LSI LOGIC

- Training on each I/O limits performance
  - Keep option open to train less often
- Eliminating QAS limits performance
  - Keep QAS for those target / initiator pairs that want to go fast
- Incremental changes to SPI-4 eventually add up

# Performance impact depends upon :

- I/O size → Significant for ≤ 8 KB I/Os
- I/O direction (read / write) → Higher impact for Write on larger transfers
- Queue Depth → Impact increases as queue depth < 8
- Target head movement → Impact increases as head movement decreases

Do we create a specification that optimizes a specific operational environment, or one that can work well over a wide range of environments ?

LSI LOGIC

- Evaluate:
  - P + QAS — baseline Packetized, QAS, no training
  - P + QAS + Training — Packetized, QAS, 2us training
  - P + Training — Packetized, no QAS, 2 us training

- This analysis does not consider Parallel option, with or without training

Training & QAS impact < 2% at 320 at large blocks, at 10 Q'd I/Os

Impact of Training 9% at 8KB, 4 Q'd

Impact of QAS another 9% at 8KB, 4 Q'd

Impact of each Training & QAS is 14% each at Ultra640 - it gets worse!

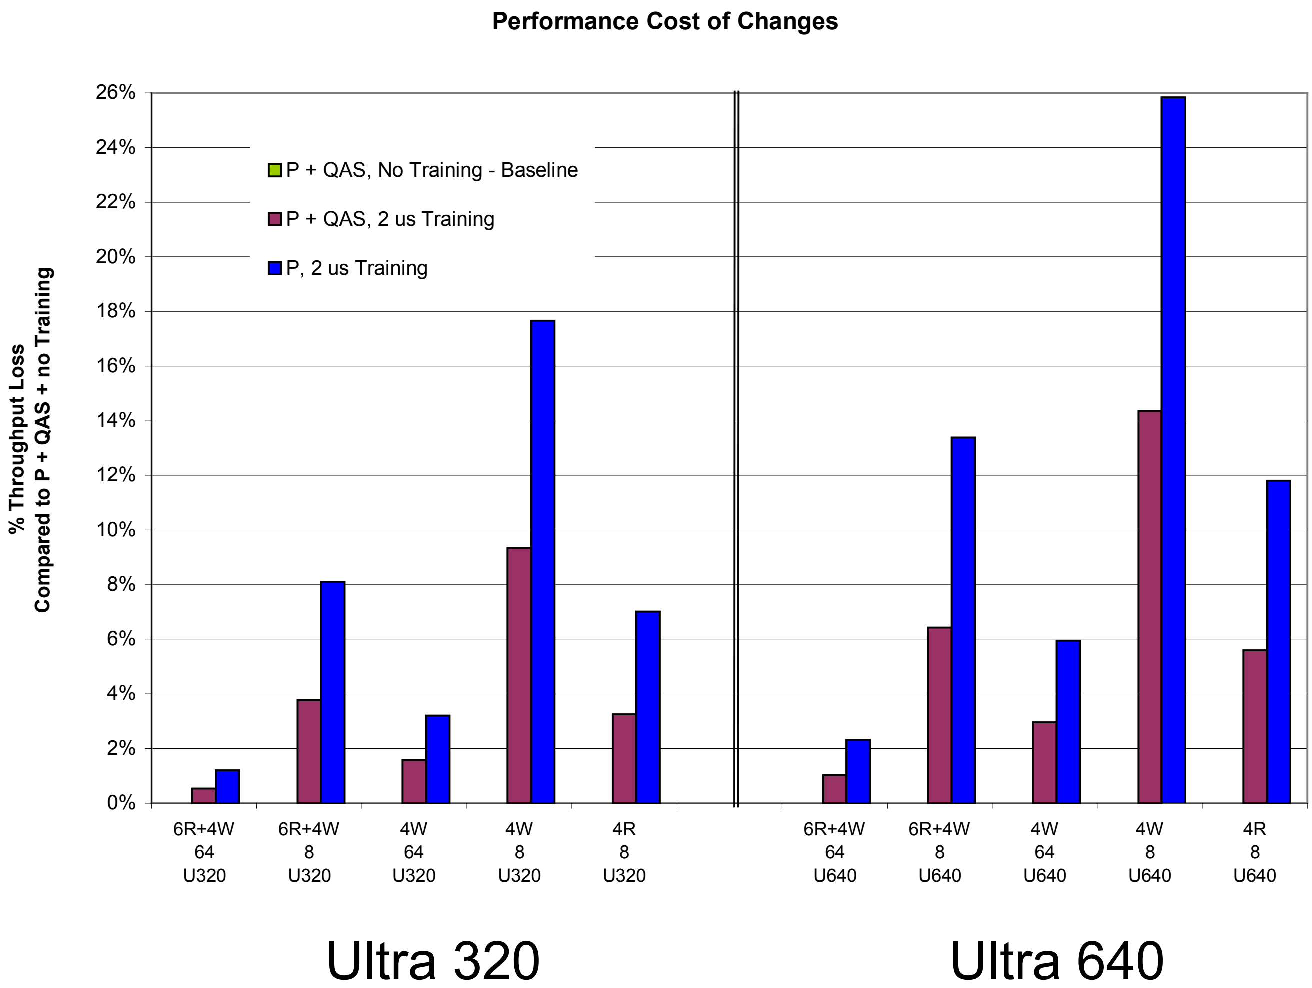


I/O size in Kbytes

4W = 4 Queued Write I/Os

LSI LOGIC

Impact of changes is small at Q Depth > 8. Today Q Depth = 1 for sequential. Which applications will change the way they issue I/Os for Ultra320?

