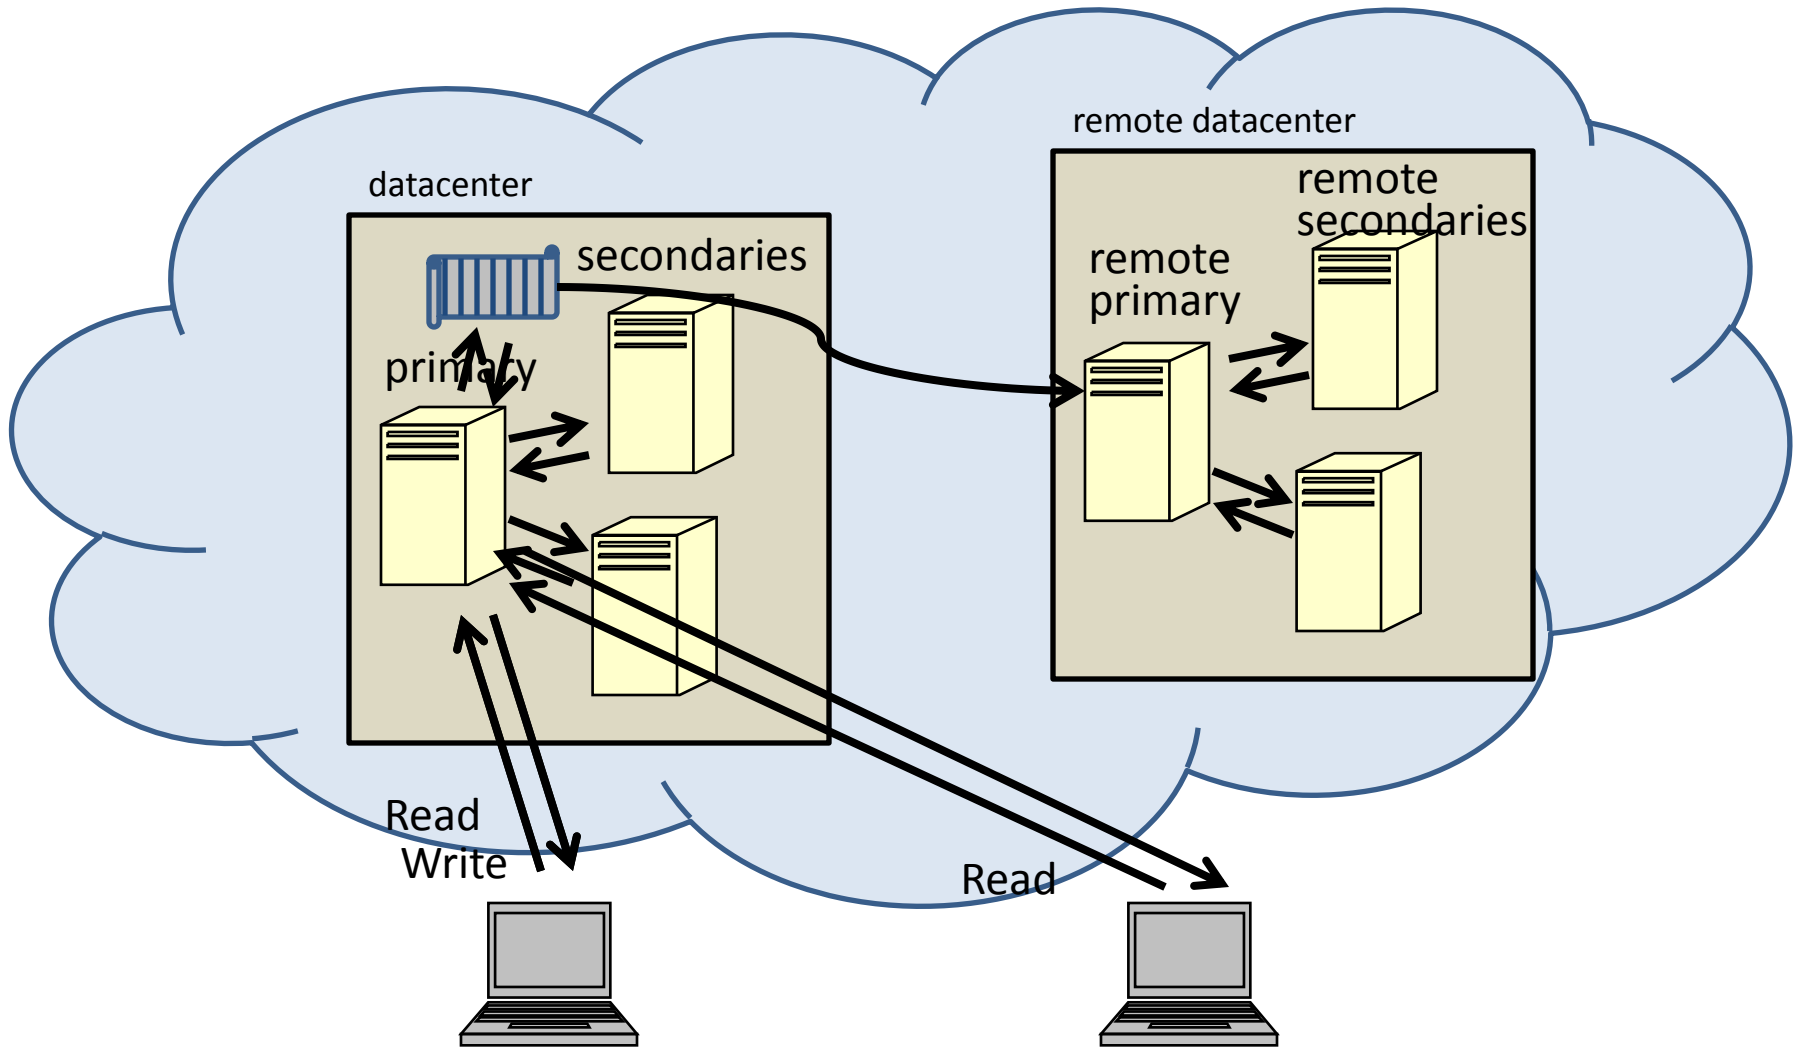


Weak Consistency and Cloud Storage Replication

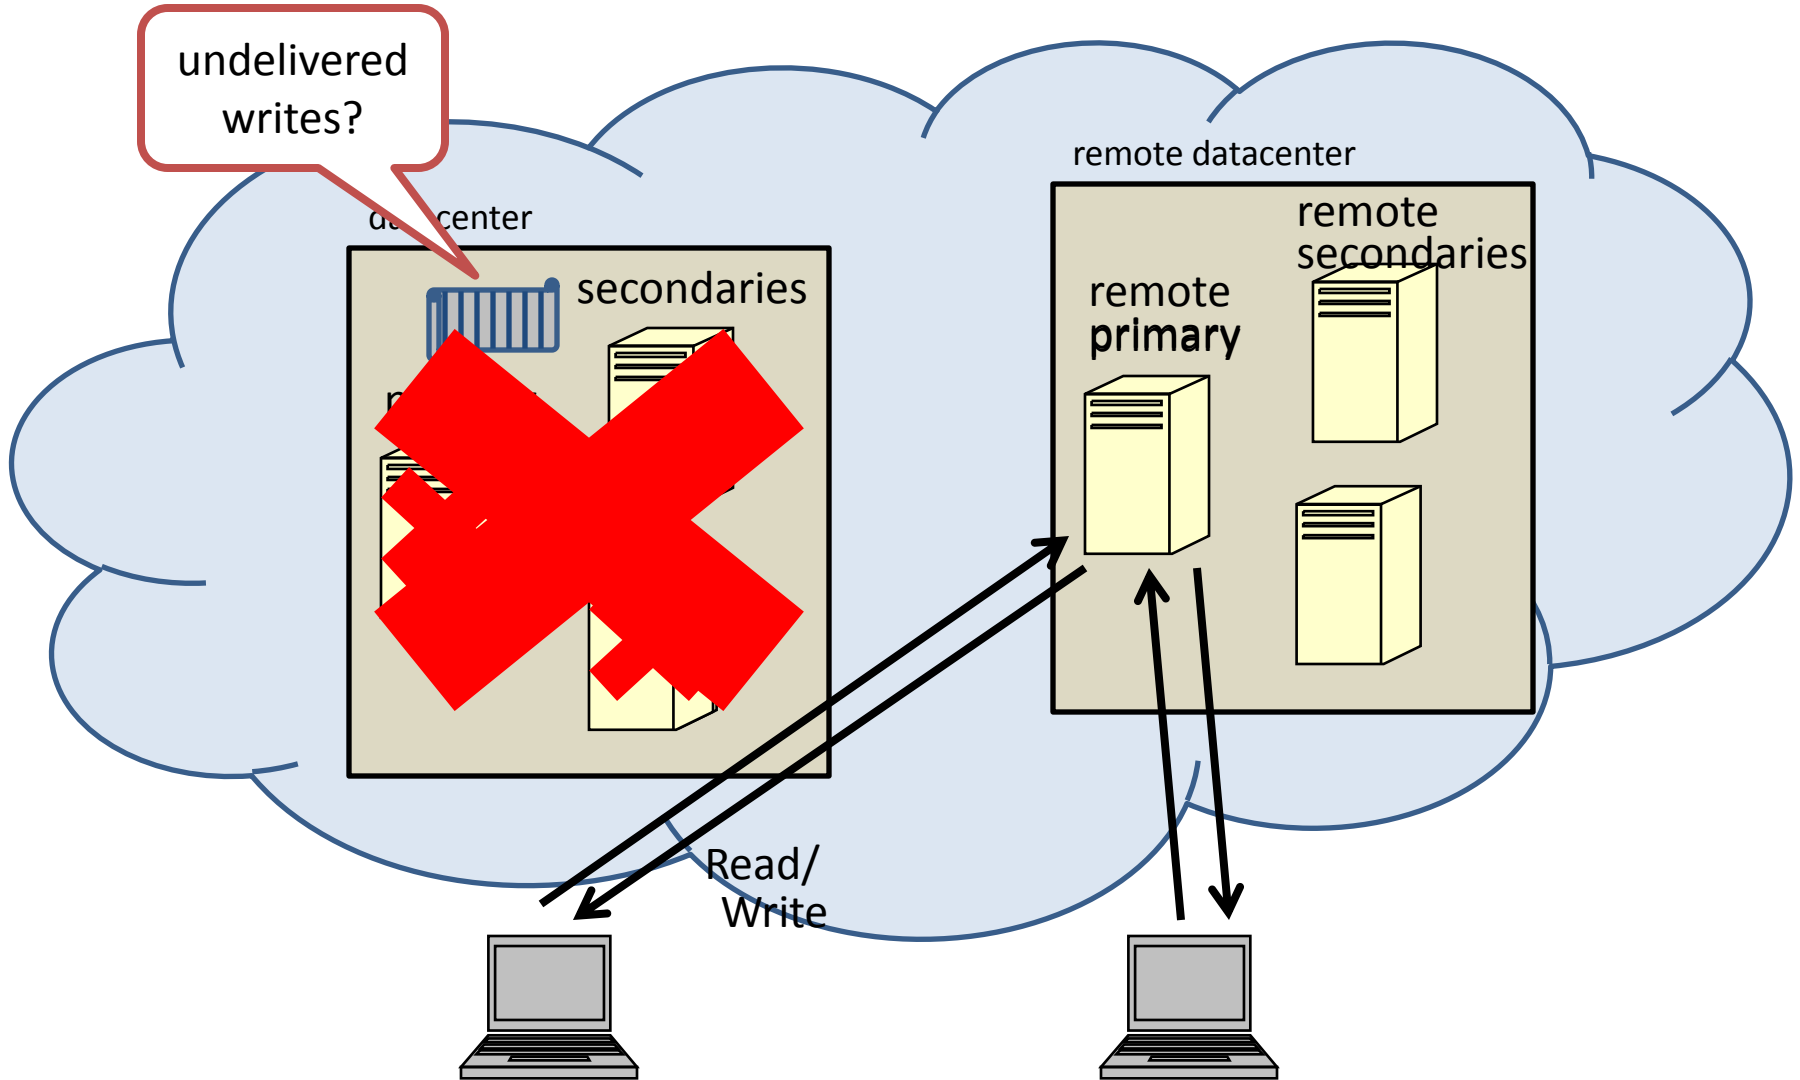
Doug Terry
MSR Silicon Valley

8/2/2010

Cloud Storage Replication

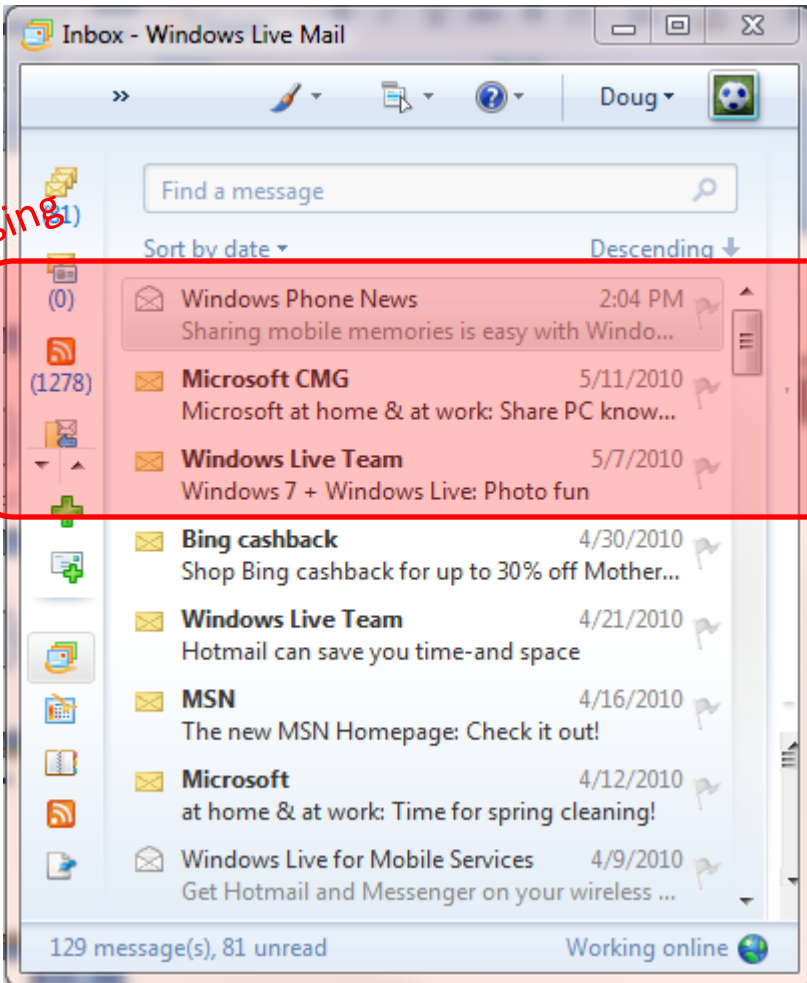


Failure Recovery

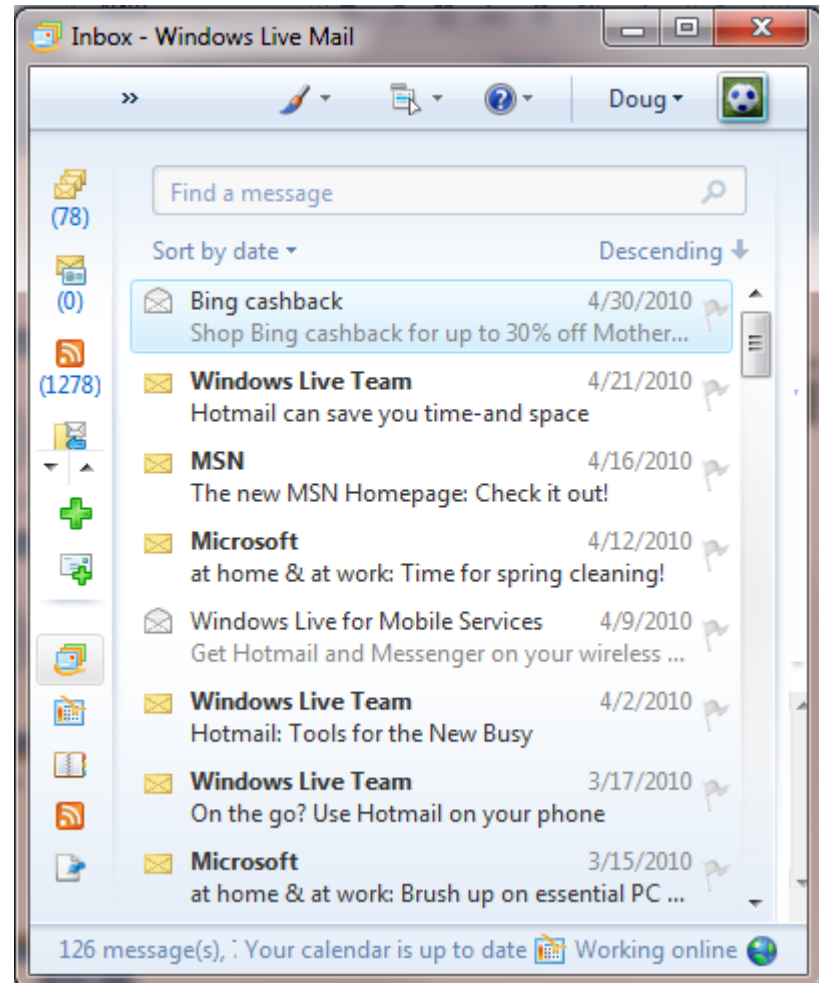


Example: Windows Live Mail

Before failover:



After failover:



Alternative Consistency Models

- Strong consistency
- Consistent reads
- Conditional writes, writes-follow-reads, etc.
- Bounded inconsistency
- Causal consistency
- Epsilon serializability, fork consistency, etc.
- Monotonic reads

Possible Approach

- Maintain eventual consistency model for replication across datacenters
- Detect if a client read operation might violate monotonicity
 - Has the client read data from writes that have not been delivered to the new primary?
- Warn, but don't block, those affected clients
 - Consistency checking vs. enforcement

Conclusions

- Can do better than best effort consistency
 - and better than strong consistency
- But choosing the best consistency requires understanding of
 - Application semantics
 - User expectations/desires
 - Common access scenarios
 - Systems properties: message latencies, failure rates, recovery times, failure correlations, partitions, etc.