



Write Caching with Reduced Durability

Tim Blackman
Sun Microsystems, Inc.



Games present challenges for persistence

- Low latency more important than high throughput
- High percentage of writes
 - > Perhaps 50%
- Large scale
 - > Many concurrent users
 - > Much data
- Need horizontal scaling
 - > No permanent, per-node state

Project Darkstar model

- Represent game logic as small, movable task objects
- Provide automatic retry for task transactions
- Support point-to-point and group communication
- Provide access to data objects via transactional persistence

Darkstar can make use of two advantages

- Can move both computation (tasks) and data
- Modest data loss is acceptable
 - > So long as consistency is maintained

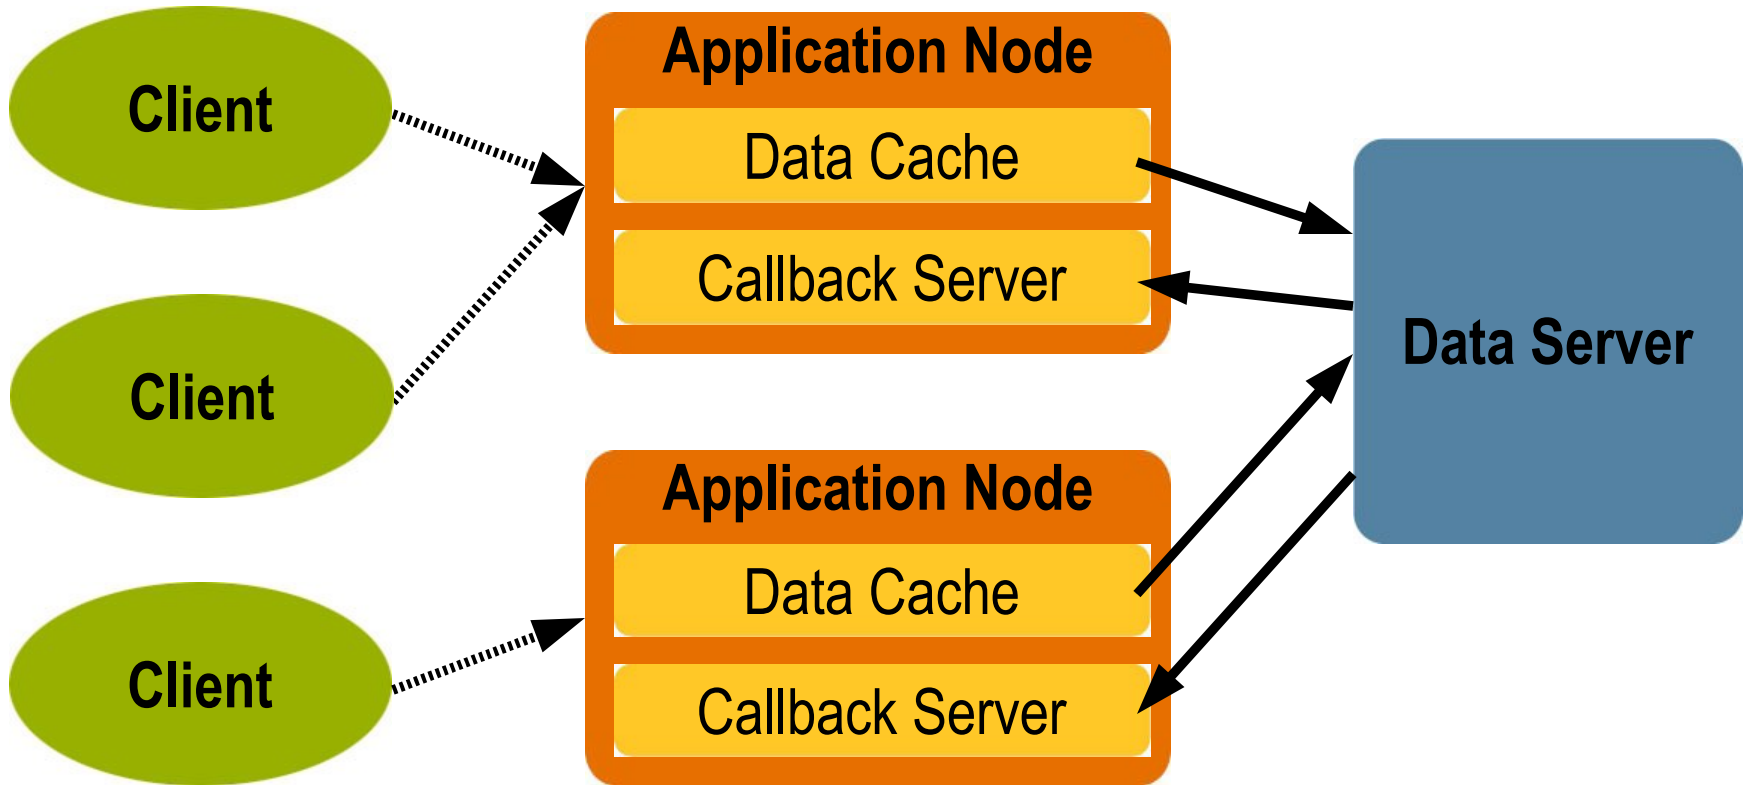
Some not-so-good possibilities

- Single node implementation doesn't scale
- Non-caching distribution is too slow
- Read caching doesn't help
- Local storage is unacceptable

Write caching provides a solution

- Modify cached data in place
- Stream updates to central server on commit
 - > Don't wait for confirmation
 - > Takes advantage of high network bandwidth
 - > Recent changes lost on failure
- Return data to central server on demand
 - > But after updates, to maintain consistency
- Good performance requires good locality

Architecture



Data Cache

- Cache recent items for read and write
- Queue updates to send to server
- Evict items not recently used

Callback Server

- Handle server callback requests
- Evict immediately if
 - > Not in use
 - > No transactions with queued updates
- Otherwise queue response

Data Server

- Track cache contents of data caches
- Call back items if needed
- Queue waiting requests
- Single point of failure
 - > Address in future work

Current status: October 2009

- Initial coding complete
- Testing underway
- Preliminary single node performance looks good
- Technical report
 - > <http://research.sun.com/techrep/2009/abstract-187.html>



Write Caching with Reduced Durability

Tim Blackman

tim.blackman@sun.com

<http://www.projectdarkstar.com>