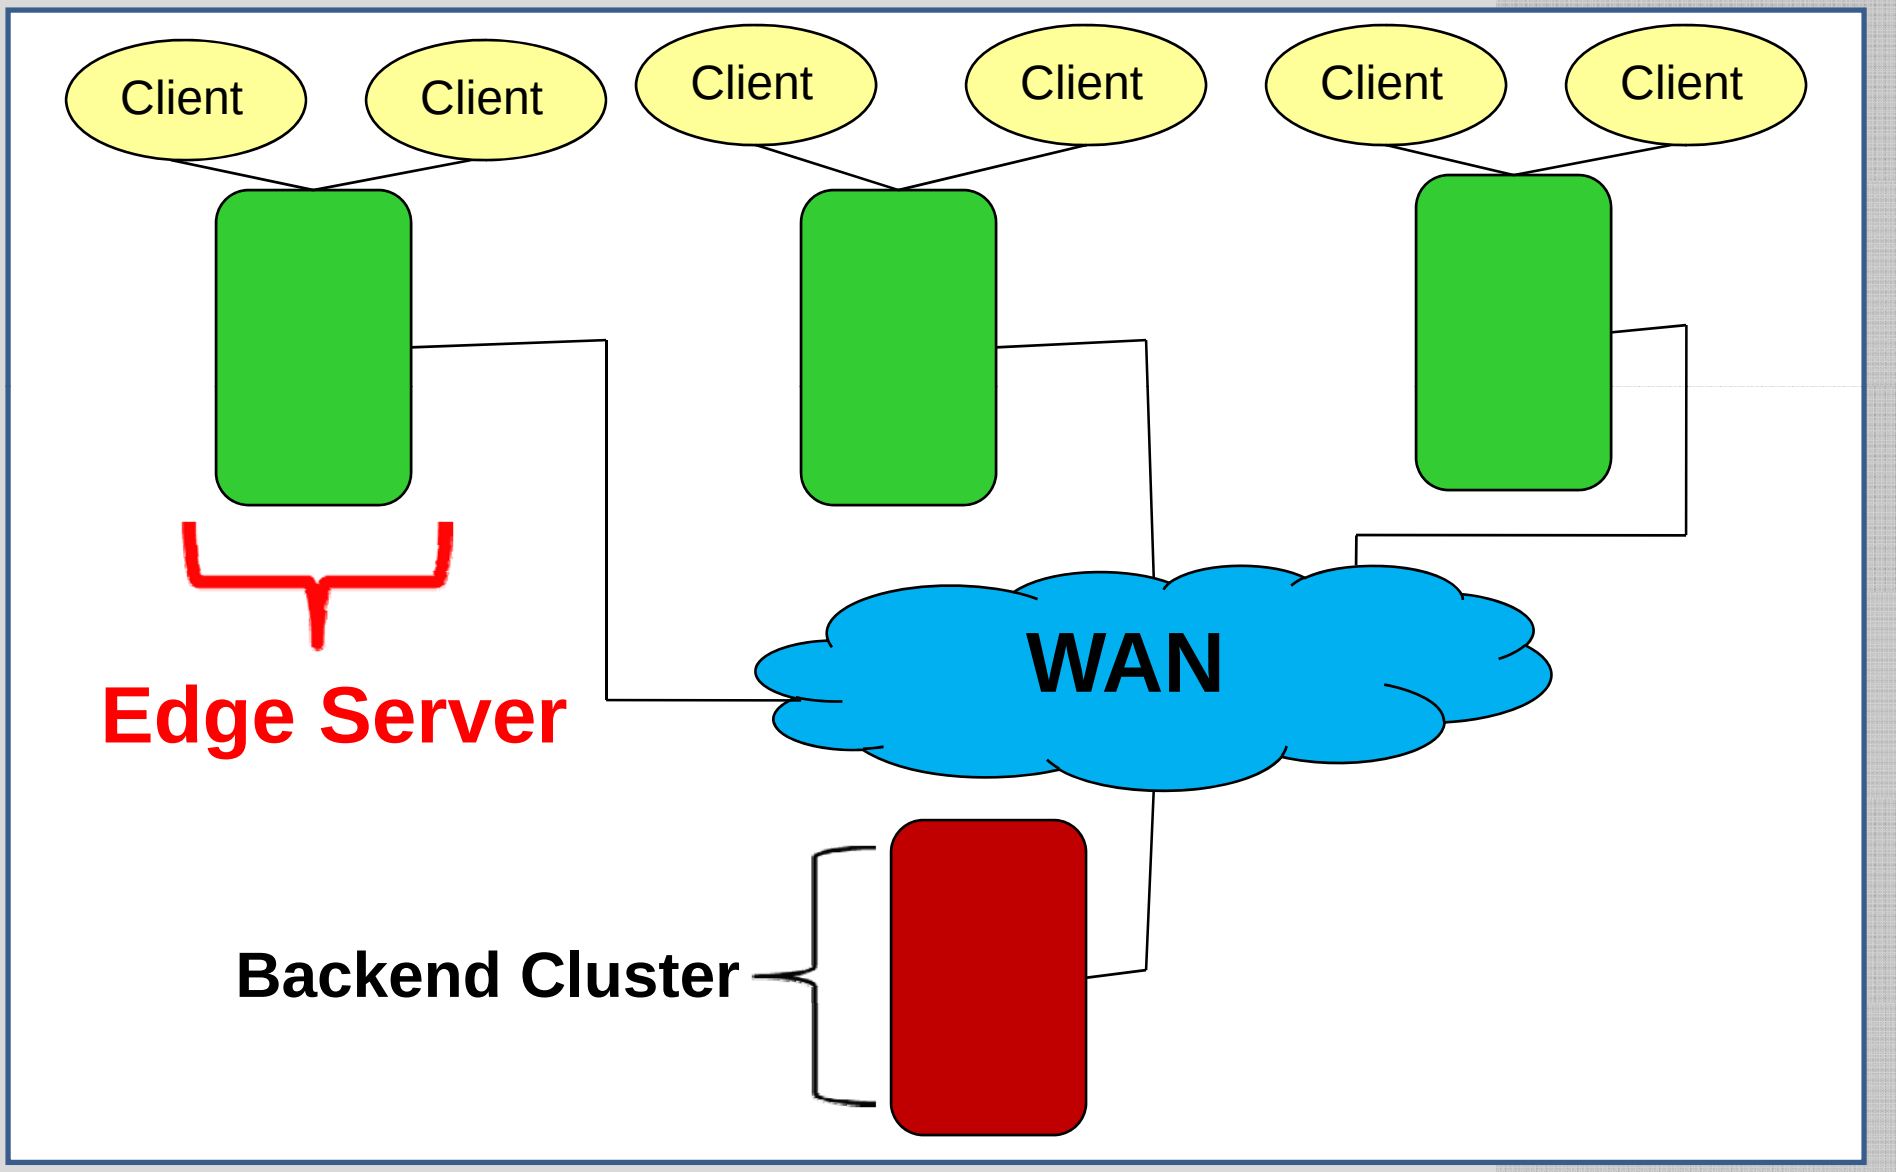

University of Waterloo

Edge Server Architecture

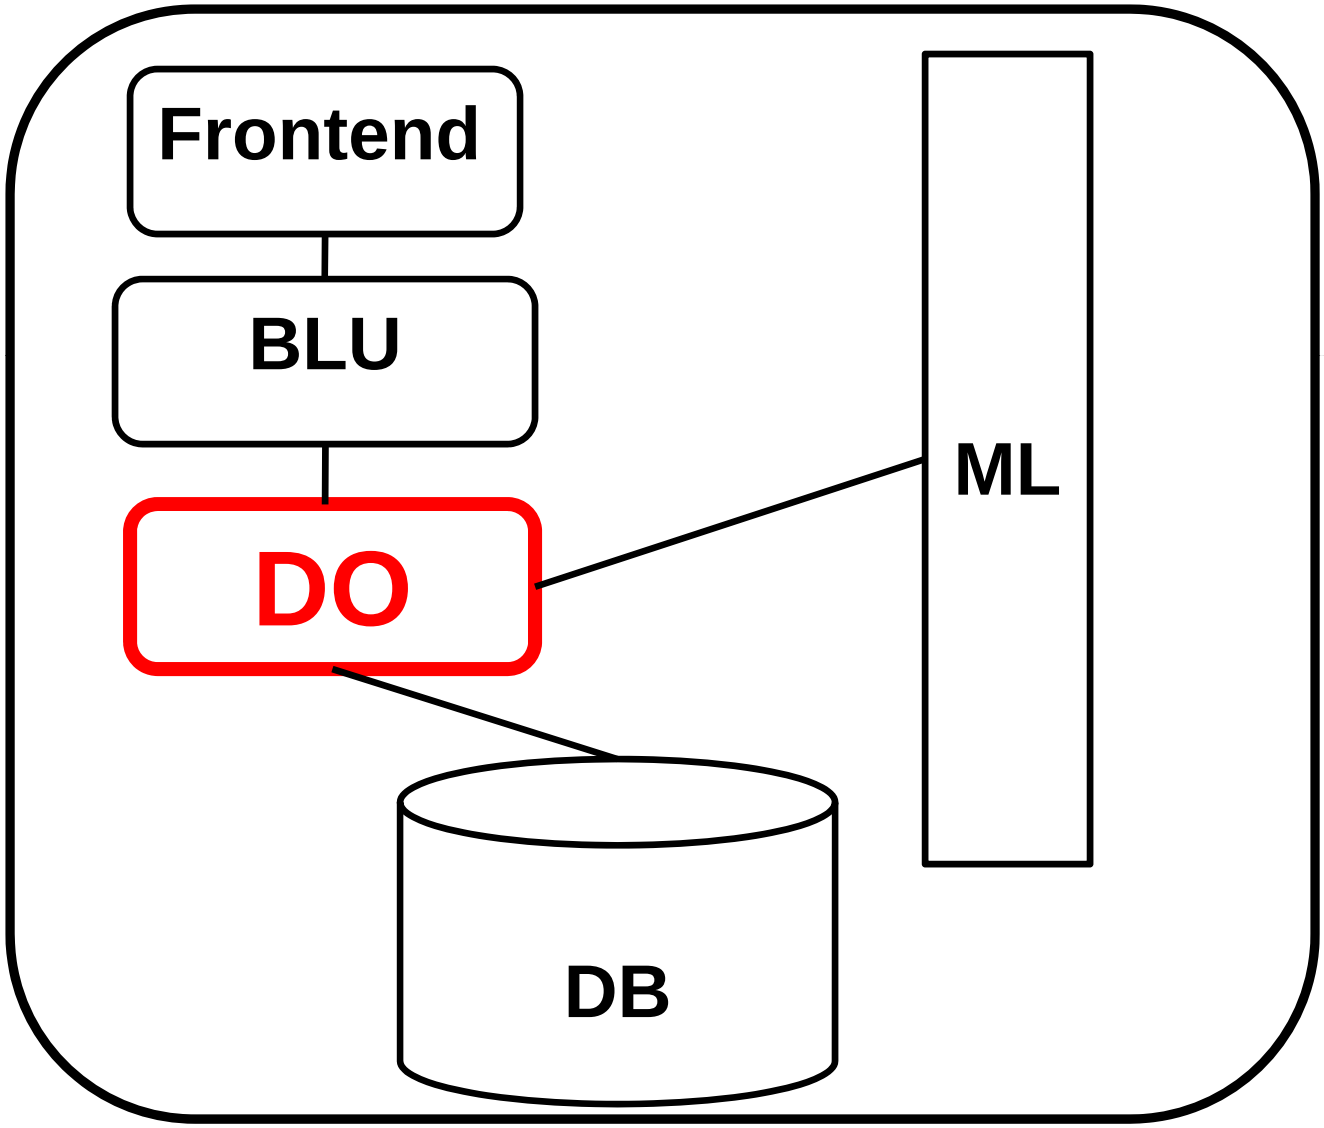

University of Waterloo

The Distributed Object Architecture

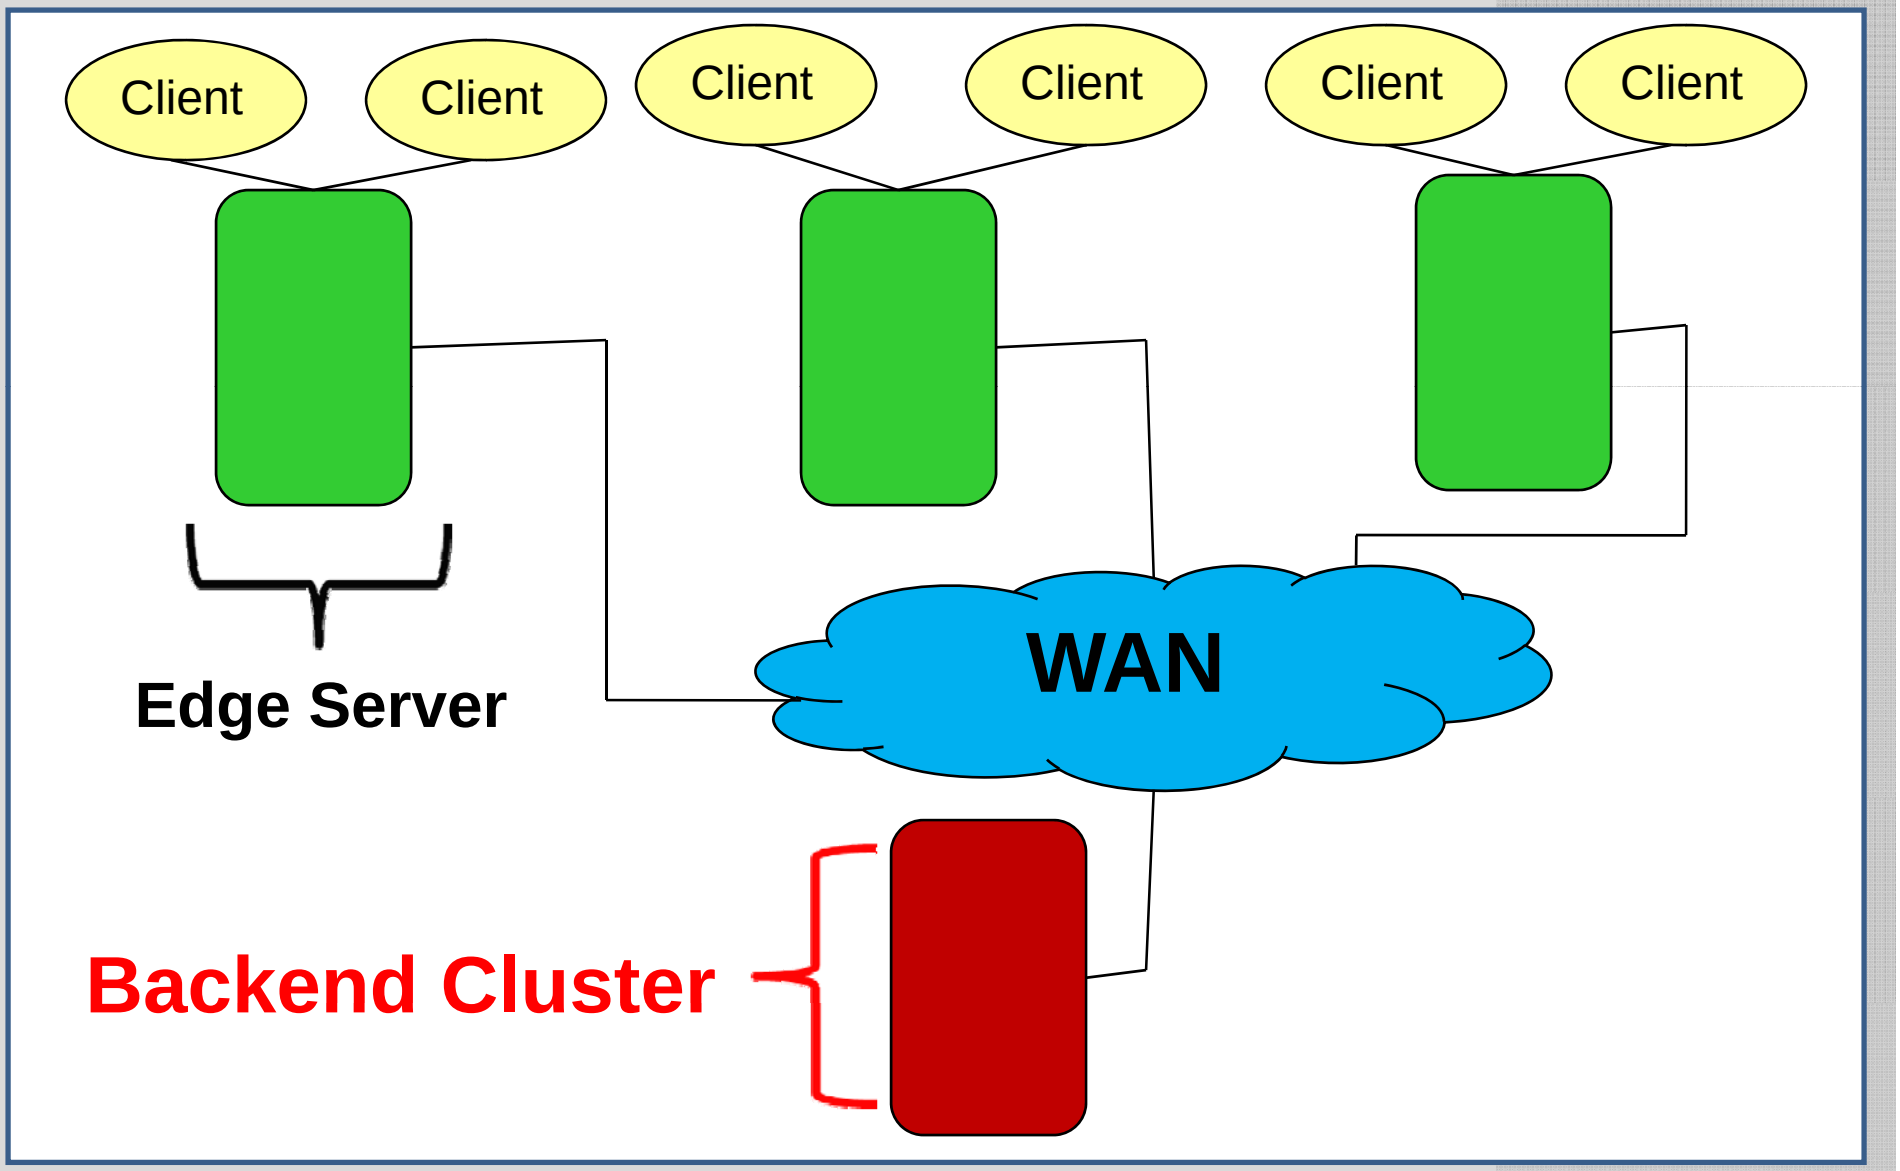

University of Waterloo

Backend Cluster Architecture

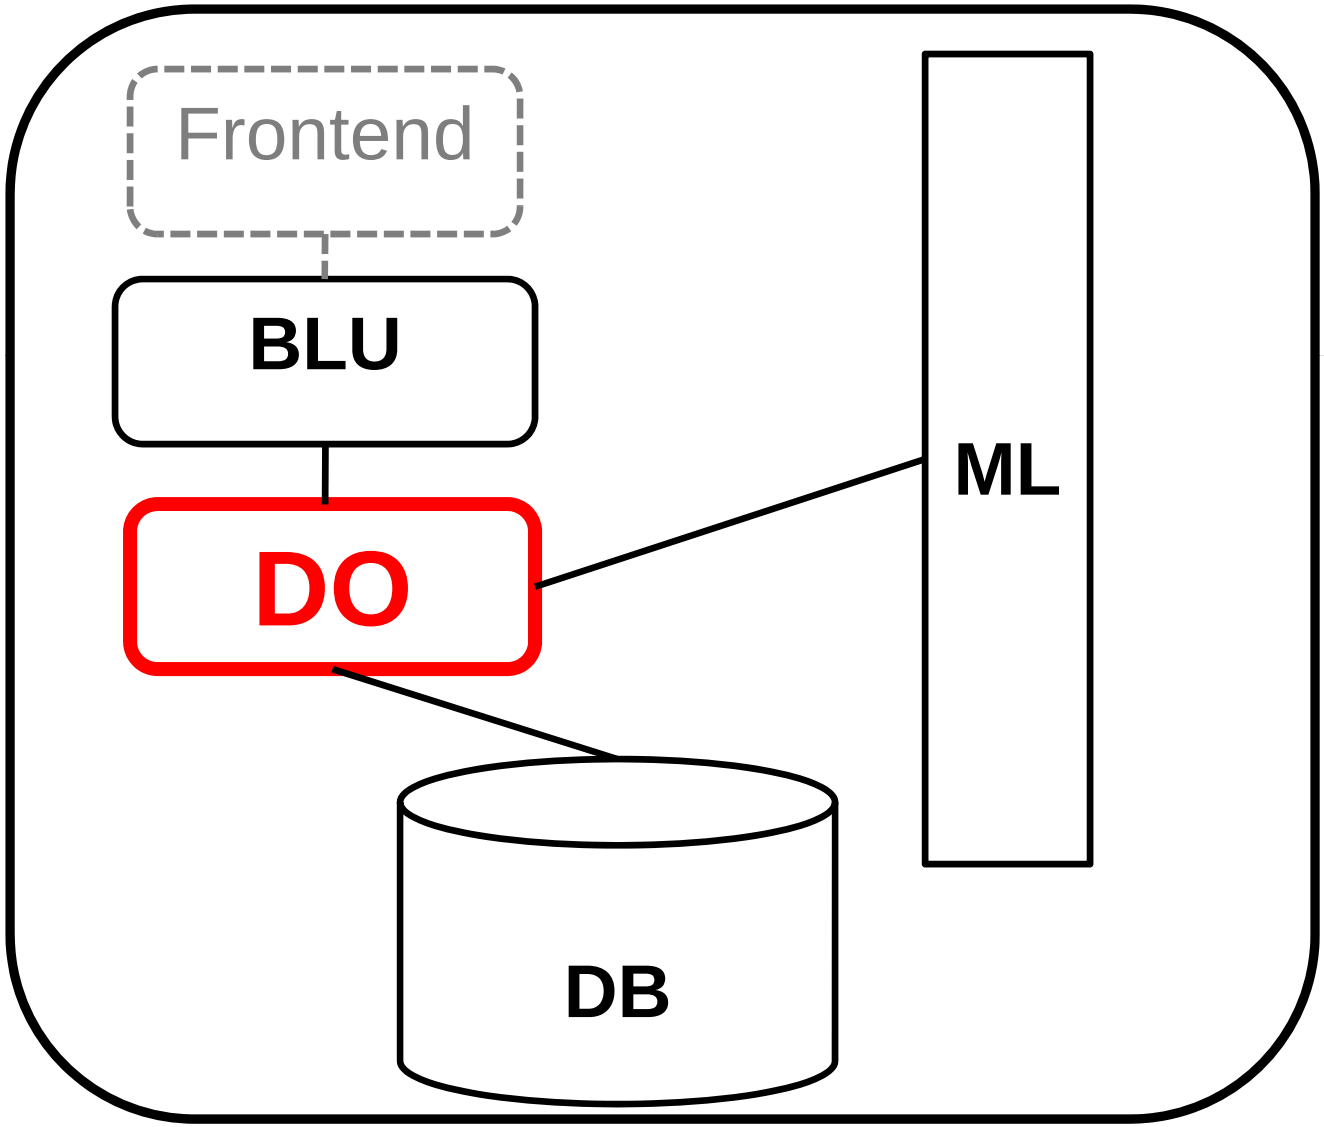

