



**FUSION-io**  
POWERING INNOVATION

# Tuning For Speed: Percona Server and Fusion-io

Torben Mathiasen - Percona Live 2011 NYC  
[tmathiasen@fusionio.com](mailto:tmathiasen@fusionio.com)

# Introduction



- Databases and Fusion-io
- The moving parts that determine performance
  - Hardware
  - Kernel
  - File systems
  - Applications



- **Benchmarks vs Applications**
  - Advertised, benchmark and application performance
  - Flexible IO tester and innosim
- **XtraDB, InnoDB and MyISAM**
  - We really like XtraDB and InnoDB
  - MyISAM is less complex, small tables may benefit from less locking overhead.

# Basic Tuning



- Hardware/Software used
  - 48 cores (AMD 6174, 2.2Ghz)
  - 128GB RAM (adjusted downwards for specific tests)
  - RHEL6
- File systems
  - EXT4, trim/discard
  - XFS, no online discard yet
  - Write barriers
    - No cache on ioMemory modules. Acknowledged writes are guaranteed to be persistent. Even during power failure, crashes.
- Improving steady-state write performance and increasing endurance

# Basic Tuning Cont.



- I/O Schedulers
  - Default with Fusion-io is NOOP
  - Merges requests but does not optimize for rotational devices
  - For MySQL, submitting requests directly to the device have shown performance improvements. `use_workqueue=0`
- MSI-X Interrupts vs PIN based
  - Performance boost for some configurations. `disable_msi=0`
  - Decreases CPU utilization
  - `tintr_hw_wait=50` to coalesce interrupts (default 10us)
- Outstanding I/O
  - Default is 128 reads and 128 writes per queue. Batching can increase this further. Can make a difference with benchmarking. No clear advantage for MySQL
- Distribute Interrupts
  - Distribute interrupts among multiple cores for multi card setups to avoid soft lockups -> `irqbalance`

- How to test? Which workloads?
  - Settle on a few, concentrate on those and then verify in production. Production data is great, but may not be available
  - We have been using Percona tpcc-mysql and tpce-mysql. Good workloads, but may not match yours
- Direct I/O vs Buffered
  - Use direct-I/O for data. Logs are usually buffered
  - Buffered I/O introduces an extra data copy. Memory for page cache better spent in buffer pool. Less complex write path
  - When does buffered I/O makes sense?
- Block sizes
  - Linux MD on RHEL6 seems to favor 4K sectors and blocks
  - Marginal better efficiency of the Virtual Storage Layer with 4K sectors

# MySQL Tuning Cont.




- Check pointing -> heavy writes
  - Adaptive, push I/O constantly.  
InnoDB\_adaptive\_checkpoint=keep\_average, keep short loop cycle
  - Buffer Pool vs Dirty pages
    - Large buffer pool can mean a lot of dirty pages. A LOT!
- Mount points
  - Sequential vs Random I/O
    - Spindles are ok at sequential platter access. Flash is better
    - A 500GB database is a big database. But a 640GB io-duo can fit it all
    - If possible, put everything on flash. If needed, put logs on hard drives
  - Spindle hosted database logs
    - Make sure that the spindle does not receive other I/O. Multiple database logs may be sequential individually, but still cause the heads to seek
- Recovery time vs Log file sizes
  - Large logs to improve performance and lessen checkpoint activity

