



High Velocity Data

HPTS
10/23/11

Brian Bulkowski

CEO & Founder
Citrusleaf



Why you need 1M TPS



Winner – Zynga - \$20B

Loser – EA - \$5B





High Velocity Data Apps

Display Advertising

Real-time bidding applications

Cookie matching

Server side user profiles

Frequency capping

Social - Online & social game data

Retrieval of select user histories in seconds

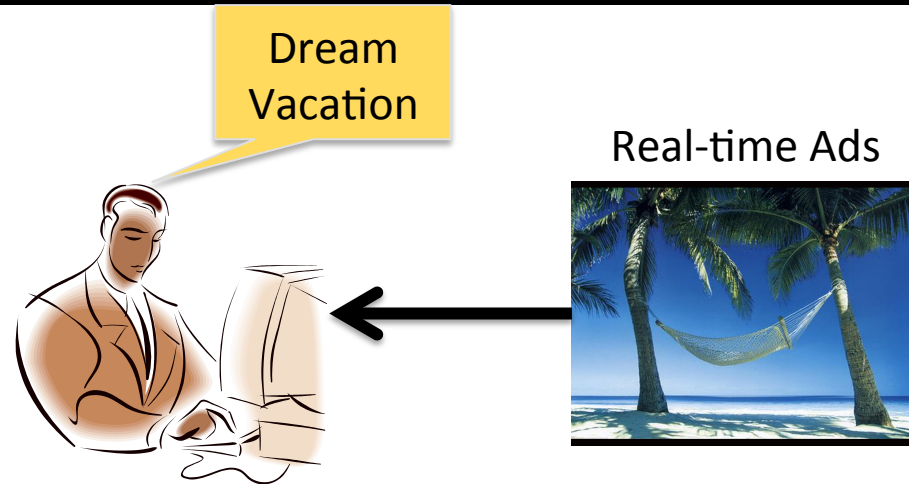
User ID storage & access

High Traffic Web Sites

Session Management



Citrusleaf: Mission Critical NoSQL DB



ACID Compliant

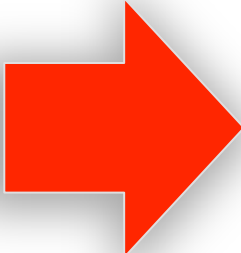
- But only small transactions

High Performance

- 200K+ TPS per node
- Latency < 1ms

Unparalleled Customer Support

- Optional active monitoring
- Quick turn-around fixes



Commercially GA > 1 year
NoSQL solution that
enables mission-critical
internet applications





Digital advertising use cases

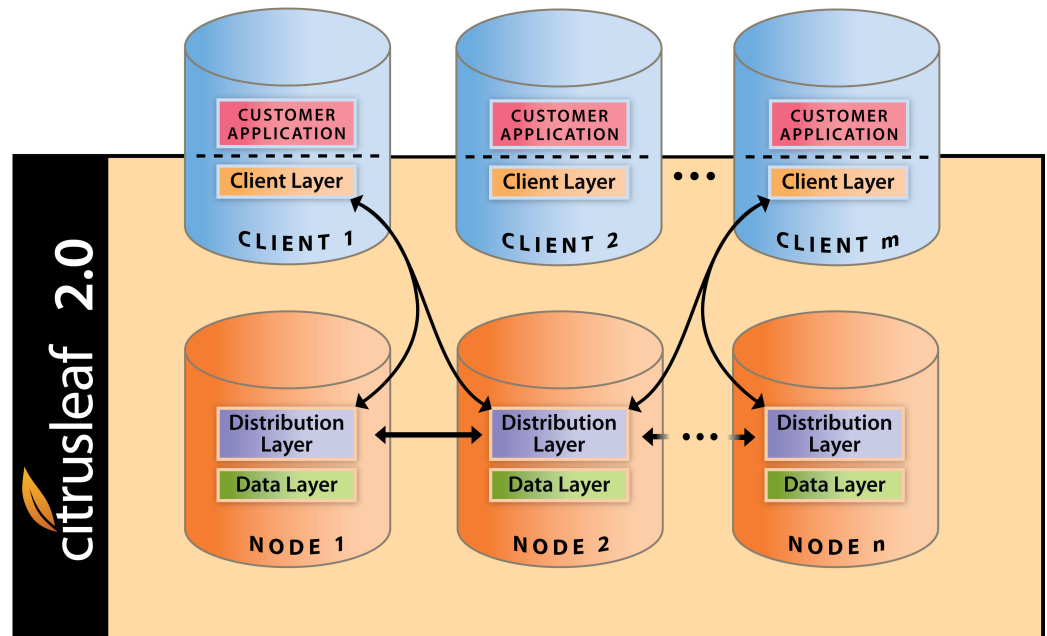
Use Case:	Major Real-Time Advertisement Company	Display Management Platform	Video Advertisement Company
Applications:	<ul style="list-style-type: none">• User Profile Store• Real-Time Bidding	<ul style="list-style-type: none">• User Profile Store• Campaign data	<ul style="list-style-type: none">• User Profile Store• Real-Time Bidding
Environment:	<ul style="list-style-type: none">• 20 servers per DC• 2.5B objects• 2.0T data (6T)• 100K TPS per DC• 3 data centers• private cloud• SSD - userdata• DRAM/HDDs – map	<ul style="list-style-type: none">• 10 servers per DC• 1.2B objects• 3.0T data (11T)• 75K TPS per DC• 3 data centers• own datacenter• SSD - userdata• DRAM/HDDs – map	<ul style="list-style-type: none">• Public cloud (EC2)• 200M objects• 300M data• 20K TPS per DC• 2 data centers• Initial deployment on DRAM/HDD to be migrated to own datacenter