Impact of Training & QAS each are significant at smaller Queue depths, together they can cripple SCSI

Lots of I/O happens with blocks of 4K to 8KB

These problems get worse at Ultra640

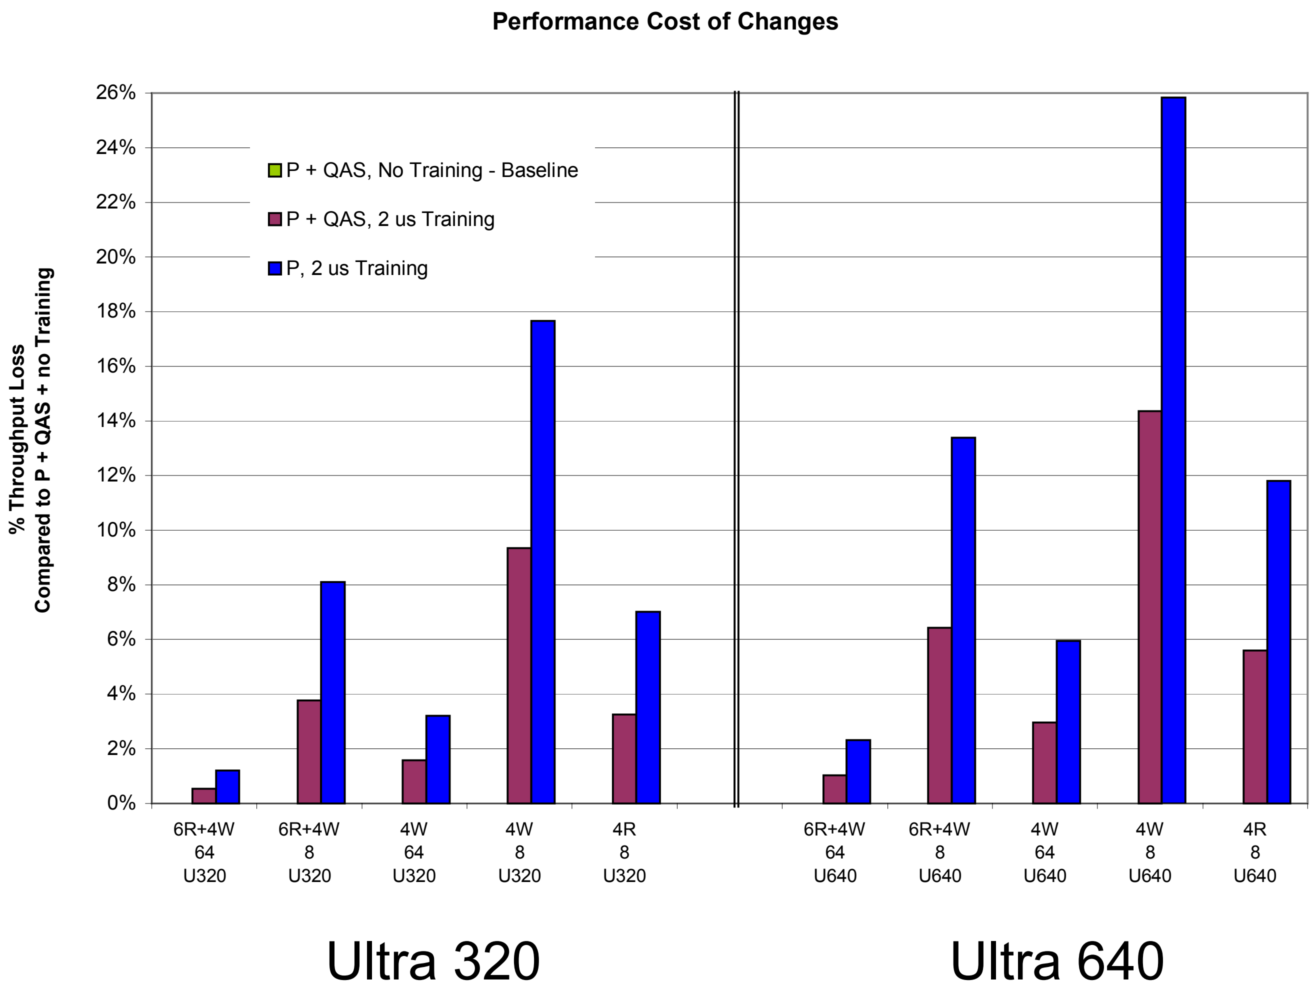


I/O size in Kbytes

4W = 4 Queued Write I/Os

LSI LOGIC

Theoretical results - only counts SCSI protocol timings

Realistic IOPs values are 40 - 50 % lower than these theoretical results at 0.5 KB

Ultra320 with Packetized & training will be slower than Ultra160 Non Packetized at 1 Q'd I/O per drive - same test as we use today !

