

STREAM Memory Bandwidth result Submission on IBM Power 8 server

Date result produced : Thu Oct 22 2015
Questions? : Gnanakumar.Rajaram@Oracle.com

System Configuration :

Server Name	IBM S824
CPU	Power 8 3.52 GHz, 4 chips, 24 cores, 96 threads in SMT4
Memory	512 GB
OS	AIX 7.1
OS patch applied	AIX 7.1 TL 3 SP3 # oslevel -s 7100-03-03-1415
Compiler	XL C/C++ 13.1 for AIX 7.1
System SMT Enabled	SMT4
Large Pages	4000 x 16MB Pages configured
STREAM Benchmark Array size used	1000000276
Compiler flags used to build STREAM	xlc_r stream_5-10_posix_memalign.c -o stream-p8.lpg-5-10posix.exe -DSTREAM_ARRAY_SIZE=1000000276 -DNTIMES=40 -O5 - qsimd -qsmp=omp -qstrict -qlargepage -qthreaded -blpdata -q64
Disclaimer	Oracle made reasonable efforts to ensure the accuracy and validity of its testing and suggest that knowledgeable database and software engineers run it themselves to validate these results.

Environment Variable:

```
XLSMPOPTS=startproc=0:stride=4  
LDR_CNTRL=MAXDATA=0xD0000000@LARGE_PAGE_DATA=M  
OMP_NUM_THREADS=24  
MEMORY_AFFINITY=MCM
```

STREAM Benchmark

<https://www.cs.virginia.edu/stream>

STREAM Source Code (OMP):

```
/* Revision: $Id: stream.c,v 5.10.1 2014/06/17 08:16:08 mccalpin Exp mcalpin $ */  
https://www.cs.virginia.edu/stream/FTP/Code/Versions/stream\_5-10\_posix\_memalign.c
```

The source code change in the file stream_5-10_posix_memalign.c to do the reverse allocation:

```
    for (j=STREAM_ARRAY_SIZE-1; j>=0; j--) {  
        a[j] = 1.0;  
        b[j] = 2.0;  
        c[j] = 0.0;  
    }  
---  
    for (j=0; j<STREAM_ARRAY_SIZE; j++) {  
        a[j] = 1.0;  
        b[j] = 2.0;  
        c[j] = 0.0;  
    }
```

Results:

Maximum STREAM Benchmark Performance:

```
SMT 4, 24 OMP Threads used with XLSMPOPTS=startproc=0:stride=4  
-----  
STREAM version $Revision: 5.10 $
```

This system uses 8 bytes per array element.

Array size = 100000276 (elements), Offset = 0 (elements)
Memory per array = 7629.4 MiB (= 7.5 GiB).
Total memory required = 22888.2 MiB (= 22.4 GiB).
Each kernel will be executed 40 times.
The *best* time for each kernel (excluding the first iteration)
will be used to compute the reported bandwidth.

Number of Threads requested = 24
Number of Threads counted = 24

Your clock granularity/precision appears to be 1 microseconds.
Each test below will take on the order of 63173 microseconds.
(= 63173 clock ticks)

Increase the size of the arrays if this shows that
you are not getting at least 20 clock ticks per test.

WARNING -- The above is only a rough guideline.
For best results, please be sure you know the
precision of your system timer.

Function	Best Rate MB/s	Avg time	Min time	Max time
Copy:	253710.1906	0.0634	0.0631	0.0677
Scale:	253000.4770	0.0642	0.0632	0.0913
Add:	319276.2247	0.0755	0.0752	0.0768
Triad:	326176.6201	0.0765	0.0736	0.1649

Solution Validates: avg error less than 1.000000e-13 on all three arrays

Bisection Bandwidth Benchmark Performance (Nonlocal STREAM):

SMT 4, 24 OMP Threads used with XLSMPOPTS=startproc=0:stride=4

STREAM version \$Revision: 5.10 \$

This system uses 8 bytes per array element.

Array size = 100000276 (elements), Offset = 0 (elements)
Memory per array = 7629.4 MiB (= 7.5 GiB).
Total memory required = 22888.2 MiB (= 22.4 GiB).
Each kernel will be executed 40 times.
The *best* time for each kernel (excluding the first iteration)
will be used to compute the reported bandwidth.

Number of Threads requested = 24
Number of Threads counted = 24

Your clock granularity/precision appears to be 1 microseconds.
Each test below will take on the order of 228363 microseconds.
(= 228363 clock ticks)

Increase the size of the arrays if this shows that
you are not getting at least 20 clock ticks per test.