University of Waterloo

Table of Content

- Architecture
- **Motivation**
- TPC-W Objects
- System Evaluation
- Conclusion
- Issues and Observations

Motivation

- Encapsulating database access behind object specific interfaces
  - Satisfaction of client requests locally
  - Better response time
  - Consistency maintenance according to the semantics of objects
  - Encapsulation the complexity of data replication

University of Waterloo

Table of Content

- Architecture
- Motivation
- **TPC-W Objects**
- System Evaluation
- Conclusion
- Issues and Observations

TPC-W Benchmark

- Feasibility of the edge services architecture for TPC-W
  - Catalog
  - Order
  - Profile
  - Inventory
  - Best-seller-list

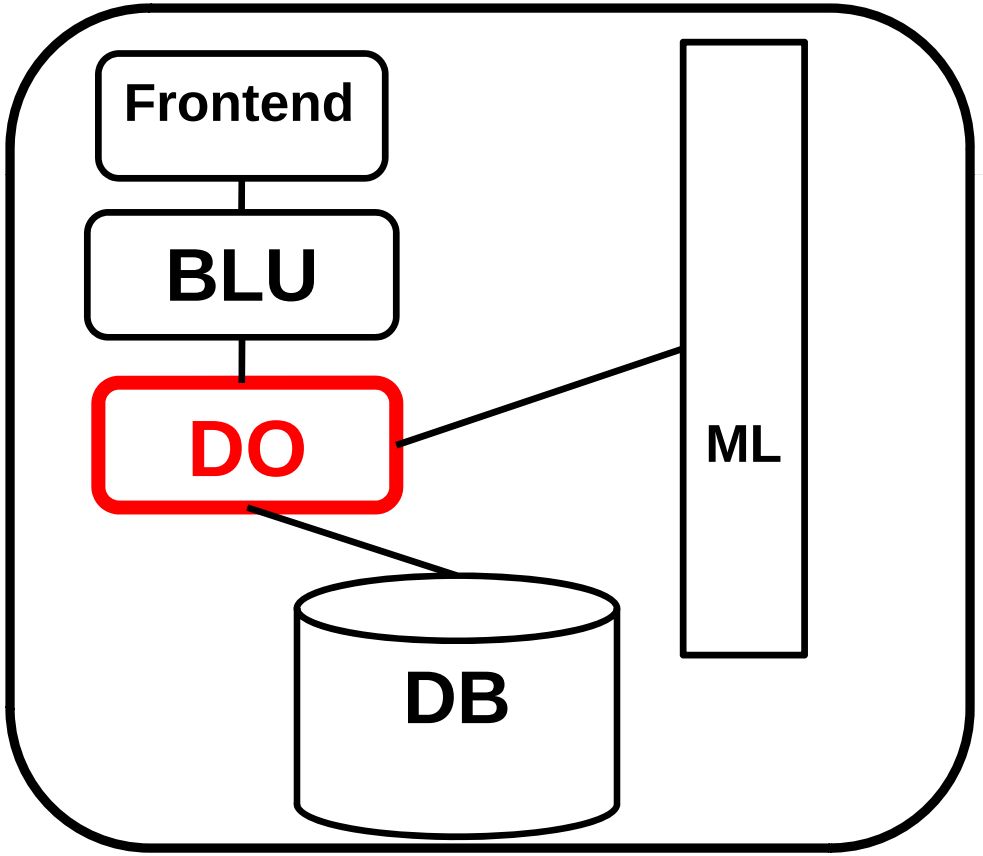

University of Waterloo

TPC-W Benchmark

- **Catalog**
- Order
- Profile
- Inventory
- Best-seller-list

The Catalog Object

- Used for managing catalog information
- Abstraction of one-to-many updates
- Push-all update strategy

