

SPECweb99 Result

© Copyright 1999, Standard Performance Evaluation Corporation

IBM: IBM eServer pSeries 690 HPC
Zeus Technology Ltd.: Zeus 4.0

SPECweb99 = 21000

Test Date: Oct-2001
Tester: IBM

Hardware Avail: Dec-2001 OS Avail: Jan-2002 HTTP Software Avail: Oct-2001 Sup. Software Avail: Dec-2001 SPEC license #: 11

Hardware

Vendor: IBM
Model: IBM eServer pSeries 690 HPC
Processor: 1.3 GHz POWER-4
Processors: 16 cores, 8 chips, 2 cores/chip, 4 chips/MCM
Primary Cache: 64KBI+32KBD (on chip) per core
Secondary Cache: 1440KB unified (on chip) per chip
Other Cache: 32MB * 4 unified off chip
Memory: 192 GB
Disk Subsystem: 2x18.2GB 16Bit SCSI Disk
23x18.2GB SSA Disk(fast-write enabled)
Disk Controllers: Wide/Ultra-2 SCSI I/O Controller,
2 x IBM SSA 160 SerialRAID Controller
Other Hardware: 16xAlteon ACESwitch 180

Software

Operating System: AIX 5L V5.1 (64-bit)
File System: JFS and JFS2
Other Software: None

HTTP Software

Vendor: Zeus Technology Ltd.
HTTP Software: Zeus 4.0
API: ISAPI
Server Cache: FRCA (Fast Response Cache Accelerator)
Log Mode: Binary CLF for Zeus, CLF for FRCA

Test Sponsor

Test Date: Oct-2001
Tested By: IBM
SPEC License: 11

Network

of Controllers: 32
Network Controllers: 32 IBM Gigabit Ethernet-SX PCI Adapter
of Nets: 32
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: 96
Model: RS/6000 44P-270
Processor: 375 MHz POWER3-II
of Processors: 4
Memory: 2048MB
Network Controller: 1 IBM Gigabit Ethernet-SX PCI Adapter
Operating System: AIX 4.3.3
Compiler: IBM Cset++ Version 3.1.3

Notes/Tuning Information

SUT Notes

APAR IY 24907 was applied to AIX 5L V5.1 to enhance TCP/IP performance
no(network option) tuning:
no -o tcp_sendspace=131072(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)
bindintcpu [interrupt level] [cpu] (Interrupt affinity - bind 32 Ethernet interrupts to 16 cpu's(2 interrupt to 1 cpu))
bindprocessor [process] [cpu] (Process affinity - bind 32 Zeus child processes to 16 cpu's(2 processes 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%)
16 different logs for 16 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, 20 disks striped for file_set, 3 disks striped for web logfile)
Zeus doc_root using JFS2; OS, webserver, weblog using JFS
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 large_send=yes (turn on large_send capability of TCP/IP, default=no)
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)
4 external 7040-61D I/O drawers (8 adapters each).

SPECweb99 Result

© Copyright 1999, Standard Performance Evaluation Corporation

Notes/Tuning Information (Continued)

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

Client Notes

APAR IY17356 was applied to AIX 4.3.3
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)

Other Notes

Starts 16 unique instances of Zeus sharing a common docroot
The 2 child iprocesses of each Zeus instance bind to two unique IP address
Zeus 4.0 global.cfg (one per instance)
tuning!bind_any no
tuning!cache_files 74017
tuning!cache_large_file 5120
tuning!cache_small_file 5120
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 10240
tuning!cbuff_size 65536
tuning!timeout 600
tuning!multiple_accept yes
tuning!sendfile yes
tuning!sendfile_maxsize 1048576
tuning!sendfile_minsize 5120
tuning!sendfile_reservefd 31991
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
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