

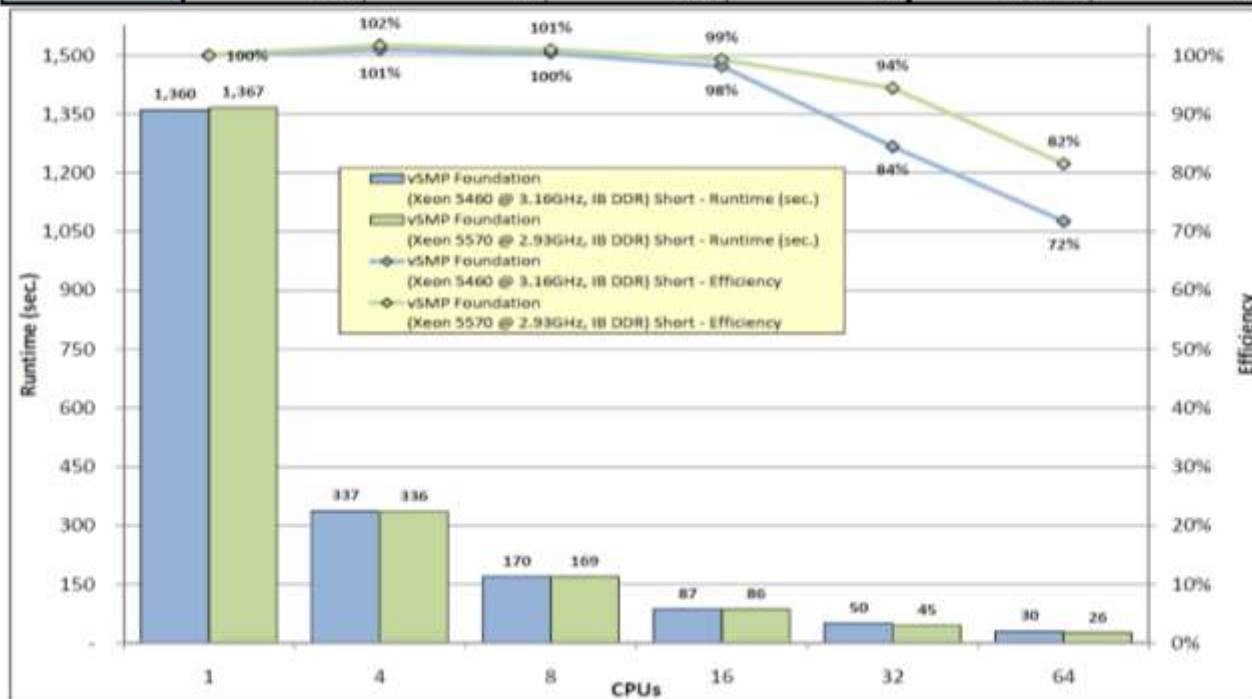
Scale-Up: Shared Memory

LANCZOS: OPENMP (SHARED MEMORY) ALGORITHM TO CALCULATE EIGENVALUES

Demonstrating close to linear scalability with OpenMP code

82% Efficiency with 64 CPUs (X5570) vs. 72% Efficiency with previous generation (5460)

System & Processor Model	vSMP Foundation (Xeon 5460 @ 3.16GHz, 1B DDR)		vSMP Foundation (Xeon 5570 @ 2.93GHz, 1B DDR)		vSMP Foundation (Xeon 5460 @ 3.16GHz, 1B DDR)	
	Short Runtime (sec.)	Short Efficiency	Short Runtime (sec.)	Short Efficiency	Long Runtime (sec.)	Long Efficiency
2 CPUs	1,360.49	100%	1,366.85	100%		
4 CPUs	336.94	101%	335.94	102%		
8 CPUs	169.51	100%	169.16	101%		
16 CPUs	86.67	98%	86.03	99%	96,095.00	100%
32 CPUs	50.33	84%	45.22	94%	48,124.00	100%
64 CPUs	29.61	72%	26.18	82%	24,293.95	99%