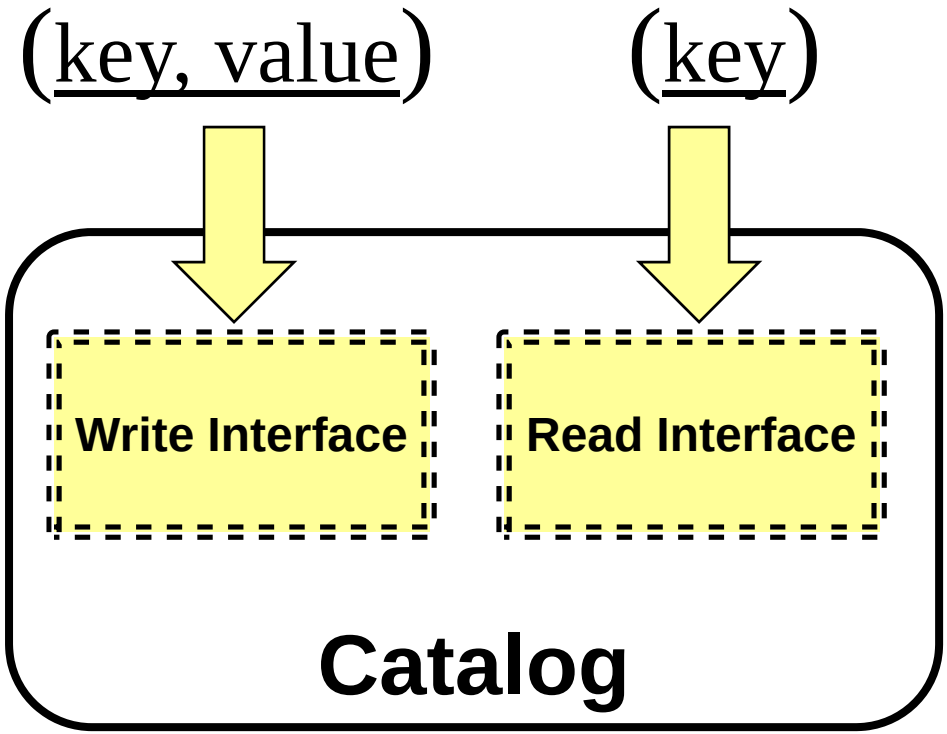

University of Waterloo

TPC-W Benchmark

- Catalog
- **Order**
- Profile
- Inventory
- Best-seller-list

The Order Object

- Used for managing the propagation of completed orders
- Abstraction of many-to-one updates
- FIFO consistency

(order, order sequence ID, edge server ID, message handler)

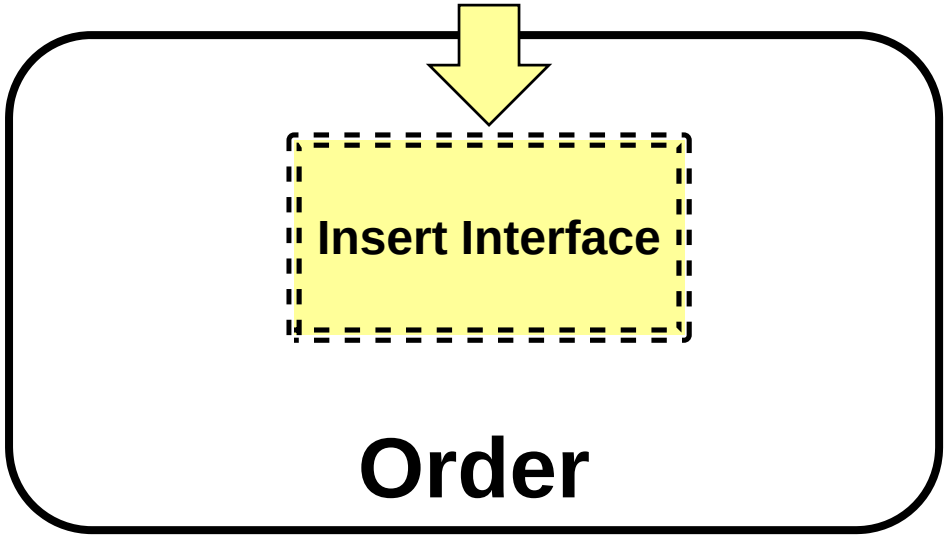

TPC-W Benchmark

- Catalog
- Order
- **Profile**
- Inventory
- Best-seller-list

The Profile Object

- Used for handling reads/writes with low concurrency

(User ID, field ID, *value*) (User ID)

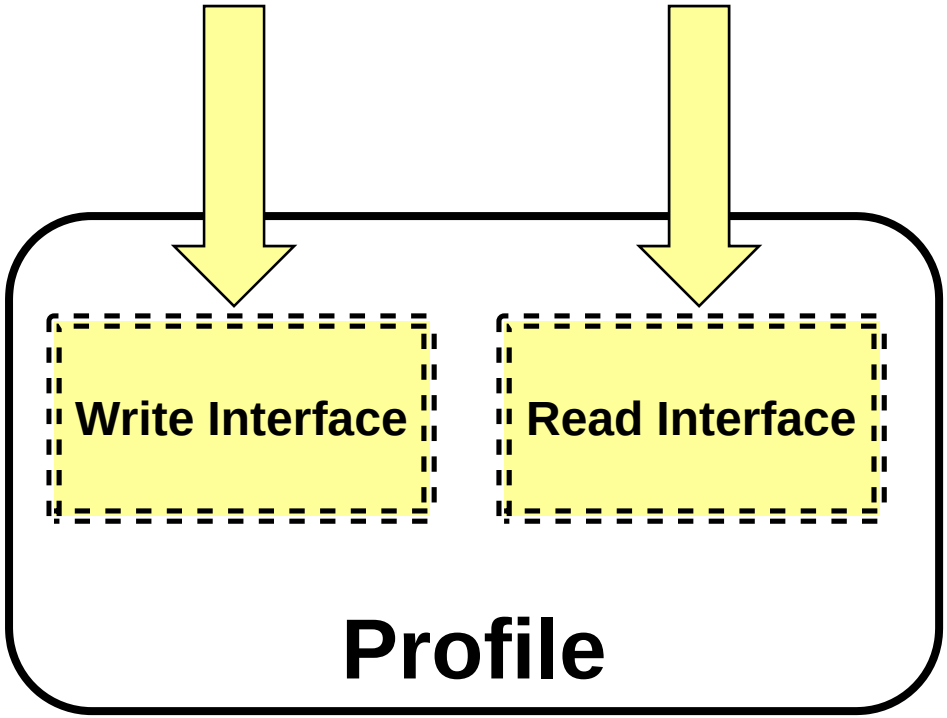

University of Waterloo

The Profile Object

- Semantics
  - Low write/read ratio
  - Sending the same user to the same edge server
- Using a write-any read-any policy
- Propagation of updates among all edge servers
- Conflict resolution using “reconciliation rules”