WARNING -- The above is only a rough guideline.
For best results, please be sure you know the
precision of your system timer.

Function	Best Rate MB/s	Avg time	Min time	Max time
Copy:	54972.6300	0.2925	0.2911	0.3073
Scale:	53914.3659	0.2995	0.2968	0.3451
Add:	60765.8219	0.3983	0.3950	0.4708
Triad:	61687.9510	0.3901	0.3891	0.3929

Solution Validates: avg error less than 1.000000e-13 on all three arrays

Large Pages Configured:

```
vmo -o lgpg_regions=4000 -o lgpg_size=16777216  
vmo -o v_pinshm=1
```

This shows NUMA Configuraion:

```
# lssrad -av
```

```
REF1  SRAD      MEM      CPU
```

```

0
    0 107862.44 0-23
    1 107938.00 24-47
1
    2 105985.00 48-71
    3 107936.44 72-95

```

This shows the available pages of different Page Sizes:

```

# svmon -G
memory      size      inuse      free      pin      virtual  mmode
pg space    131072    7098
work        pers      clnt      other
pin         977199   0         962      3761696
in use     5158593  0         217843

PageSize   PoolSize   inuse      pgspace   pin      virtual
s  4 KB     -         4160420   7098     3829585  3942577
m  64 KB    -         75745     0        56892   75745
L  16 MB    4000     1         0        4000    1
S  16 GB    -         0         0        0        0

```

System Configuration

```

# prtconf

System Model: IBM,8286-42A
Machine Serial Number: Not Available
Processor Type: Not Available
Processor Implementation Mode: POWER 8
Processor Version: PV_8_Compact
Number Of Processors: 0
Processor Clock Speed: 3525 MHz
CPU Type: 64-bit
Kernel Type: 64-bit
LPAR Info: 1 hal-p8 whole
Memory Size: 508416 MB
Good Memory Size: Not Available
Platform Firmware level: FW810.02
FW810.11
firmware
Firmware Version: IBM,SV810_087
Console Login: enable
Auto Restart: true
Full Core: false
NX Crypto Acceleration: Capable and Enabled

```

```

Network Information
Host Name:
IP Address:
Sub Netmask:
Gateway:
Name Server:
Domain Name:

```

Paging Space Information

Volume Groups Information

```

=====
Active VGs
=====
datavg:
PV_NAME      PV STATE      TOTAL PPs   FREE PPs   FREE DISTRIBUTION
hdisk1      active        528         50         00..00..00..00..50
hdisk2      active        528         527        106..105..105..105..106
=====

rootvg:
PV_NAME      PV STATE      TOTAL PPs   FREE PPs   FREE DISTRIBUTION
hdisk3      active        528         225        96..00..00..23..106
=====

```

INSTALLED RESOURCE LIST

The following resources are installed on the machine.
+/- = Added or deleted from Resource List.
* = Diagnostic support not available.

```

Model Architecture: chrp
Model Implementation: Multiple Processor, PCI bus

