VT-d PI Performance on KVM

Oct. 2014

Feng Wu

Environment

- Hardware
  - BDW
- Guest Information
  - 4 vCPUs
  - 1G Memory
  - RHEL 6.5
  - Directly-assigned 10G NIC
  - Workload: kernel build (make –j4)
- Tools
  - Netperf
  - Perf kvm

Netperf TCP_STREAM

- Purpose:
- Test the network throughput of TCP stream traffic.

Netperf TCP_STREAM

- w/o VT-d PI:
- Recv Send Send
- Socket Socket Message Elapsed
- Size Size Size Time Throughput (tried 5 times) Average throughput
- bytes bytes bytes secs. 10^6bits/sec 10^6bits/sec
- 87380 16384 4 60.00 95.33/95.33/87.96/94.72/95.22 93.71
- 87380 16384 8 60.00 212.55/209.16/198.93/198.57/198.47 203.54
- 87380 16384 16 60.00 418.72/398.56/409.44/396.37/402.43 405.10
- 87380 16384 32 60.00 823.11/777.13/752.84/746.71/745.37 769.03
- 87380 16384 64 60.00 1560.50/1506.64/1503.87/1518.66/1531.72 1524.28
- 87380 16384 128 60.00 3126.46/3052.63/3047.48/3035.66/3032.85 3059.02
- 87380 16384 256 60.00 5836.09/5862.76/5796.52/5842.19/5810.34 5829.58
- 87380 16384 512 60.00 7613.03/7044.06/7179.34/7422.97/7312.17 7314.31
- 87380 16384 1024 60.01 7253.58/7453.53/7584.26/7220.44/7120.34 7326.43
- 87380 16384 2048 60.00 7388.51/7577.36/7143.30/7266.59/7079.94 7291.14
- 87380 16384 4096 60.00 7440.51/7536.63/7363.35/7410.44/7257.41 7401.67

Netperf TCP_STREAM

- w/ VT-d PI:
- Recv Send Send
- Socket Socket Message Elapsed
- Size Size Size Time Throughput (tried 5 times) Average throughput
- bytes bytes bytes secs. 10^6bits/sec 10^6bits/sec
- 87380 16384 4 60.00 87.92/98.31/92.39/90.34/92.95 92.38
- 87380 16384 8 60.00 219.35/195.44/196.56/194.84/194.47 200.13
- 87380 16384 16 60.00 419.77/395.25/397.34/404.77/406.06 404.64
- 87380 16384 32 60.00 817.77/814.08/818.27/815.00/806.06 814.24
- 87380 16384 64 60.00 1644.87/1608.99/1654.31/1607.44/1628.89 1628.9
- 87380 16384 128 60.00 3151.70/3195.09/2998.46/3032.76/3172.56 3110.11
- 87380 16384 256 60.00 5892.61/5685.83/5659.95/5680.91/5673.24 5718.51
- 87380 16384 512 60.00 7765.93/7901.10/7690.43/8065.14/7837.60 7852.04
- 87380 16384 1024 60.00 7591.45/7932.04/7723.22/7781.32/7844.43 7774.49
- 87380 16384 2048 60.01 8004.03/7928.92/7848.58/7860.87/7750.78 7878.64
- 87380 16384 4096 60.00 7742.98/7793.85/7613.98/7802.63/7729.49 7736.59