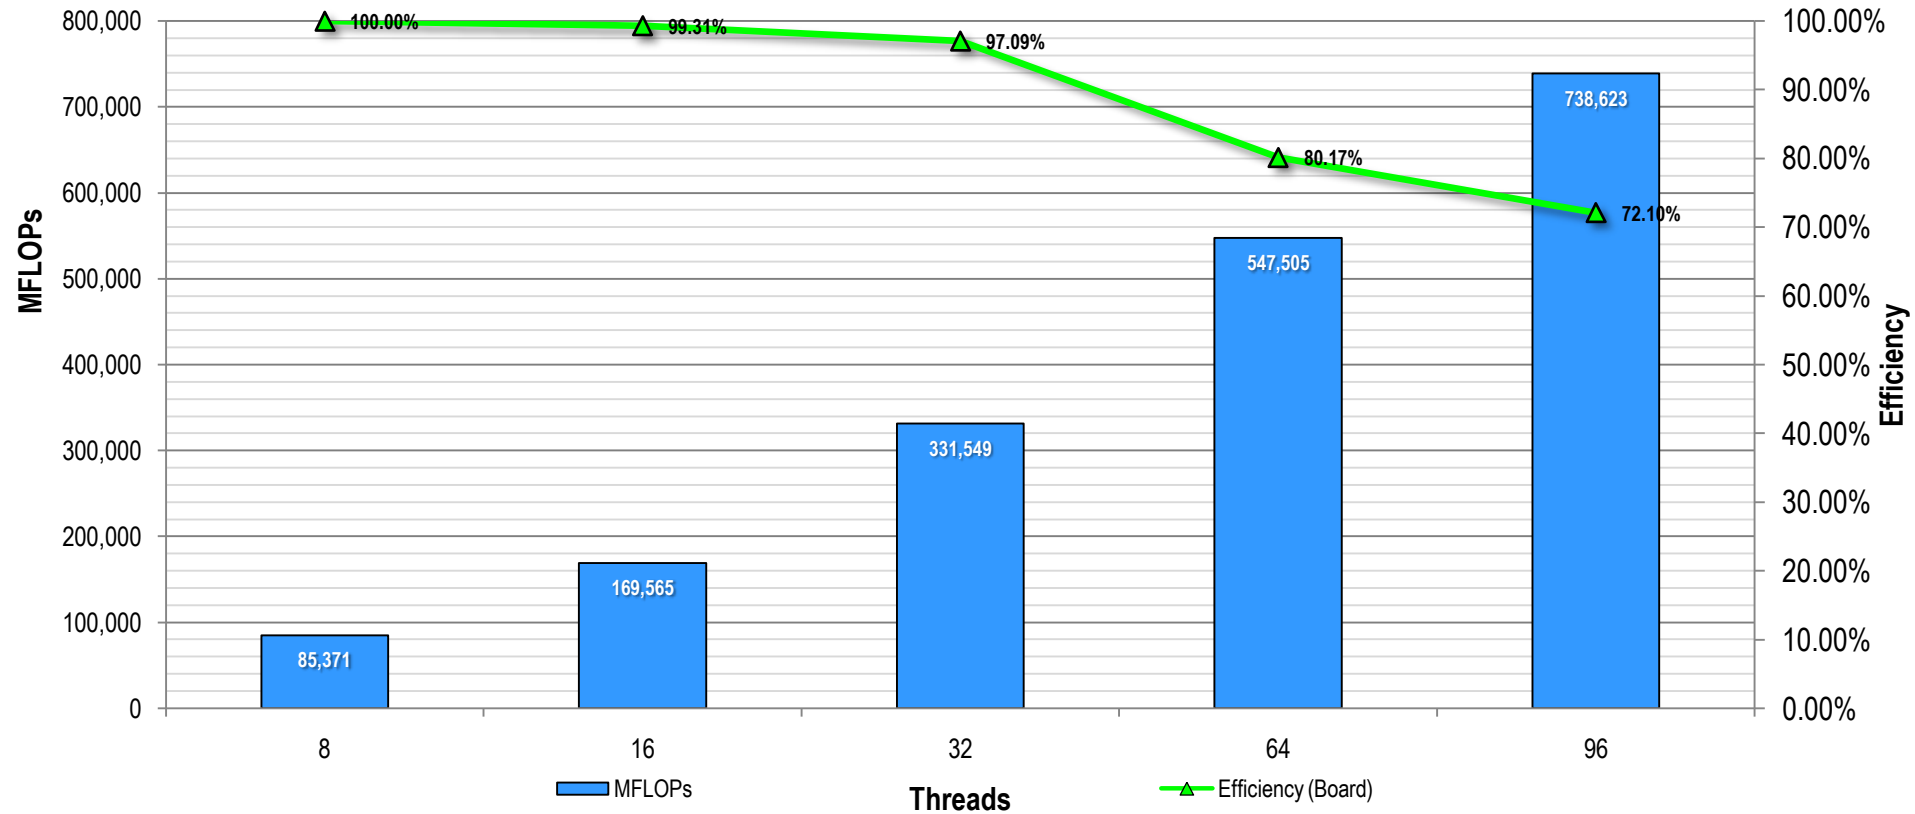


OPENMP PARALLELIZATION

DGEMM (INTEL MKL) – MATRIX SIZE: 12,000 X 12,000

The benefit of using a commercial mathematical library (with no additional modifications) which is NUMA aware

80%% Efficiency with 64 CPUs (X5570), 72% Efficiency with 96 CPUs



HW Characteristics:

12 x 2 x Intel XEON X5570 QC (2.93GHz, 6.4GT/s), HT off, Turbo Boost off; 83GB/192GB (4 x 6 x 4GB (1333MHz) + 4 x 6 x 4GB (1066MHz)); DDR; vSMP Foundation 2.0

(Source: ScaleMP)



Aggregate. Scale. Simplify. Save.

Scale-Up: Large Memory vs. I/O

NASTRAN: XL0TDF1 - CAE APPLICATION, PERFORMING EXTENSIVE I/O

Leveraging aggregated RAM results in

ScaleMP 3 X Faster runtime compared to HDDs, 2X compared to SSDs

