

SPECweb99 Result

© Copyright 1999, Standard Performance Evaluation Corporation

IBM: IBM eServer pSeries 655
Zeus Technology Ltd.: Zeus 4.2r3

SPECweb99 = 10291

Test Date: May-2003
Tester: IBM

Hardware Avail: Jul-2003 OS Avail: Jul-2003 HTTP Software Avail: Jul-2003 Sup. Software Avail: Jul-2003 SPEC license #: 11

Hardware

Vendor: IBM
Model: IBM eServer pSeries 655
Processor: 1700MHz POWER4+
Processors: 4 cores, 4 chips, 1 core/chip, 4 chips/MCM
Primary Cache: 64KBI+32KBD (on chip) per core
Secondary Cache: 1536KB(I+D) unified (on chip) per chip
Other Cache: 128MB unified (off chip)/MCM, 4 chips/MCM
Memory: 64GB (4x16GB)
Disk Subsystem: 2x36GB (10 KRPM) 16Bit SCSI Disk
20x18.2GB (10 KRPM) SSA Disks
(fast-write enabled)
Disk Controllers: Wide/Ultra-2 SCSI I/O Controller
2 IBM SSA 160 SerialRAID Adapters
Other Hardware: 9 x Alteon ACESwitch 180

Software

Operating System: AIX 5L V5.2(64bit)
File System: JFS2
Other Software: None

HTTP Software

Vendor: Zeus Technology Ltd.
HTTP Software: Zeus 4.2r3
API: ISAPI
Server Cache: FRCA (Fast Response Cache Accelerator)
Log Mode: Binary CLF for Zeus, CLF for FRCA

Test Sponsor

Test Date: May-2003
Tested By: IBM
SPEC License: 11

Network

of Controllers: 8
Network Controllers: 8 IBM Gigabit Ethernet-SX PCI-X Adapter(14106802)
of Nets: 8
Type of Nets: Gigabit Ethernet
Network Speed: 1 Gb/sec
MSL (sec): 30 (Non RFC1122)
Time-Wait (sec): 60 (Non RFC1122)
MTU: 1500

Clients

of Clients: 40
Model: IBM pSeries 630
Processor: 1000 MHz POWER4
of Processors: 4
Memory: 8GB
Network Controller: 1 Gigabit Ethernet-SX PCI-X Adapter (14106802)
Operating System: AIX 5.2
Compiler: IBM Cset++ Version 5.1

Notes/Tuning Information

SUT Notes

APAR IY44268 (misc AIX 5.2 code enhancements, include tcpip and Gb ethernet driver performance) was applied to AIX 5L V5.2
bindintcpu [interrupt level] [cpu] (Interrupt affinity - bind 8 Ethernet interrupts to 4 cpu's(2 interrupts to 1 cpu))
maxpgahead=64(the maximum number of pages to be read ahead,default=8)
minperm=5%(the point below which the page-stealer will steal file of computational pages,default=20%)
maxperm=95%(the point above which the page stealing algorithm steals only file page,default=80%)
4 different logs for 4 zeus instances (1 log per zeus instance) for dynamic requests
frca-log.N (once the frca-log.N reaches 1GB, a new frca-log.N+1 is generated) for static requests
Disk layout: 1 disk for OS, 1 disk for webserver, 14 disks striped for file_set, 6 disks striped for web logfile)
Set Gigabit Ethernet driver attribute rxbuf_pool_sz=2048 (set Receive buffer pool size, default=1024)
Set Gigabit Ethernet driver attribute txdesc_queue_sz=1024 (set Transmit descriptor queue size, default=512)
Set Gigabit Ethernet driver attribute copy_bytes=256 (copy packet into contiguous buffer on transmit if this many or less bytes, default=2048)
Set Gigabit Ethernet driver attribute intr_rate=2000 (maximum number of receive interrupts per second that a Gb ethernet adapter will generate, default=0 (no limit))
fractrl pconintr 100 (frca set to allow 100% cpu time on interrupt, default=90)
fractrl revaltimeout 100000 (set frca revalidation timeout to 100,000 seconds, default=300 seconds)
no(network option) tuning:
no -o tcp_sendspace=65536(default socket send buffer size, default=16384)
no -o tcp_recvspace=65536(default socket receive buffer size, default=16384)
no -o send_file duration=100000(cached send_file expiration time, default=300)
no -o somaxconn=10240(max listen backlog, default=1024)
no -o nbc_max_cache=409600(max file size for network buffer cache, default=1048576)
no -o nbc_pseg=800000(max file entries in private segment,default=0)
no -o tcp_timewait=5(ensure that TIME_WAIT is at least 60 seconds, default=1)
no -o tcp_newreno=0 (disables the modification to TCP's Fast Recovery algorithm as described in RFC 2582, default=1)
no -o tcp_inpcb_hashtab_sz=128021 (increase TCP PCB hash table size, default=24499)
One external 7040-61D I/O drawer used for 8 Gigabit ethernet adapters