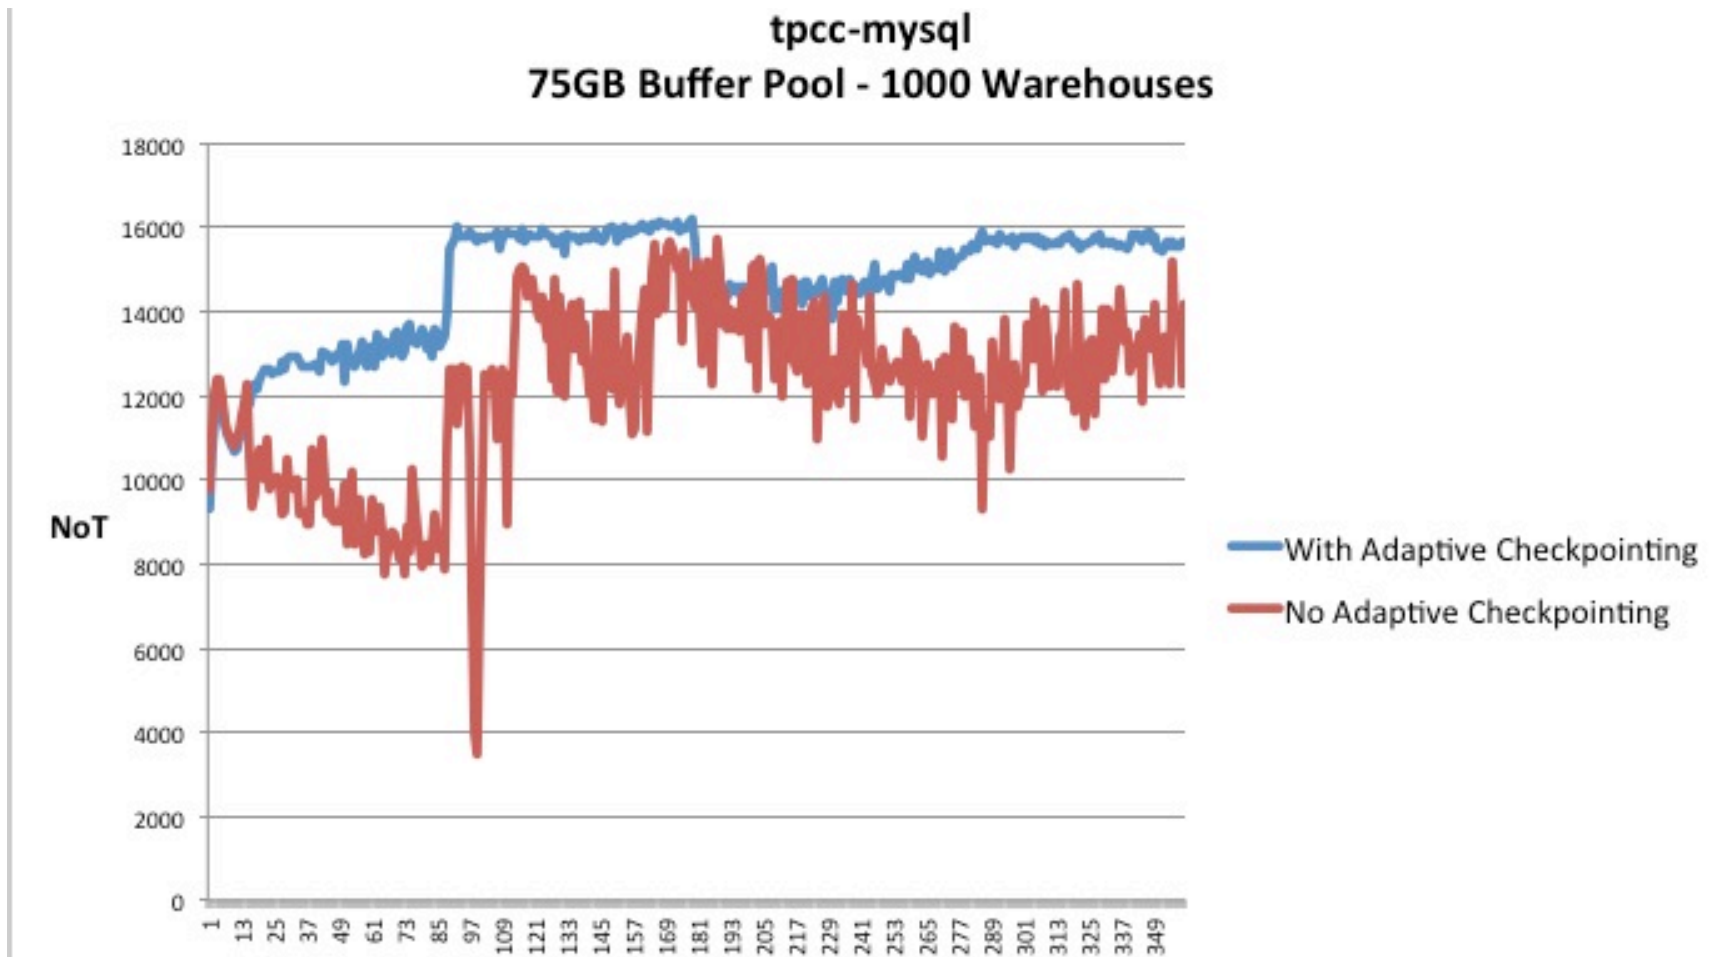
# What simple optimization can



Percona 5.1.53  
TPC-C 1000W  
75GB Buffer Pool



# Adaptive check pointing/




# Sector and Block sizes



tpcc-mysql

8GB Buffer Pool - 500 Warehouses



- Configuration Parameters
  - `innodb_thread_concurrency=0`
  - `innodb_read_ahead=0`
  - `innodb_read_io_threads=8` (16 for writes)
  - `innodb_adaptive_flushing=1`
  - `innodb_adaptive_checkpoint=keep_average`
  - `innodb_max_dirty_pages_pct=60`
  - `innodb_flush_method=O_DIRECT`
  - `innodb_io_capacity=10000`
  - `InnoDB_flush_neighbor_pages=0`

# Closing remarks



- The Challenge
- Community efforts, development and future prospects
  - Linux Multiqueue (2.6.x)
  - Linux Explicit I/O unplug (2.6.39)

Lots of improvements being made all the time.

- Percona/MySQL 5.5 – Looking forward to



**FUSION-io**  
POWERING INNOVATION

**Questions? – Thank You!**

Torben Mathiasen – Percona Live 2011 NYC  
[tmathiasen@fusionio.com](mailto:tmathiasen@fusionio.com)