Netperf TCP_STREAM

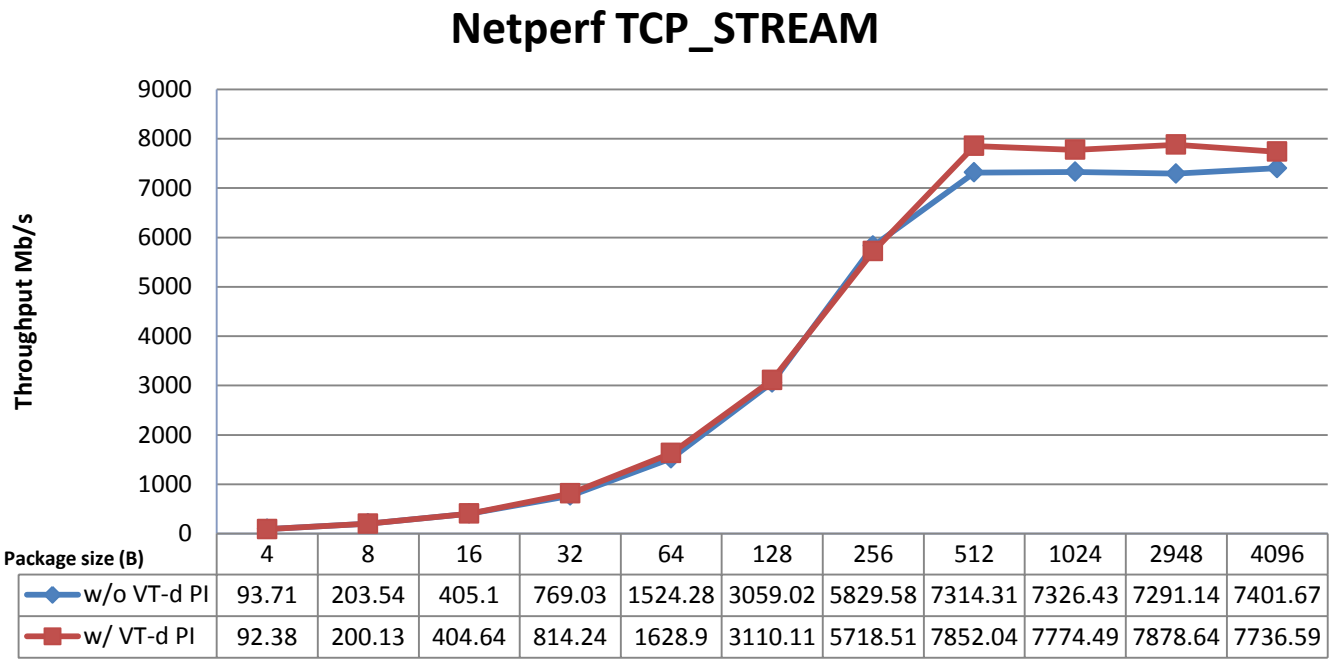

Package size (B)	4	8	16	32	64	128	256	512	1024	2948	4096
w/o VT-d PI	93.71	203.54	405.1	769.03	1524.28	3059.02	5829.58	7314.31	7326.43	7291.14	7401.67
w/ VT-d PI	92.38	200.13	404.64	814.24	1628.9	3110.11	5718.51	7852.04	7774.49	7878.64	7736.59

Netperf TCP_RR

- Purpose:
- Test the TCP Request/Response rate. It can reflex the round trip network latency.

Netperf TCP_RR

- w/o VT-d PI:
- Recv Send Send
- Socket Socket Message Elapsed
- Socket Size Request Resp. Elapsed Trans. (Tried 5 times) Average throughput
- Send Recv Size Size Time Rate 10^6bits/sec
- bytes Bytes bytes bytes secs. per sec
- 16384 87380 4 4 60.00 23445.65/23508.45/18911.26/23941.89/20745.19 22110.49
- 16384 87380 8 8 60.00 14651.09/23901.10/17873.24/22515.74/23297.59 20447.75
- 16384 87380 16 16 60.00 18216.59/22500.58/22426.60/20078.41/22581.86 21160.81
- 16384 87380 32 32 60.00 23896.26/24051.48/18178.72/23633.17/19168.76 21785.69
- 16384 87380 64 64 60.00 19170.48/19393.69/18401.55/18033.50/18143.50 18628.54
- 16384 87380 128 128 60.00 13764.54/18951.09/18919.43/14593.40/17654.23 16776.54
- 16384 87380 256 256 60.00 18228.27/11779.83/11574.21/13432.31/14212.34 13845.39
- 16384 87380 512 512 60.00 13316.00/15216.93/13763.91/15049.56/13452.21 14159.72
- 16384 87380 1024 1024 60.00 12646.65/13033.35/11094.06/8485.43/12378.34 11527.57
- 16384 87380 2048 2048 60.00 5169.25/5886.51/5415.23/4616.82/6108.57 5439.28
- 16384 87380 4096 4096 60.00 3967.51/5039.20/4334.42/3364.83/4263.54 4193.9