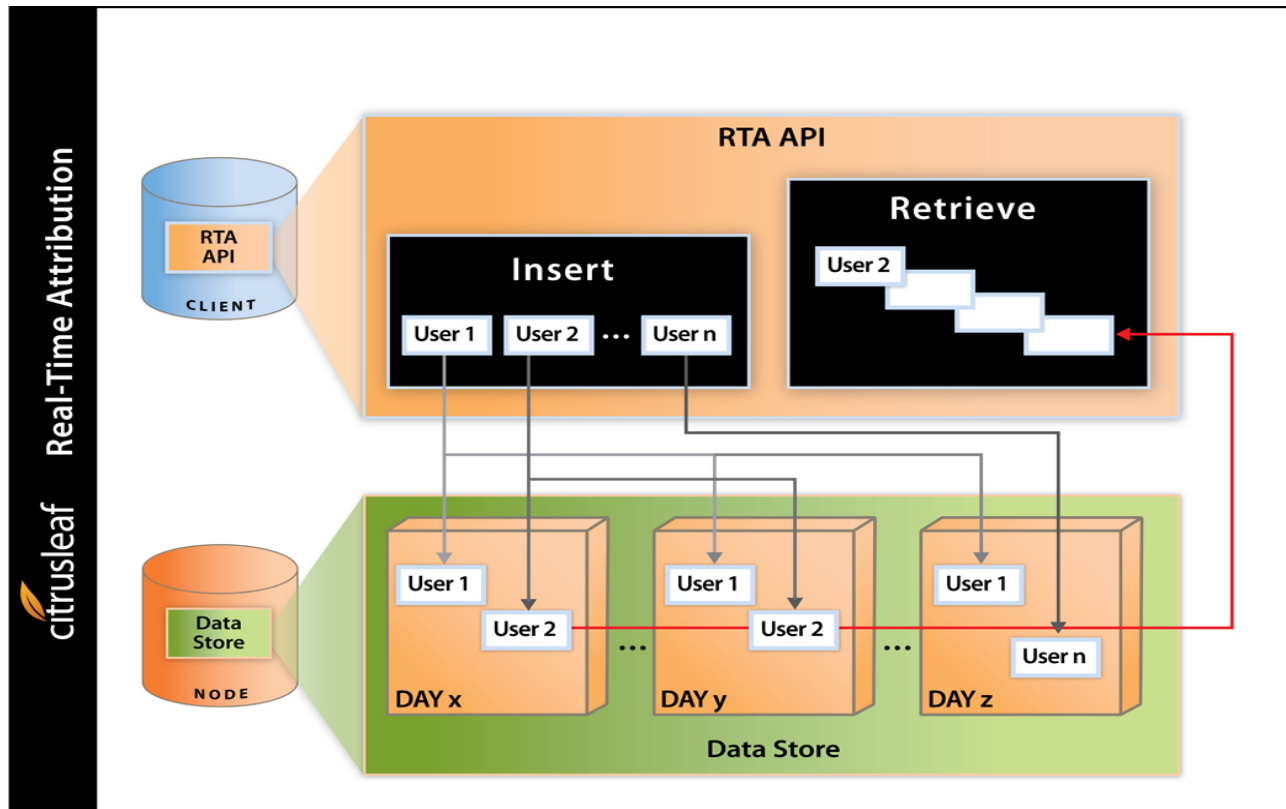
Citrusleaf Architecture

- Primary key index
(document model)
- Client based balancing
- CAP “trick”
(make A small)
- Automatic rebalancing
- 24X7 uptime today
(enterprise support)
- Cross data center (BCP)
- Commodity hardware (TCP)





Use Case: Real-time Attribution



100 billion objects


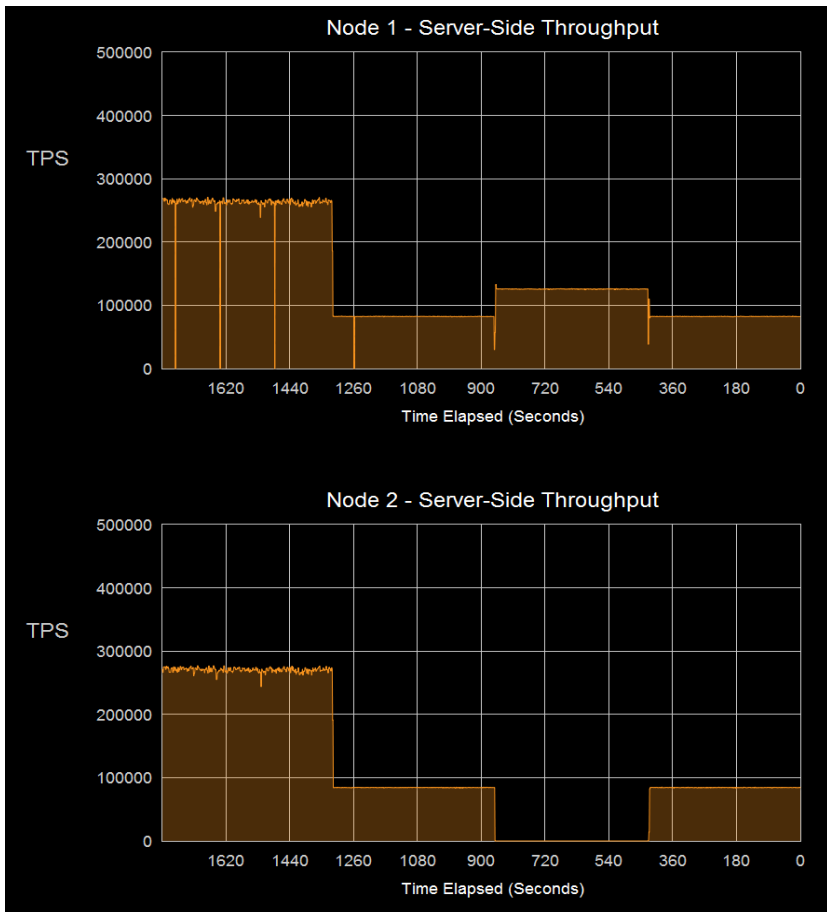
latency < 1 second

instant report





“Live” Demo



Intel i5 2400 – 8GB – \$600 per server
20 byte keys, 8 byte values, in memory – 250K tps per server





Final thoughts

- If you can do 1M TPS on four servers, why do you need a 100 node cluster?
- Intel dual sockets bottleneck on QPI latency (and Amazon doesn't allow NUMA)
- For SSDs, measure read latency under write load (ask us about our SSD benchmark tools)
- Did I mention no downtime?





Citrusleaf

Brian Bulkowski

CEO & Founder
Citrusleaf

Free trial downloads at citrusleaf.net