SPECweb99 Result

© Copyright 1999, Standard Performance Evaluation Corporation

Notes/Tuning Information (Continued)

SUT Notes (Continued)

Two external 7133-D40 SSA drawers used for 20 SSA disks
MCM stands for Multiple Chip Module

HTTP API Notes

Zeus PEPP-06 SPECweb99 toolkit used for dynamic content
post.log in /weblog directory instead of doc_root
./Configure --sendfile=no --locking=semop
Compiled with IBM VAC 6.0.0
CFLAGS=-q64 -qlargetype -O2 -qmaxmem=8192
LDLFLAGS=-q64 -qlargetype -bnoentry -Wl -G -bexpall -bM:SRE -lc -lbsd

Client Notes

no (network option) tuning:
no -o delayack=3(delay ack for connection setup and shutdown,default=0)
no -o delayackports={80}(delay ack ports,default={ })
no -o tcp_timewait=5(ensure that TIME_WAIT is at least 60 seconds,default=1)
tcp_sendspace=65536(default socket send buffer size, default=16384)
tcp_recvspace=65536(default socket receive buffer size, default=16384)
AIX (thread-scheduling option) tuning:
AIXTHREAD_SCOPE=S (AIX pthread environment variable, set all threads have system based contention scope, {1:1}, default=P{8:1})
AIXTHREAD_SLPRATIO=1:1 (AIX pthread environment variable, controls the number of kernel threads that should be held in reserve for sleeping threads, default={12:1})
YIELDLOOPTIME=1000 (AIX pthread environment variable, controls the number of times to yield the processor before blocking on a busy lock, default=0)

Other Notes

Starts 4 unique instances of Zeus sharing a common docroot
Each Zeus instance bound to processor corresponding to its NIC
Zeus 4.2r3 global.cfg (one per instance)
tuning!bind_any no
tuning!cache_files 50021
tuning!cache_large_file 9216
tuning!cache_small_file 4096
tuning!cache_stat_expire 180000
tuning!cache_flush_interval 180000
tuning!cache_max_bytes 0
tuning!num_children 2
tuning!keepalive yes
tuning!keepalive_timeout 1000
tuning!keepalive_max -1
tuning!listen_queue_size 9216
tuning!cbuff_size 1048576
tuning!multiple_accept yes
tuning!sendfile yes
tuning!sendfile_maxsize 1048576
tuning!sendfile_minsize 100
tuning!sendfile_reservefd 31993
tuning!so_nagle_off no
tuning!so_rbuff_size 32768
tuning!softservers no
tuning!unique_bind yes
tuning!use_poll no
tuning!cache_cooling_time 0
tuning!modules!cgi!cleansize 0
tuning!modules!stats!enabled no
tuning!modules!nsapi!enabled no
tuning!accelerator!frca!enabled yes
tuning!accelerator!frca!log!enabled yes
tuning!accelerator!frca!log!filename /weblog/log/frca-log
tuning!accelerator!frca!log!format CLF

SPECweb99 Result

© Copyright 1999, Standard Performance Evaluation Corporation

Notes/Tuning Information (Continued)

Other Notes (Continued)

tuning!accelerator!frca!maxsize 1048576
tuning!accelerator!frca!minsize 100
Zeus_Site_Configuration (one per instance):
modules!cgi!enabled yes
modules!isapi!enabled yes
bindaddr ip_addr1, ip_addr2
Tuning Disclosure: IBM-20011019.txt
Dynamic API: IBM-20011019_2-API.tar.gz