



Using Flash Memory in Low-Latency Trading

Matt Key
IBM

About Bids Trading



- Alternative Trading System provider
 - 'Dark Trading' to mask the impact of block trading
- Founded 2007 as consortium of 13 leading firms
 - Now over 40 sponsoring brokers
- Average Trade Volume at 85M Shares
 - ~\$500K trade revenue in relatively short window



Built on Solid State

- Relatively small dataset (fits in single-server RAM)
 - But persistence required
- Quickly moved to solid-state infrastructure for 70% quicker I/Os and 50% faster batches
- Growth lead to expand into a flash-backed infrastructure
 - Could not compromise performance of previous RAM-based storage

Shared infrastructure

- Originally built on clustered servers and mirrored storage
- Evolved to include clustered servers as well as application level mirroring
- Flash naturally presented higher capacities which is used in off-peak time for batch

- Flash in low-latency trading is supplemental to DRAM
- Most exchanges operate in single-digit millisecond response times at the app level
 - Native inconsistencies of flash media prohibit direct use
 - Flash allows for smaller footprints of DRAM as time-to-persist requires little backup power

Major US Exchange

- Built on persistent DRAM SAN technology
- Evolving to distributed DRAM+flash
- Would love to be on NV-DIMM technology
- Primary barrier is commercialized software
 - Loves open source but standardized on Oracle/SQL

- Latency, Latency, Latency
- Most apps are low-thread count and synchronous
- Solid state has raised scrutiny of “shaving microseconds”
 - RDMA vs IP
 - SRP vs FC