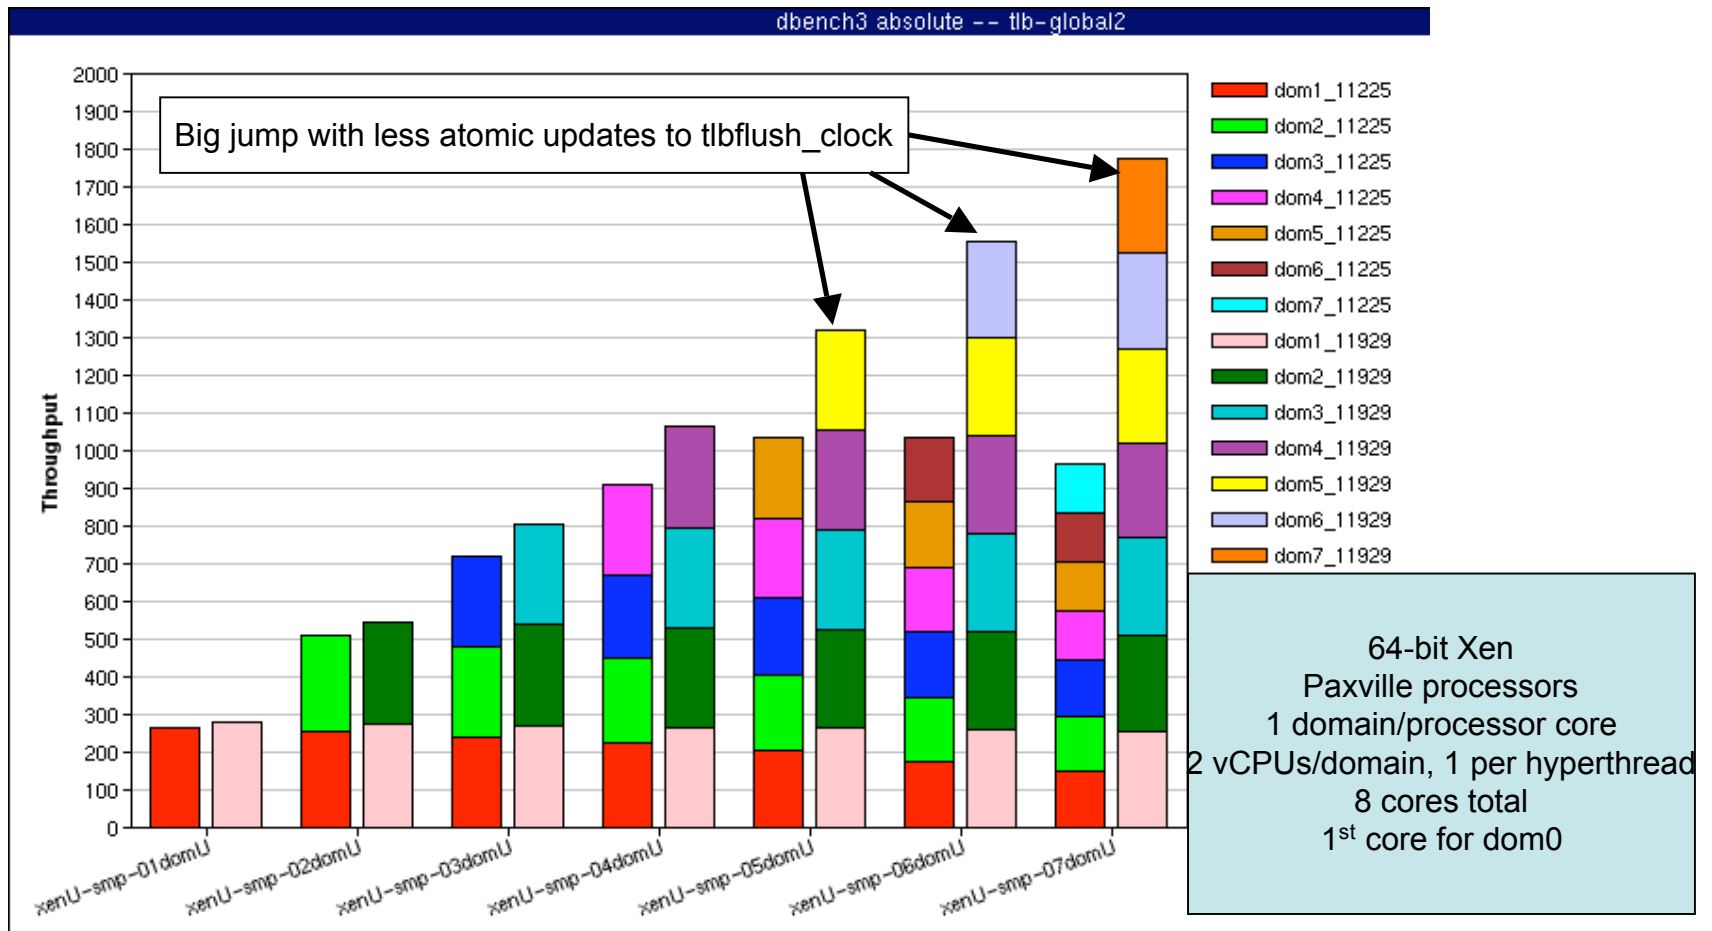


# Xen Summer 2006 Summit Performance Update

Andrew Theurer, [habanero@us.ibm.com](mailto:habanero@us.ibm.com)