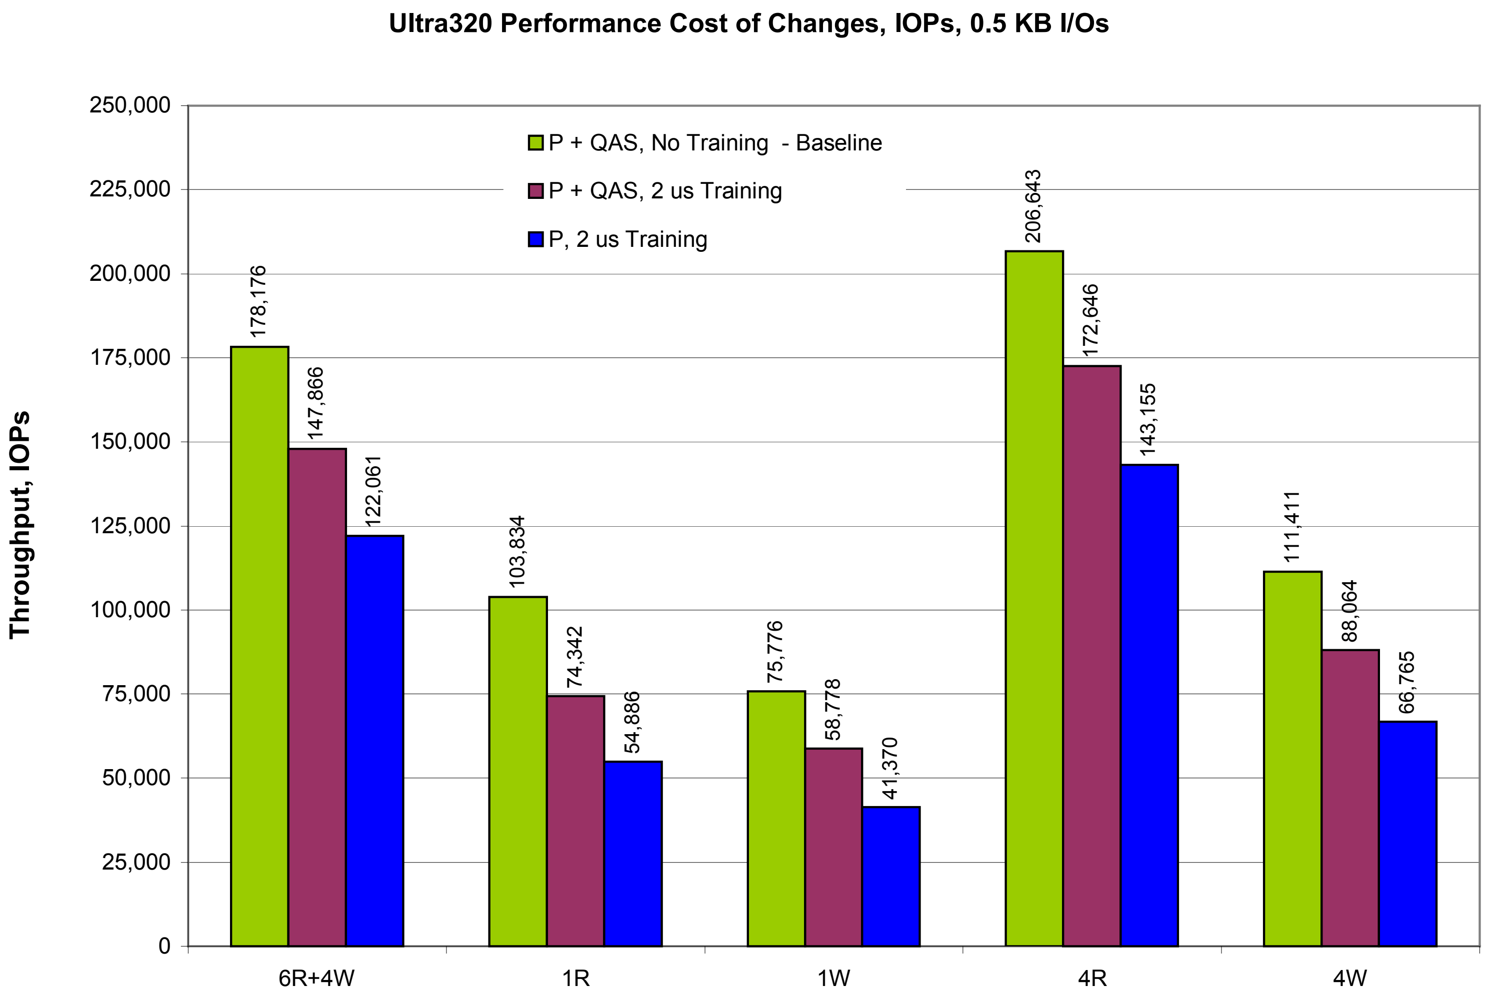


6R+4W = 10 Q'd I/Os, 6 read, 4 write

4W = 4 Q'd I/Os, all write

LSI LOGIC

- Removing QAS simplifies design and validation
    - …. But ...
  - Reducing impact of no QAS requires high queue depths for sequential I/O
  - High queue depths increases I/O latency
  - Heroic seek optimization may not be as effective as expected
  - Will applications be rewritten to make use of large queue depths ?
  - Will early implementations use queue depths > 1 for sequential I/O ?
  - Assuming high queue depths also assumes specific workload profiles - not a general solution
    - …. This trades 1 problem for a new set