Netperf TCP_RR

- w/ VT-d PI:
- Recv Send Send
- Socket Socket Message Elapsed
- Socket Size Request Resp. Elapsed Trans. (Tried 5 times) Average throughput
- Send Recv Size Size Time Rate 10^6bits/sec
- bytes Bytes bytes bytes secs. per sec
- 16384 87380 4 4 60.00 28078.91/28404.39/28795.37/27183.40/27123.30 27917.07
- 16384 87380 8 8 60.00 28239.22/27836.33/28682.98/28008.17/28496.20 28252.58
- 16384 87380 16 16 60.00 28285.45/27864.88/28532.80/28194.81/28853.70 28346.33
- 16384 87380 32 32 60.00 27280.95/28014.28/28212.85/28486.82/17046.90 25808.36
- 16384 87380 64 64 60.00 16990.88/24136.35/24880.73/26704.97/17277.66 21998.12
- 16384 87380 128 128 60.00 19999.45/20687.68/20478.93/20388.73/20387.19 20388.70
- 16384 87380 256 256 60.00 24663.34/14811.94/13432.55/24986.10/16754.32 18929.65
- 16384 87380 512 512 60.00 16743.25/16221.31/16871.19/16851.59/16415.59 16620.59
- 16384 87380 1024 1024 60.00 9714.08/15221.48/14701.62/15142.79/15197.91 13995.58
- 16384 87380 2048 2048 60.00 5785.06/7358.04/6888.13/6155.04/5767.41 6390.74
- 16384 87380 4096 4096 60.00 5124.92/4435.09/5033.81/5121.93/5092.46 4961.64