University of Waterloo

The Profile Object

- Availability
- Latency
  - Relaxing consistency
- Simplicity
  - Increasing storage space and update bandwidth

TPC-W Benchmark

- Catalog
- Order
- Profile
- **Inventory**
- Best-seller-list

The Inventory Object

- Used for a constraint enforcement
  - Notification of requesters, if the inventory is 0
  - Minimizing false positives

(Numeric Value, Book ID)

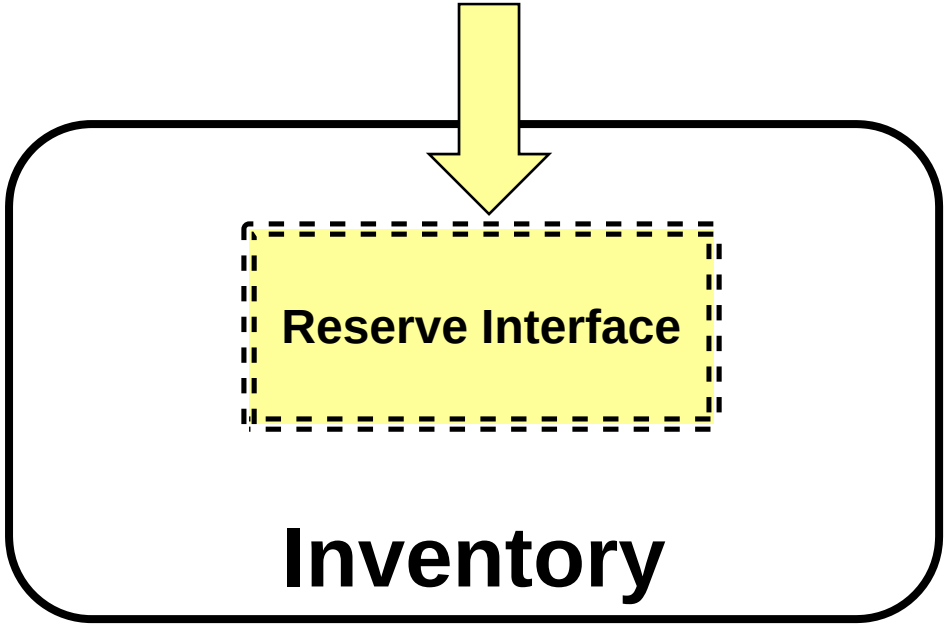

The Inventory Object

- Inventory division among edge servers
- Keeping track of the edge servers' inventories by the backend server
  - Re-distribution

The Inventory Object

- Availability
- Acceptable consistency
- System response time
- Limitation
  - False positive

University of Waterloo

Table of Content

- Architecture
- Motivation
- TPC-W Objects
- **System Evaluation**
- Conclusion
- Issues and Observations