LSI LOGIC

Theoretical results - only counts SCSI protocol timings

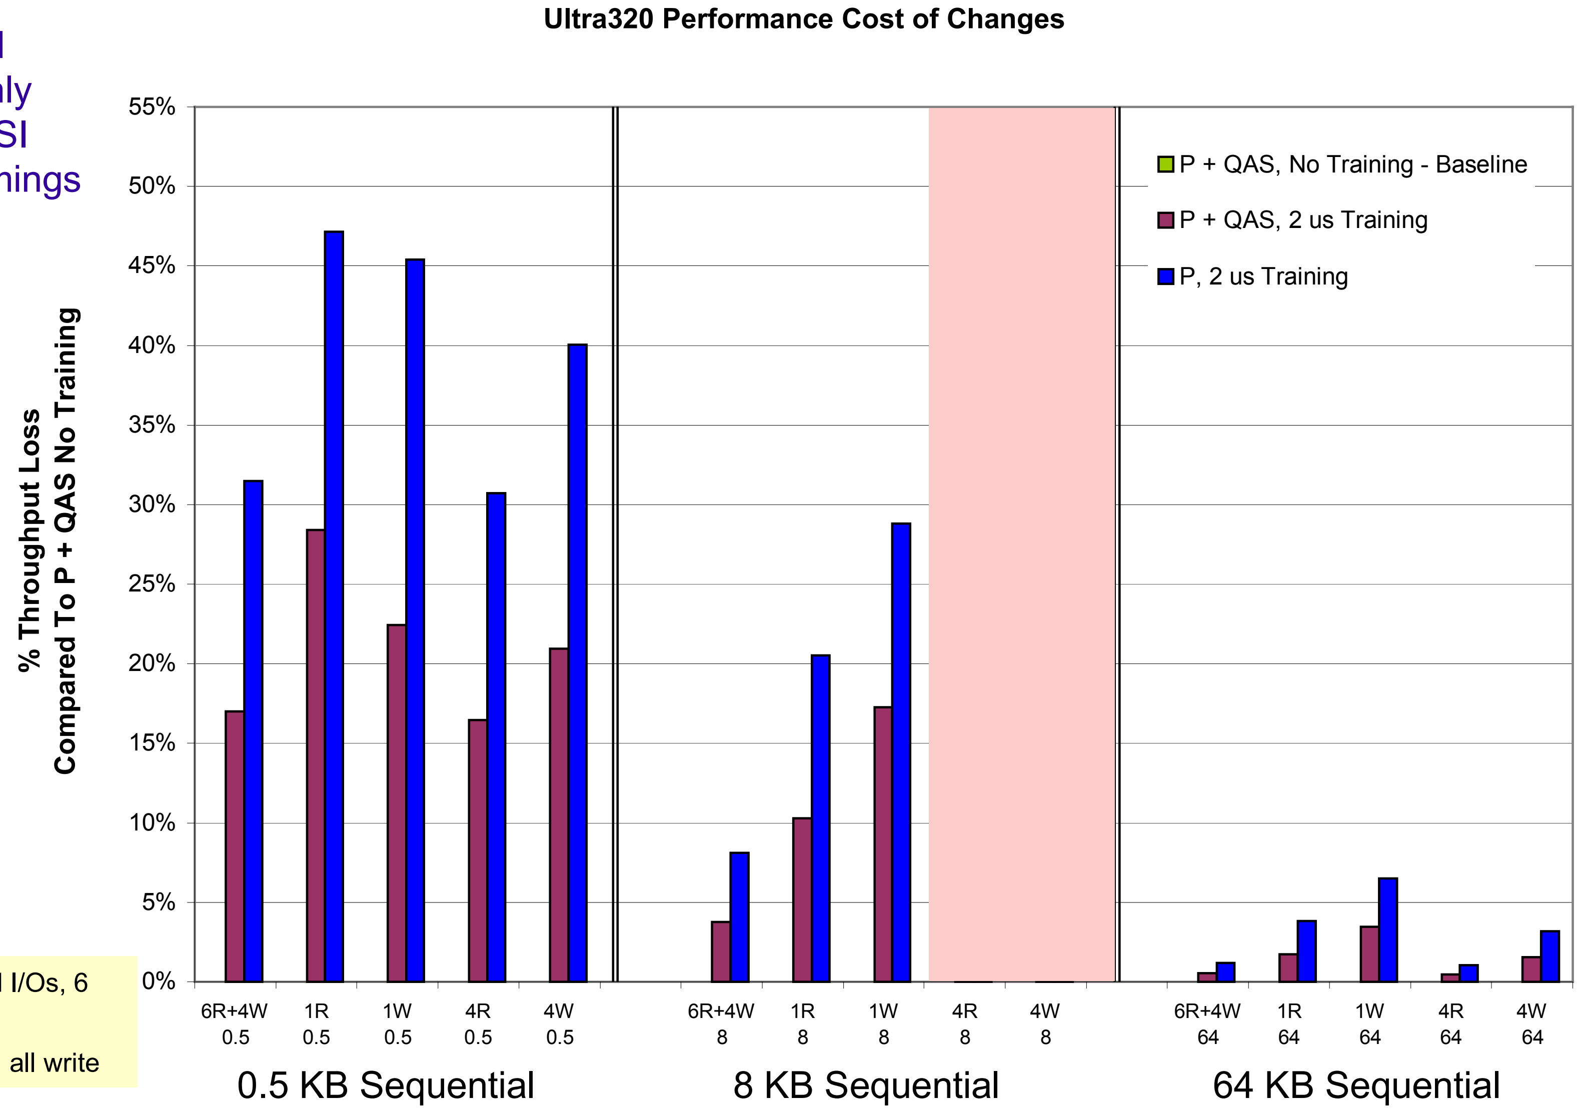


6R+4W = 10 Q'd I/Os, 6 read, 4 write

4W = 4 Q'd I/Os, all write

Theoretical results - only counts SCSI protocol timings