Netperf TCP_STREAM

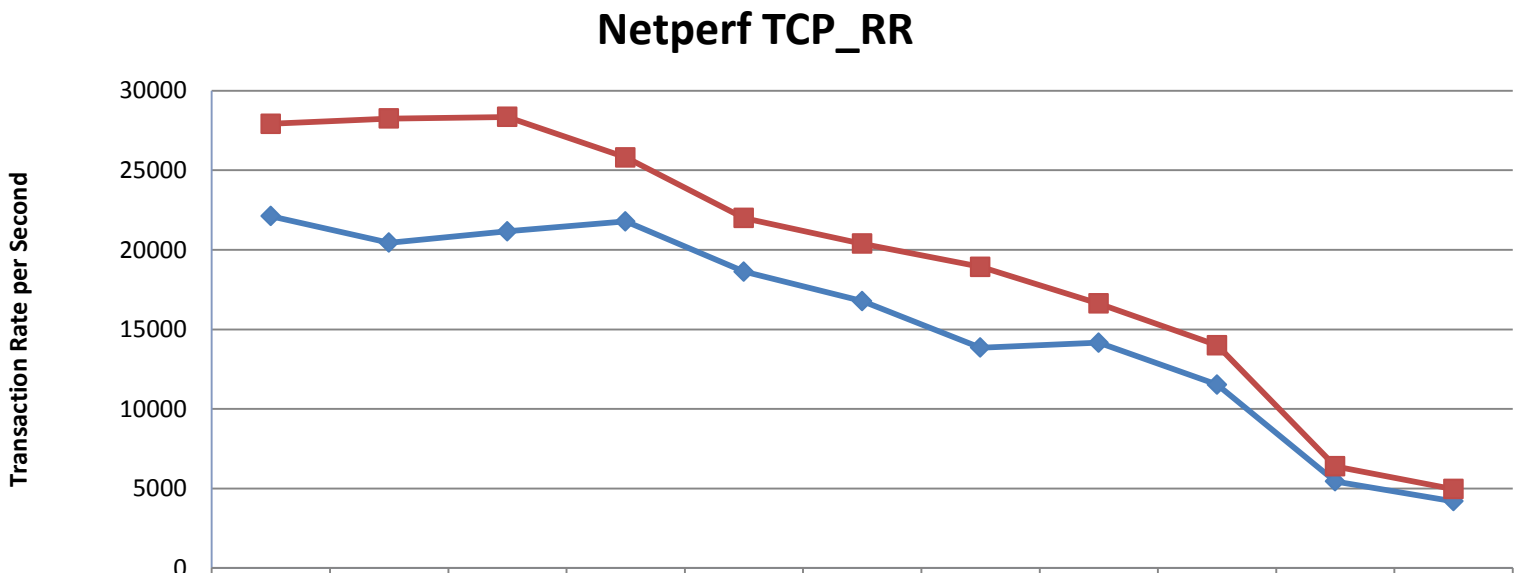

Package size (B)	4	8	16	32	64	128	256	512	1024	2948	4096
w/o VT-d PI	22110.49	20447.75	21160.81	21785.69	18628.54	16776.54	13845.39	14159.72	11527.57	5439.28	4193.9
w/ VT-d PI	27917.07	28252.58	28346.33	25808.36	21998.12	20388.7	18929.65	16620.59	13995.58	6390.74	4961.64

Host CPU Utilization

- Purpose:
- - Test the host CPU utilization (used by KVM)

- Method:
- Use netperf to send TCP stream package to guest. (This will generate lots of external interrupts)
- Meanwhile, use perf kvm top to measure the CPU utilization.

Host CPU Utilization

- w/o VT-d PI:
- Package
- Size (B)
- 4: kernel:17.1% us: 1.7% guest kernel:14.5% guest us:66.6%
- 8: kernel:13.5% us: 1.7% guest kernel:13.8% guest us:71.1%
- 16: kernel:13.6% us: 1.6% guest kernel:10.4% guest us:74.4%
- 32: kernel:12.5% us: 2.0% guest kernel:13.1% guest us:72.4%
- 64: kernel:12.8% us: 1.8% guest kernel:15.1% guest us:70.3%
- 128: kernel:12.8% us: 2.0% guest kernel:15.7% guest us:69.4%
- 256: kernel:13.2% us: 1.1% guest kernel:20.3% guest us:65.4%
- 512: kernel:12.7% us: 1.3% guest kernel:26.0% guest us:60.0%
- 1024: kernel:12.6% us: 2.2% guest kernel:30.8% guest us:54.4%
- 2048: kernel:13.0% us: 1.9% guest kernel:25.0% guest us:60.1%
- 4096: kernel:14.0% us: 1.3% guest kernel:22.9% guest us:61.8%

Host CPU Utilization

- w/ VT-d PI:
- Package
- Size (B)
- 4: kernel:10.3% us: 1.7% guest kernel:12.7% guest us:75.3%
- 8: kernel: 8.5% us: 1.4% guest kernel:13.1% guest us:77.0%
- 16: kernel: 8.5% us: 1.1% guest kernel:11.7% guest us:78.7%
- 32: kernel: 8.2% us: 1.8% guest kernel:14.5% guest us:75.5%
- 64: kernel: 8.7% us: 1.0% guest kernel:15.2% guest us:75.1%
- 128: kernel: 9.7% us: 1.2% guest kernel:18.1% guest us:70.9%
- 256: kernel: 9.0% us: 0.8% guest kernel:28.7% guest us:61.5%
- 512: kernel: 9.1% us: 0.8% guest kernel:28.5% guest us:61.6%
- 1024: kernel: 9.3% us: 1.2% guest kernel:30.7% guest us:58.8%
- 2048: kernel: 8.8% us: 1.9% guest kernel:28.0% guest us:61.3%
- 4096: kernel: 9.4% us: 1.2% guest kernel:27.1% guest us:62.4%