System Evaluation

- Performance
- Availability
- Consistency

System Evaluation

- **Performance**
- Availability
- Consistency

Performance Evaluation

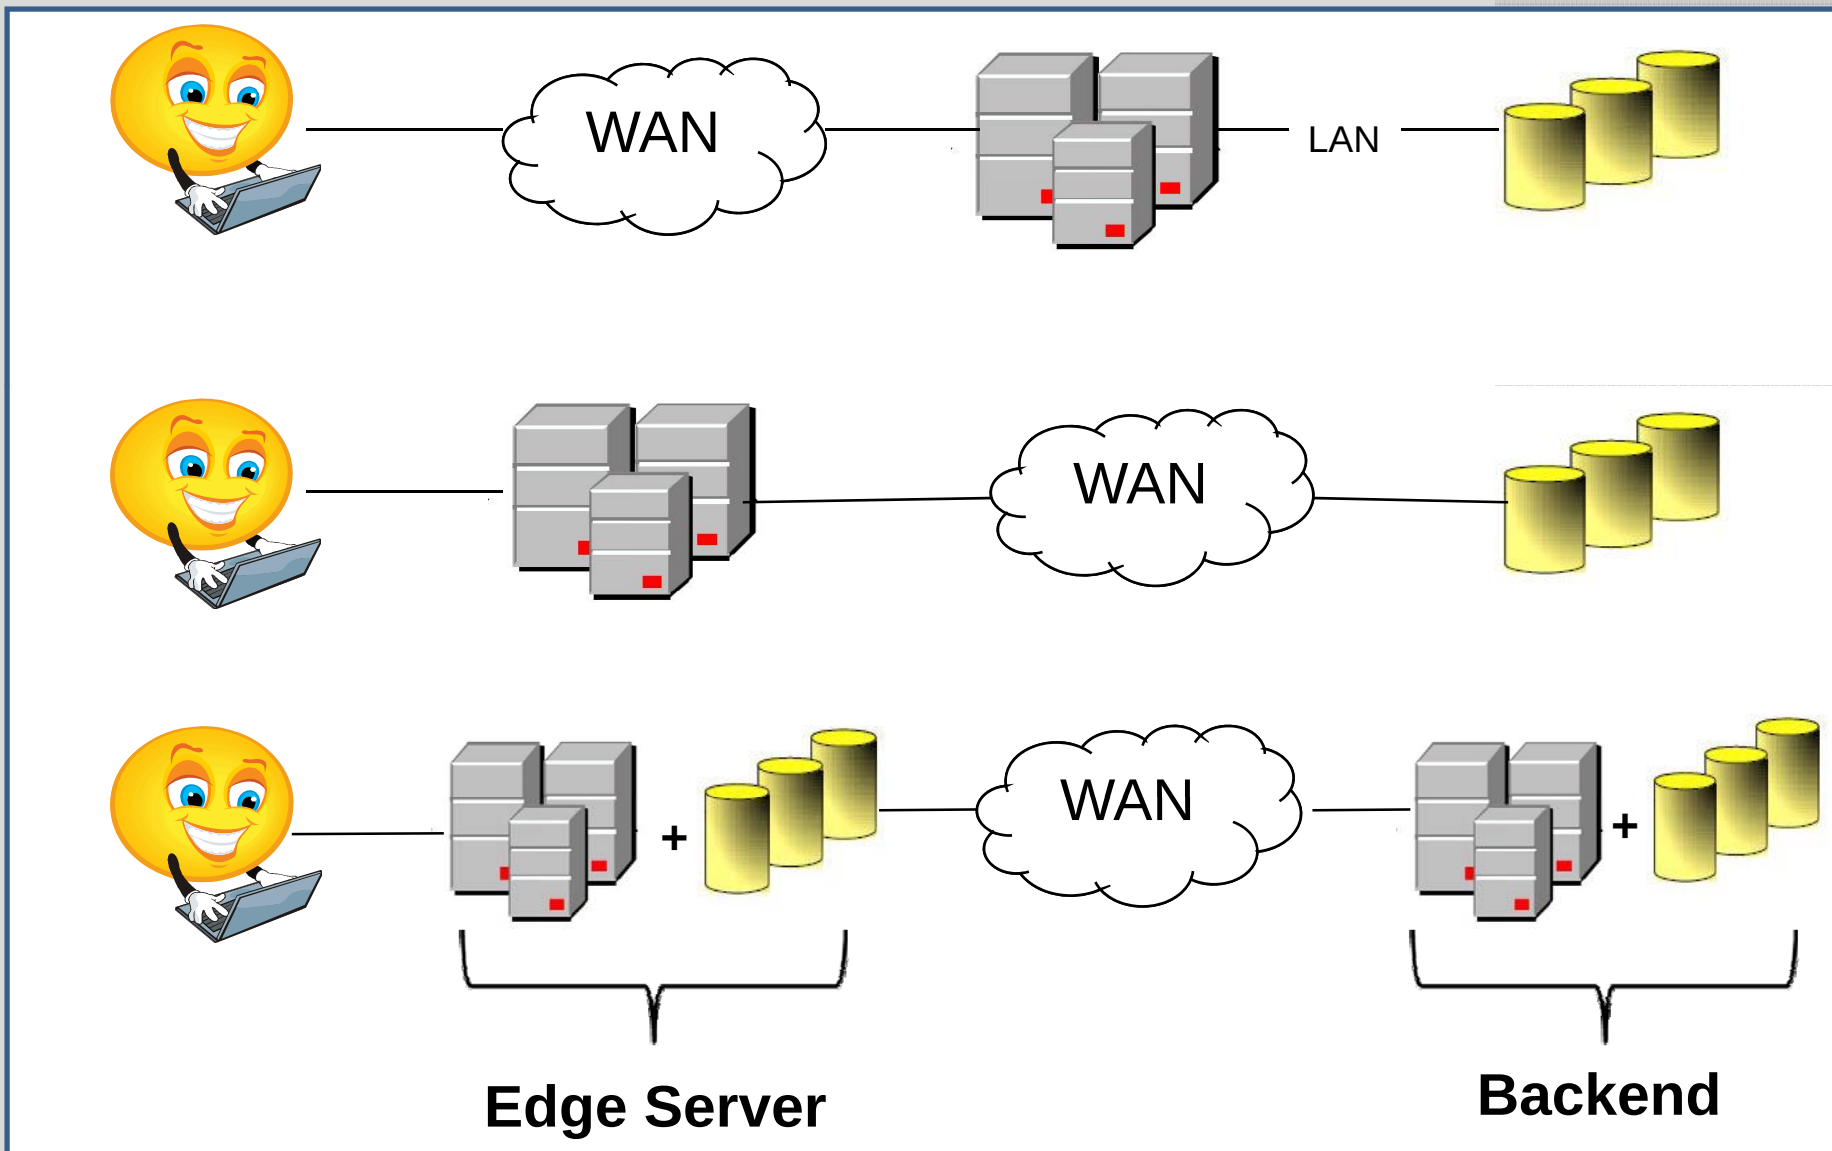

System Evaluation

- Performance
- **Availability**
- Consistency

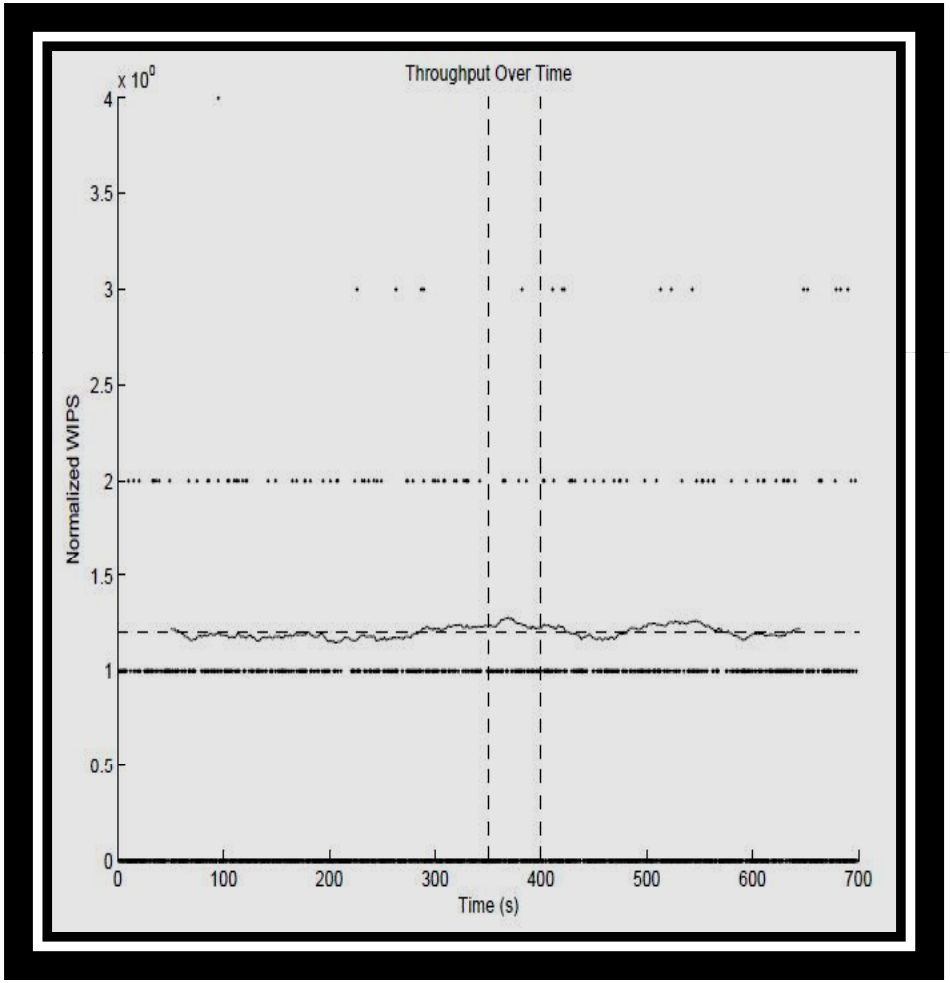

System Evaluation

- Performance
- Availability
- **Consistency**

Consistency Evaluation

- Examining the staleness of local inventory
  - Unbalanced heavy workload

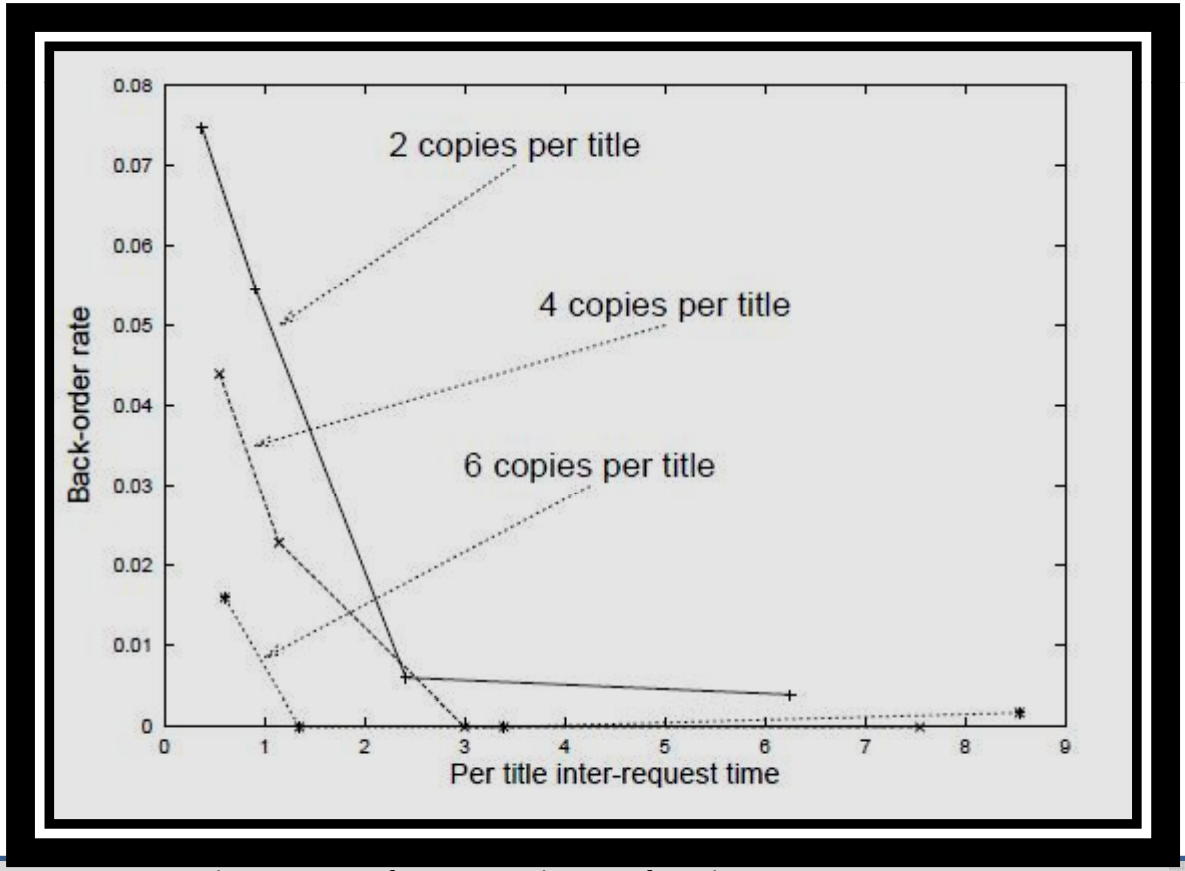

University of Waterloo

Table of Content

- Architecture
- Motivation
- TPC-W Objects
- System Evaluation
- **Conclusion**
- Issues and Observations

University of Waterloo

Conclusion

- Higher system availability and performance
  - Specific application semantics
- A quantitatively evaluation of the distributed objects model

University of Waterloo

Table of Content

- Architecture
- Motivation
- TPC-W Objects
- System Evaluation
- Conclusion