```

+ sys0		System Object
+ sysplanar0		System Planar
* vio0		Virtual I/O Bus
* vsa0	U8286.42A.216B52V-V1-C0	LPAR Virtual Serial Adapter
* vty0	U8286.42A.216B52V-V1-C0-L0	Asynchronous Terminal
* pci13	U78C9.001.WZS01J1-P1	PCI Express Bus
* pci12	U78C9.001.WZS01J1-P1	PCI Express Bus
+ ent8	U78C9.001.WZS01J1-P1-C2-T1	PCIe2 2-Port 10GbE Base-T Adapter (e4148e1614109204)
+ ent9	U78C9.001.WZS01J1-P1-C2-T2	PCIe2 2-Port 10GbE Base-T Adapter (e4148e1614109204)
* pci11	U78C9.001.WZS01J1-P1	PCI Express Bus
* pci10	U78C9.001.WZS01J1-P1	PCI Express Bus
* pci9	U78C9.001.WZS01J1-P1	PCI Express Bus
+ sissas1	U78C9.001.WZS01J1-P1-C15-T1	PCIe3 x8 SAS RAID Internal Adapter 6Gb
* sas1	U78C9.001.WZS01J1-P1-C15-T1	Controller SAS Protocol
* sfwcomm3		SAS Storage Framework Comm
+ ses1	U78C9.001.WZS01J1-P2-Y1	SAS Enclosure Services Device
+ pdisk2	U78C9.001.WZS01J1-P2-D1	Physical SAS Disk Drive
+ pdisk3	U78C9.001.WZS01J1-P2-D2	Physical SAS Disk Drive
* hdisk2	U78C9.001.WZS01J1-P1-C15-T1-L205DA5C600-L0	SAS RAID 0 Disk Array
* hdisk3	U78C9.001.WZS01J1-P1-C15-T1-L405DA5C600-L0	SAS RAID 0 Disk Array
* sata1	U78C9.001.WZS01J1-P1-C15-T1	Controller SATA Protocol
* pci8	U78C9.001.WZS01J1-P1	PCI Express Bus
+ ent4	U78C9.001.WZS01J1-P1-C10-T1	4-Port Gigabit Ethernet PCI-Express Adapter (e414571614102004)
+ ent5	U78C9.001.WZS01J1-P1-C10-T2	4-Port Gigabit Ethernet PCI-Express Adapter (e414571614102004)
+ ent6	U78C9.001.WZS01J1-P1-C10-T3	4-Port Gigabit Ethernet PCI-Express Adapter (e414571614102004)
+ ent7	U78C9.001.WZS01J1-P1-C10-T4	4-Port Gigabit Ethernet PCI-Express Adapter (e414571614102004)
* pci7	U78C9.001.WZS01J1-P1	PCI Express Bus
* pci6	U78C9.001.WZS01J1-P1	PCI Express Bus
* pci5	U78C9.001.WZS01J1-P1	PCI Express Bus
+ usbhc0	U78C9.001.WZS01J1-P1	Integrated USB 3.0 xHCI Adapter (4c1041821410b204)
* pci4	U78C9.001.WZS01J1-P1	PCI Express Bus
+ fcs0	U78C9.001.WZS01J1-P1-C6-T1	PCIe2 2-Port 16Gb FC Adapter (df1000e21410f103)
* fscsi0	U78C9.001.WZS01J1-P1-C6-T1	FC SCSI I/O Controller Protocol Device
* hdisk28	U78C9.001.WZS01J1-P1-C6-T1-W210000E1E1D73E0-L0	Other FC SCSI Disk Drive
* hdisk29	U78C9.001.WZS01J1-P1-C6-T1-W210000E1E1D73E1-L0	Other FC SCSI Disk Drive
* hdisk5	U78C9.001.WZS01J1-P1-C6-T1-W210000E1E1D73E0-L1000000000000	Other FC SCSI Disk Drive
* sfwcomm0	U78C9.001.WZS01J1-P1-C6-T1-W0-L0	Fibre Channel Storage Framework Comm
+ fcs1	U78C9.001.WZS01J1-P1-C6-T2	PCIe2 2-Port 16Gb FC Adapter (df1000e21410f103)
* fscsi1	U78C9.001.WZS01J1-P1-C6-T2	FC SCSI I/O Controller Protocol Device
* hdisk17	U78C9.001.WZS01J1-P1-C6-T2-W210000E1E1DA121-L1000000000000	Other FC SCSI Disk Drive
* hdisk22	U78C9.001.WZS01J1-P1-C6-T2-W210000E1E1DA120-L0	Other FC SCSI Disk Drive
* hdisk23	U78C9.001.WZS01J1-P1-C6-T2-W210000E1E1DA120-L1000000000000	Other FC SCSI Disk Drive
* hdisk4	U78C9.001.WZS01J1-P1-C6-T2-W210000E1E1DA121-L0	Other FC SCSI Disk Drive
* sfwcomm1	U78C9.001.WZS01J1-P1-C6-T2-W0-L0	Fibre Channel Storage Framework Comm
* pci3	U78C9.001.WZS01J1-P1	PCI Express Bus
+ sissas0	U78C9.001.WZS01J1-P1-C14-T1	PCIe3 x8 SAS RAID Internal Adapter 6Gb
* sas0	U78C9.001.WZS01J1-P1-C14-T1	Controller SAS Protocol
* sfwcomm2		SAS Storage Framework Comm
+ ses0	U78C9.001.WZS01J1-P2-Y1	SAS Enclosure Services Device
+ pdisk0	U78C9.001.WZS01J1-P2-D7	Physical SAS Disk Drive
+ pdisk1	U78C9.001.WZS01J1-P2-D8	Physical SAS Disk Drive
* hdisk0	U78C9.001.WZS01J1-P1-C14-T1-L205DA5E100-L0	SAS RAID 0 Disk Array
* hdisk1	U78C9.001.WZS01J1-P1-C14-T1-L405DA5E100-L0	SAS RAID 0 Disk Array
* sata0	U78C9.001.WZS01J1-P1-C14-T1	Controller SATA Protocol
+ cd0	U78C9.001.WZS01J1-P2-D27	SATA DVD-RAM Drive
* pci2	U78C9.001.WZS01J1-P1	PCI Express Bus
+ ent0	U78C9.001.WZS01J1-P1-C12-T1	4-Port Gigabit Ethernet PCI-Express Adapter (e414571614102004)
+ ent1	U78C9.001.WZS01J1-P1-C12-T2	4-Port Gigabit Ethernet PCI-Express Adapter (e414571614102004)
+ ent2	U78C9.001.WZS01J1-P1-C12-T3	4-Port Gigabit Ethernet PCI-Express Adapter (e414571614102004)
+ ent3	U78C9.001.WZS01J1-P1-C12-T4	4-Port Gigabit Ethernet PCI-Express Adapter (e414571614102004)
* pci1	U78C9.001.WZS01J1-P1	PCI Express Bus
* pci0	U78C9.001.WZS01J1-P1	PCI Express Bus
+ L2cache0		L2 Cache
+ mem0		Memory
+ proc0		Processor
+ proc8		Processor
+ proc104		Processor
+ proc112		Processor
+ proc120		Processor
+ proc128		Processor
+ proc136		Processor
+ proc144		Processor
+ proc152		Processor
+ proc16		Processor
+ proc160		Processor
+ proc168		Processor
+ proc176		Processor
+ proc184		Processor
+ proc24		Processor
+ proc32		Processor
+ proc40		Processor
+ proc48		Processor
+ proc56		Processor
+ proc64		Processor
+ proc72		Processor
+ proc80		Processor
+ proc88		Processor
+ proc96		Processor