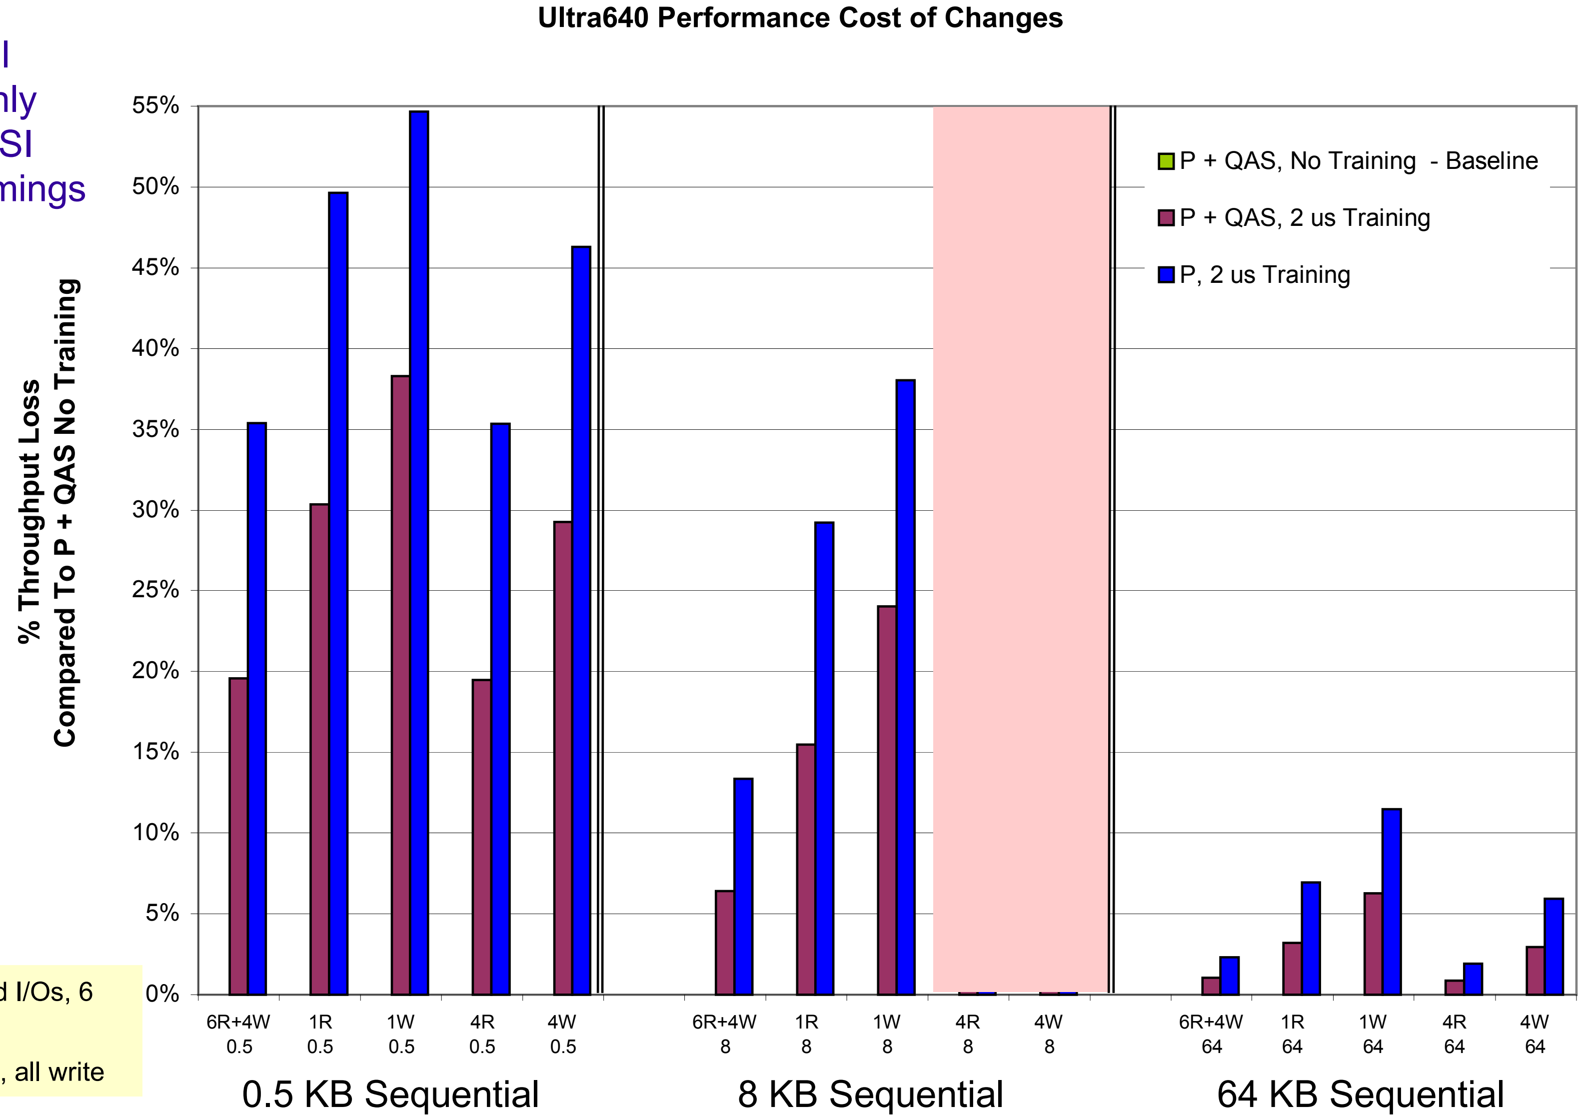


6R+4W = 10 Q'd I/Os, 6 read, 4 write

4W = 4 Q'd I/Os, all write

LSI LOGIC

Theoretical results - only counts SCSI protocol timings

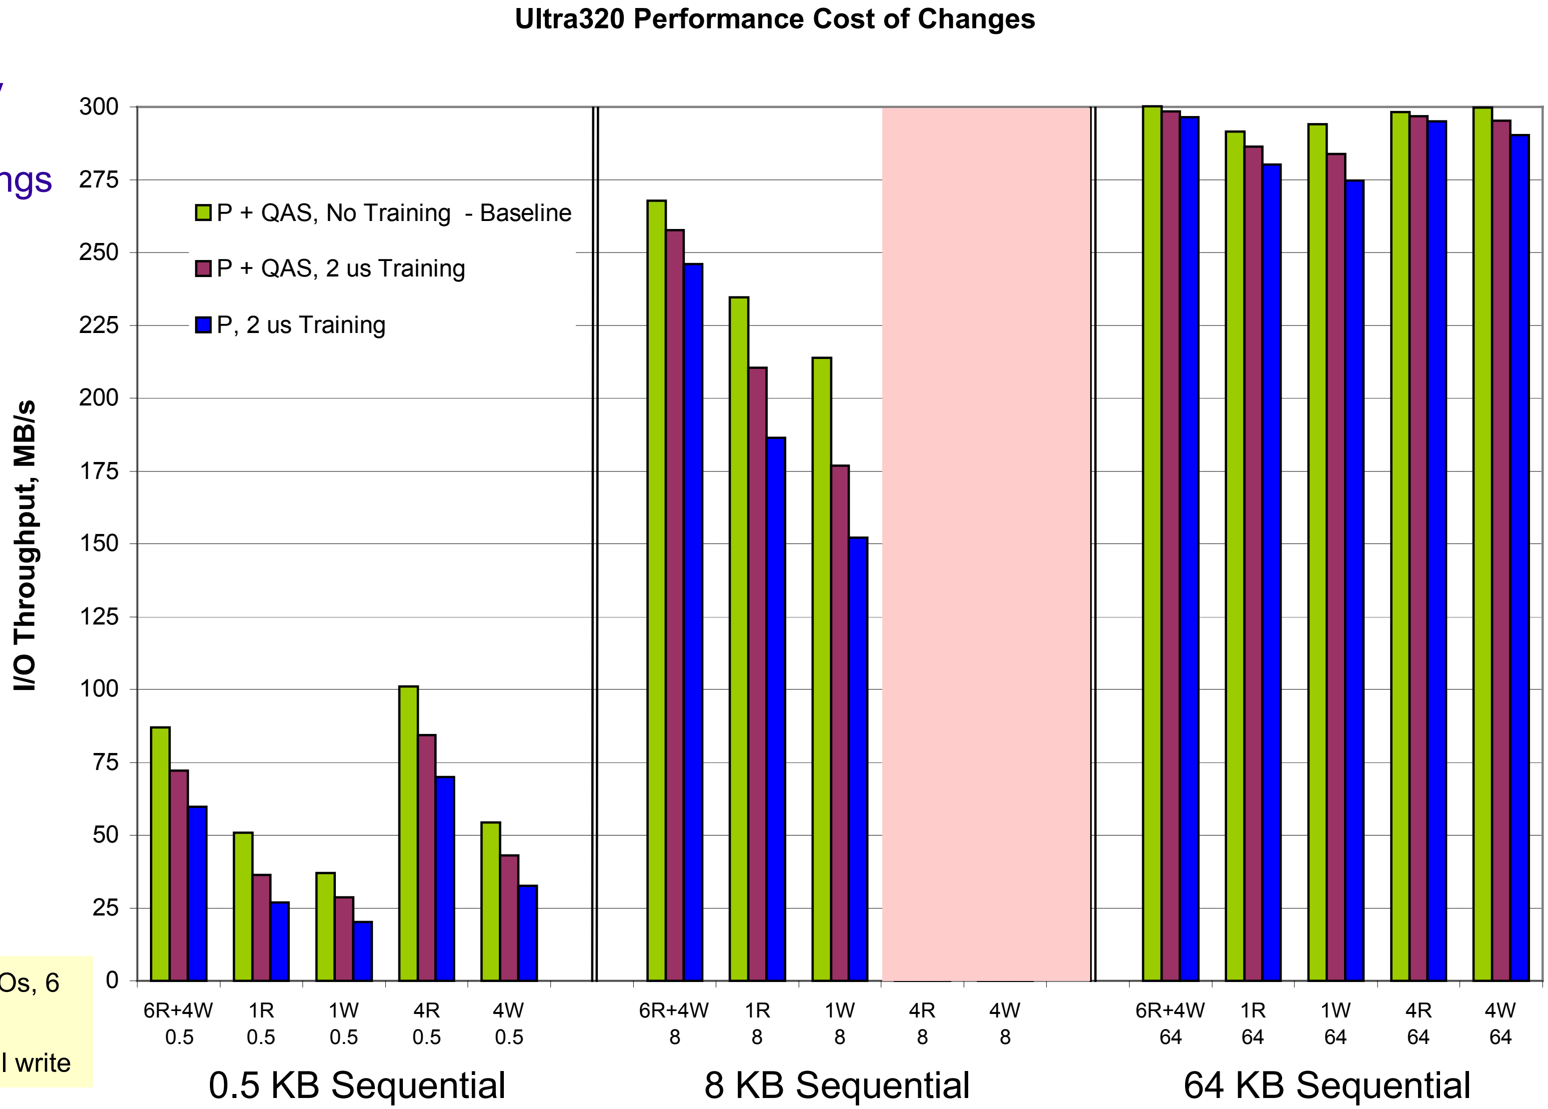


6R+4W = 10 Q'd I/Os, 6 read, 4 write