Host CPU Utilization

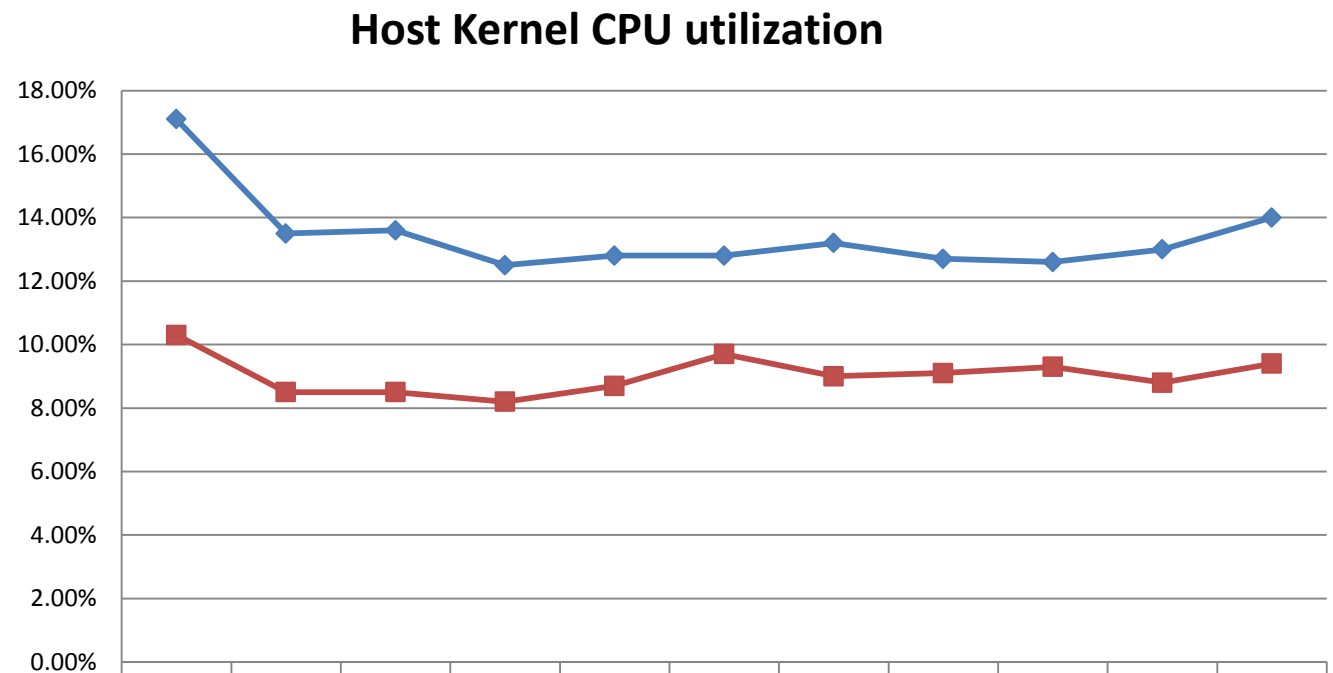

TCP package size (B)	4	8	16	32	64	128	256	512	1024	2948	4096
w/o VT-d PI	17.10%	13.50%	13.60%	12.50%	12.80%	12.80%	13.20%	12.70%	12.60%	13.00%	14.00%
w/ VT-d PI	10.30%	8.50%	8.50%	8.20%	8.70%	9.70%	9.00%	9.10%	9.30%	8.80%	9.40%

THANKS!