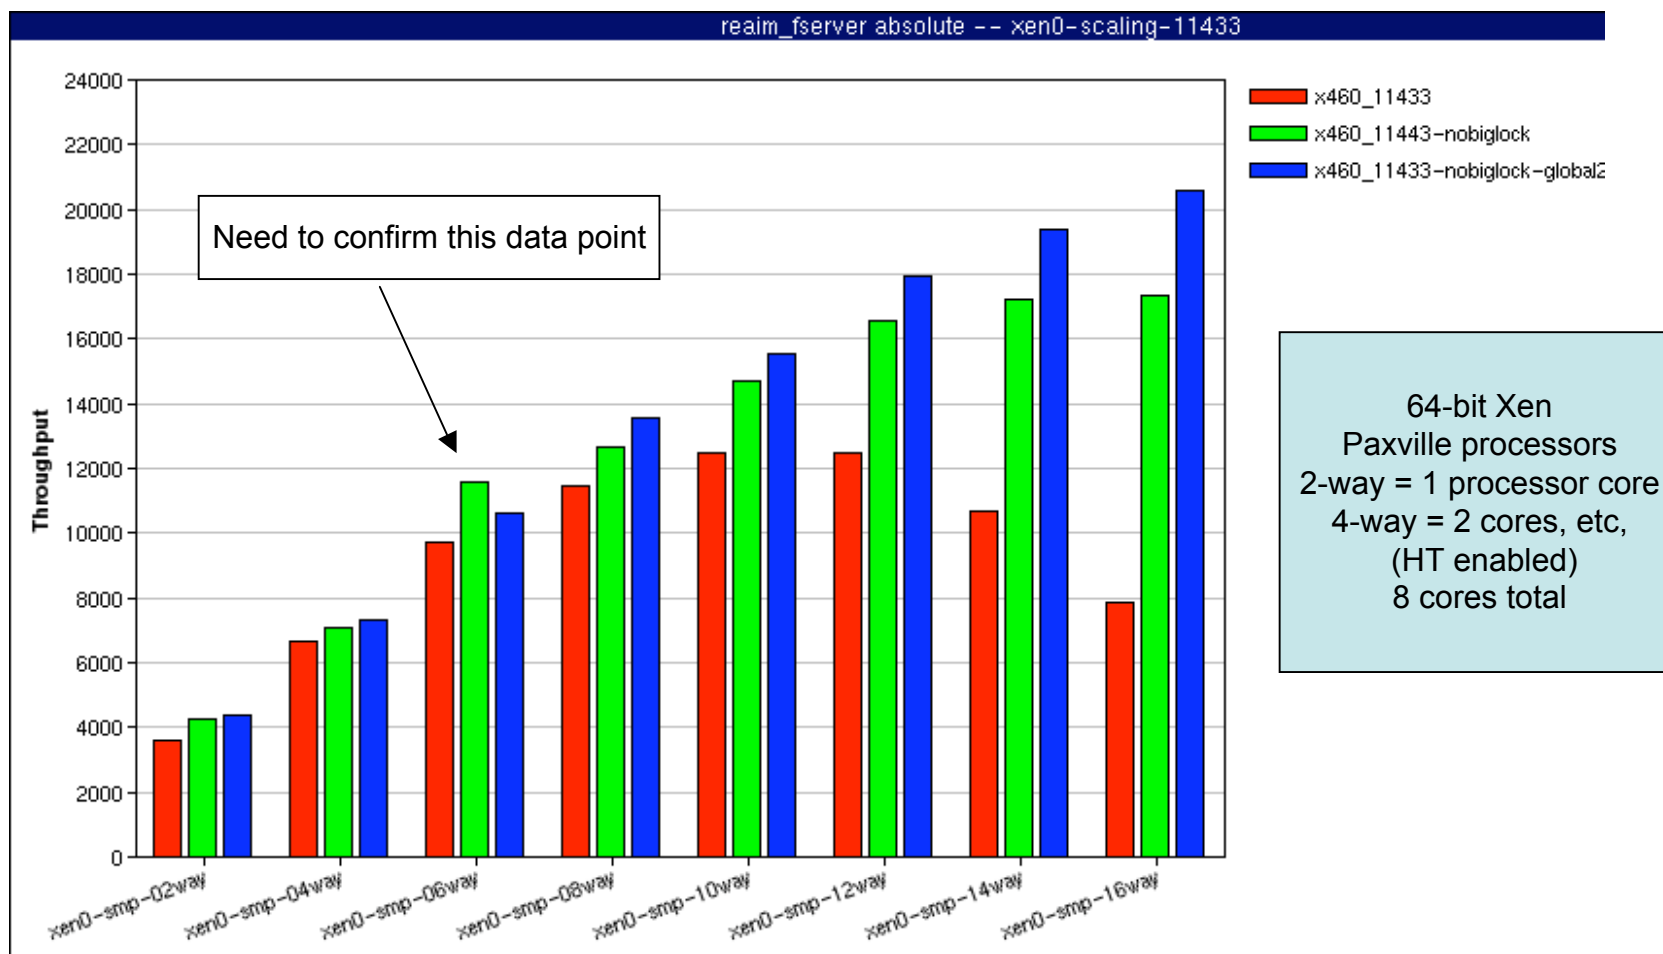
# Scalability

- Scaling many domains
  - One domain per processor core, 1 to many cores/domains
  - Presented data showing problems at OLS on 64-bit domains
    - Since then, tested 32-bit domains, which look fine
    - 64-bit did not scale because of excessive TLB flushes, which put too much pressure on global update of `tlbflush_clock` in Xen. Jun has fixed with TLB global bit patch –no need to flush all TLB entries for syscalls, etc, which also does not update `tlbflush_clock`. Patch also improves single domain performance by keeping user TLB entries.




# Scalability

- Scaling a large domain, 1 to many vCPUs
  - Elimination of writable page-tables in xen-unstable helps a lot (red bars below) but still have significant scaling problems beyond 6-way
  - Now need fine grain locking in mm functions; should be easier to implement with writable page-tables gone. (green bars have domain big-lock removed)
  - TLB global bit optimization also helps scalability (blue bars)



# Full Virtualization

- Shadow2 code shows vast improvement in processor performance/scaling
  - Boosts single domain throughput significantly -Dbench (running in tmpfs filesystem, **NO I/O**)
  - Many HVM domain scalability also improved dramatically
    - 1 to 7 HVMs scaled 1 : 6.47 with shadow2. Same test scaled 1 : 2.63 with old shadow code



# Full Virtualization

- Para-virt drivers in HVM domain significantly better than IO emulation
  - Can achieve same throughput as domU (dom0 has dedicated CPU)
  - But still lots of CPU in dom0 (HW interrupt emulation?)
  - Different versions of drivers show different results
    - Driver rev 2 had performance equal to domU
    - Driver rev 9 took significant hit on networking
    - Unsure of performance of driver in xen-unstable, not “stable”
    - Need to track down source of inconsistency