4W = 4 Q'd I/Os, all write

The Communications Company ™

LSI LOGIC

Theoretical results - only counts SCSI protocol timings

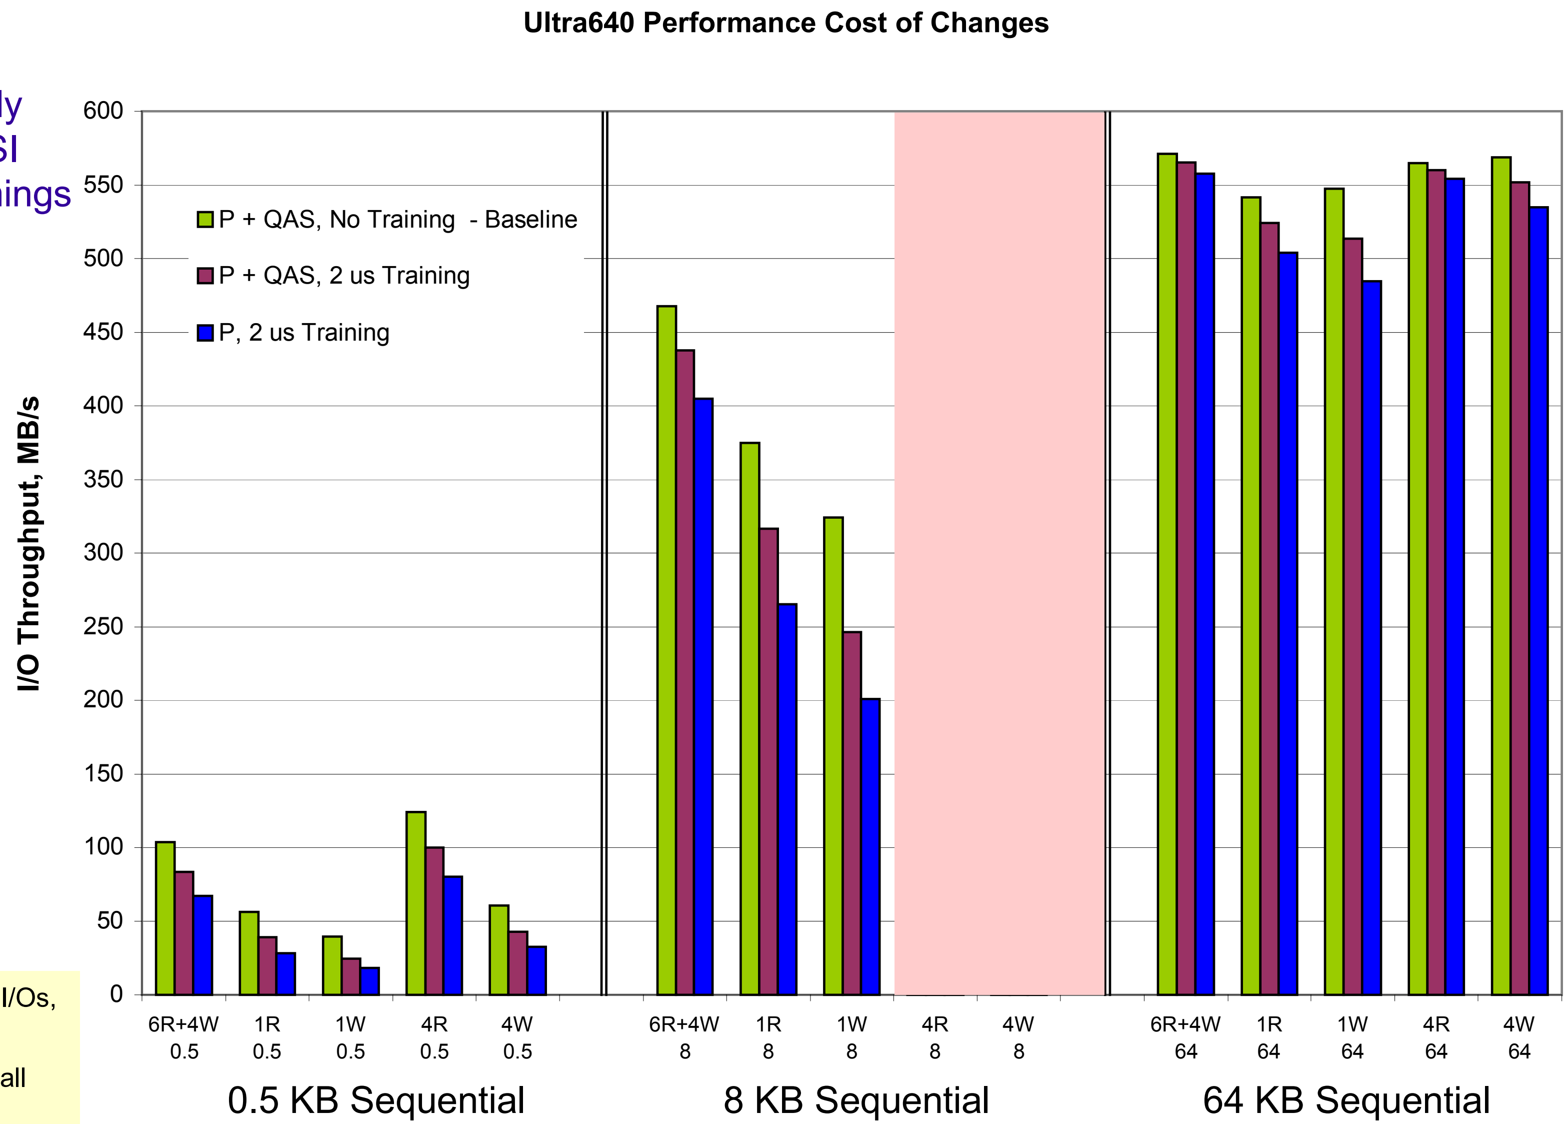


6R+4W = 10 Q'd I/Os, 6 read, 4 write

4W = 4 Q'd I/Os, all write

Theoretical results - only counts SCSI protocol timings

Realistic IOPs values are 40 - 50 % lower than these theoretical results at 0.5 KB

Ultra640 with Packetized & training will be slower than Ultra160 Non Packetized at 1 Q'd I/O per drive - same test as we use today !