oslevel -s

7100-03-03-1415

uname -a

AIX hal-p8 1 7 00F96B524C00

ulimit -a

```
# ulimit -a
time(seconds)      unlimited
file(blocks)       unlimited
data(kbytes)       unlimited
stack(kbytes)      32768
memory(kbytes)     32768
coredump(blocks)   2097151
nofiles(descriptors) 2000
threads(per process) unlimited
processes(per user) unlimited
```

df -k

Filesystem	1024-blocks	Free	%Used	Iused	%Iused	Mounted on
/dev/hd4	125829120	125308184	1%	10502	1%	/
/dev/hd2	4194304	1012636	76%	56059	20%	/usr
/dev/hd9var	4194304	3766800	11%	6440	1%	/var
/dev/hd3	4194304	2580060	39%	3454	1%	/tmp
/dev/fwdump	2097152	2096504	1%	4	1%	/var/adm/ras/platform
/dev/hd1	524288	350128	34%	1234	2%	/export/home
/dev/hd11admin	524288	523848	1%	5	1%	/admin
/proc	-	-	-	-	-	/proc
/dev/hd10opt	2097152	926736	56%	12948	6%	/opt
/dev/livedump	524288	523880	1%	4	1%	/var/adm/ras/livedump
/dev/fslv00	4194304	4144696	2%	9	1%	/data
/dev/fslv01	250085376	126697240	50%	208814	1%	/bigdata
/dev/fslv03	524288	484948	8%	22	1%	/tftpboot
/dev/fslv02	2097152	231228	89%	616	2%	/export/lpp_source
/dev/fslv04	2097152	337876	84%	42414	36%	/export/spot

vmo -a

```
ame_cpus_per_pool = n/a
ame_maxfree_mem = n/a
ame_min_ucpool_size = n/a
ame_minfree_mem = n/a
ams_loan_policy = n/a
enhanced_affinity_affin_time = 1
enhanced_affinity_vmpool_limit = 10
esid_allocator = 1
force_relalias_lite = 0
kernel_heap_psize = 65536
lgpg_regions = 4000
lgpg_size = 16777216
low_ps_handling = 1
maxfree = 1088
maxperm = 98986661
maxpin = 102769629
maxpin% = 90
memory_frames = 130154496
memplace_data = 0
memplace_mapped_file = 0
memplace_shm_anonymous = 0
memplace_shm_named = 0
memplace_stack = 0
memplace_text = 0
memplace_unmapped_file = 0
minfree = 960
minperm = 3299552
minperm% = 3
msem_nlocks = 0
nokilluid = 0
```

```
    npskill = 1024
    npswarn = 4096
num_locks_per_semId = 1
    numpsblks = 131072
    pinnable_frames = 109027143
relalias_percentage = 0
    scrub = 0
    thrpGIO_inval = 1024
    thrpGIO_npages = 1024
    v_pinshM = 1
    vm_mmap_bmap = 1
vmm_default_pspa = 0
    vmm_klock_mode = 2
wlm_memlimit_nonpg = 1
```