- **Issues and Observations**

University of Waterloo

Issues and Observations

- Need to developer's knowledge of the application's semantics
- Separation of "semantics object" from "replication object" as in the Globe system
- Consistency issues

University of Waterloo

University of Waterloo

Discussion

The Catalog Object

- Push-all update strategy
- FIFO consistency
  - Time delays between occurrence of an update and its being seen.
- Concern
  - Bandwidth 

    No problem! Small writes/reads ratio

TPC-W Benchmark

- Catalog
- Order
- Profile
- Inventory
- **Best-seller-list**

The Best-Seller-List Object

- Used for maintaining lists of best seller books
- Returning slightly stale information is better than stopping serving requests.

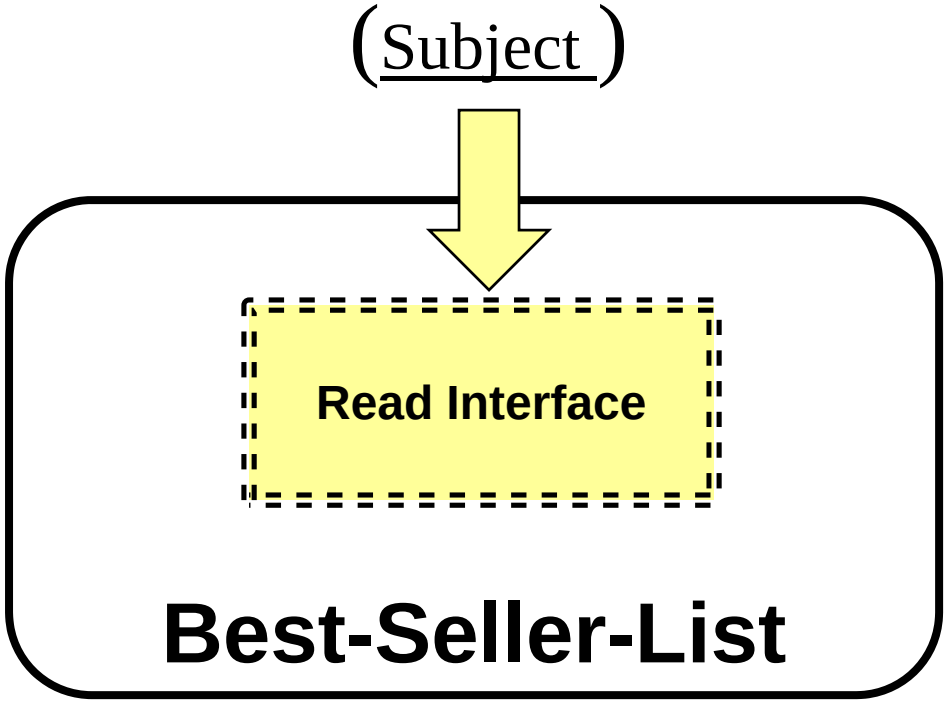

The Best-Seller-List Object

- Availability
- System response time
  - Minimizing communication
- Bandwidth consumption
  - Monitoring all orders at the backend server