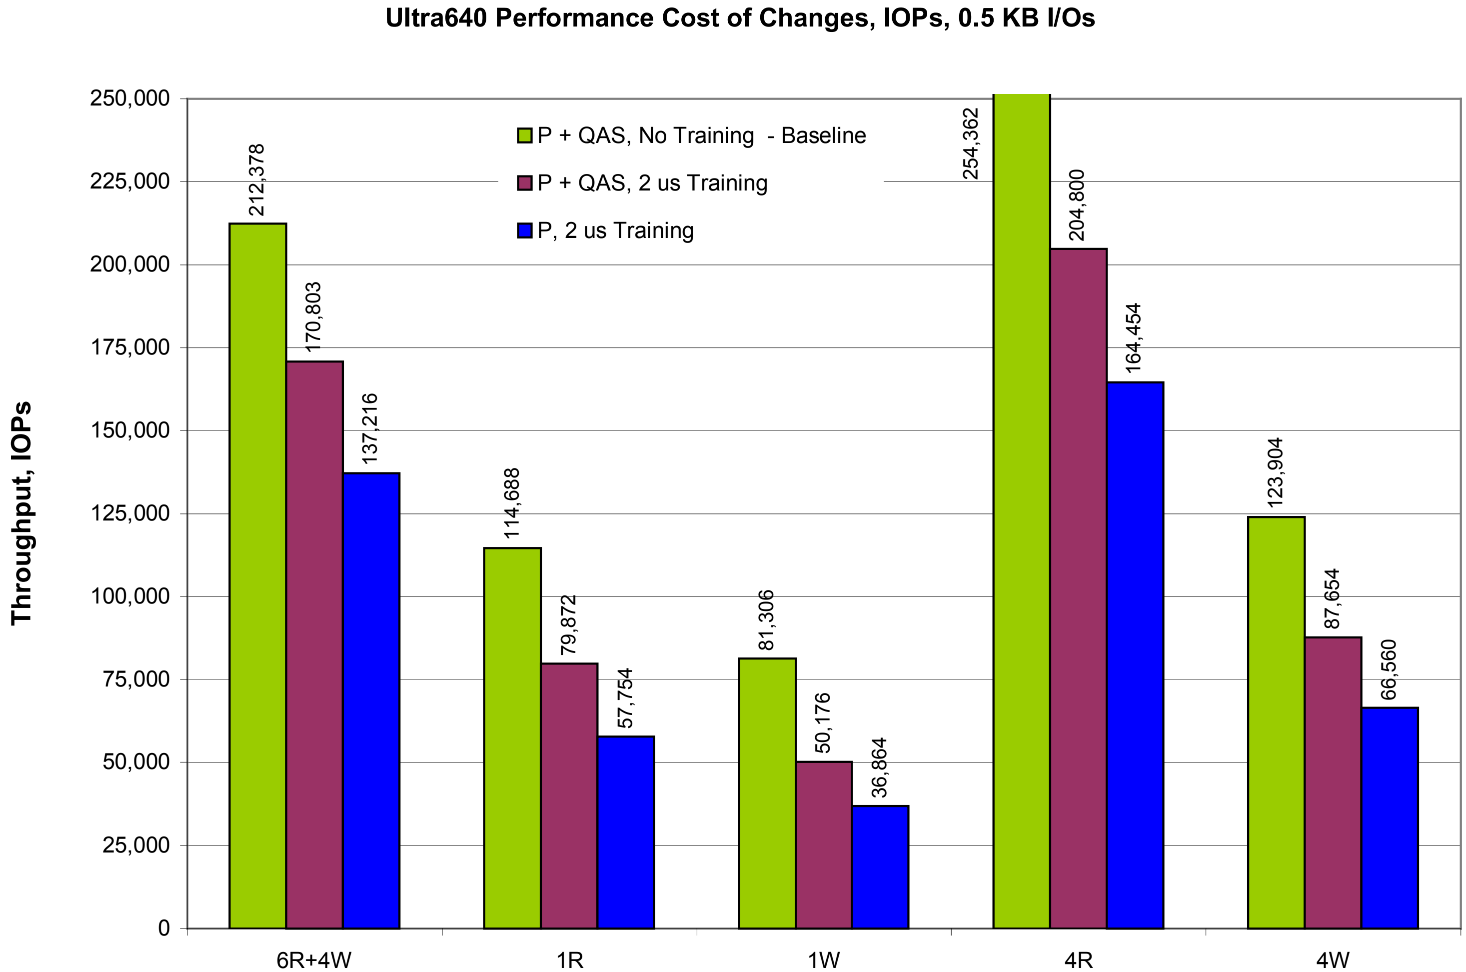


6R+4W = 10 Q'd I/Os, 6 read, 4 write

4W = 4 Q'd I/Os, all write