

Network Bottlenecks

- Jeffrey Barnard
Barnard Software, Inc.

Mon 9:15

<http://www.bsiopti.com>

jeff@bsiopti.com

407-323-4773

Simple Network Tuning

- Pick a Bottleneck
 - You are at the mercy of the slowest link
- Latency is a big problem
- Where is the problem?
 - may not be where you think it is
- Application, VSE, stack, Network Interface
- Switch, subnet, router, firewall, Internet
- Firewall, Router, subnet, PC, application

▣ Latency Example

- FLEX-ES connected to iSeries
 - VSE/ESA 2.7.2
 - LCS Ethernet on FLEX-ES
 - Bandwidth limited by I/O interrupts
 - PCI Ethernet on iSeries
 - Dedicated subnet, switch, adapters
- DB2 Client running on VSE/ESA
- DB2 UDB running on iSeries

▣ Latency Example

- Found regular spikes in VSE CPU Usage
 - Partition using 50% CPU
- BSTTINET (stack) CPU fairly low
- Suspected LCS would be bottleneck
 - Trace showed application process problem
- Found DB2 client was bottleneck
 - Base DB2 client uses LE/C API
 - CPU intensive

▣ Latency Example

- Used Opti-Analyzer (O/A, O/W)
 - PK00122/PK11298 Helped
 - Converted from LE/C to ASM SOCKET
 - Reduced CPU
 - GETVIS/FREEVIS still a problem
 - DB2 Beta client much better
 - Reduced GETVIS/FREEVIS activity

□ Latency Example II

- traceroute to 10.7.31.22 (10.7.31.22), 30 hops max, 38 byte packets
 - 1 170.153.68.16 1.777 ms 0.273 ms 0.263 ms
 - 2 170.153.180.44 1.490 ms 0.793 ms 0.886 ms
 - 3 * * *
 - 4 10.7.28.2 43.267 ms 43.373 ms 46.830 ms
 - 5 10.7.28.1 109.692 ms 108.814 ms 110.780 ms
 - 6 10.7.31.22 124.069 ms 120.785 ms 129.594 ms
-
- Local subnet transfers are fast.
 - Problems end-to-end. *Anyone want to guess why?*

□ Network Interfaces

- OSA, OSA/2, LCS, CTC, IUCV, CLAW
 - SSCH to transfer data
 - Transfer rate limited by CPU speed
 - Interrupts per second is key factor
 - Full 100BaseTX requires 120MIPS
- CPU speed is the major factor in data transfer throughput

□ Network Interfaces

- OSA Express, Hipersockets
 - \$IJBOSA provided by IBM
 - QDIO (Queued Direct I/O)
 - Memory queuing model
 - Similar to PCI DMA
 - Dramatically less interrupt activity

□ Hipersocket Ponderings

- Hipersockets is a CPU function
- Check processor and processor level
- LPARs CPU limited?
- Synchronous data transfer
- z/VM effects performance (APARs)
- VSE/ESA 2.7 DY46197
- MTU size (1500 vs. ??????)
 - Device max is 56K
 - Optimal trade off (Bulk vs. Interactive)

□ VSE/ESA

- Get a FAST CPU
- Get a FAST I/O SubSystem
- Select an Access Method
 - VSAM (Fastest)
 - Library
 - SAM
 - POWER (Slowest)
- PRTY of VSE/ESA application(s)

The Remote Host

- Server PC (usually)
- Get a FAST CPU
 - AMD Opteron
 - Low memory latency
- Max out the memory
 - 2G or 4G
- FTP server
 - FTP Server running on 256M PC ... ?
- Plus what other software?

□ The Remote Host

- Maximum Segment Size
 - Windows likes 576, 1460 is better
- Receive Window Size
 - Windows default is 8K
 - 32K or 64K is much better
 - Must be exact multiple of MSS (1460)
 - 32120 or 64240

□ The Remote Hist

- MSS and RWS are typical problems
- Up to 6x improvement in transfer rate
 - MSS up to 3x
 - RWS up to 2x
 - $3x + 2x = 5x$
 - $3x * 2x = 6x$

■ Mainframe vs. PC

- Mainframes are apples
- PC's are oranges
- Are you trying to grow oranges on an apple tree?

▣ Network Bottlenecks

- Questions ?

Jeff Barnard

Barnard Software, Inc.

<http://www.bsiopti.com>

jeff@bsiopti.com

407-323-4773