The Clear Solution for Data Center Cooling™

Extreme cooling performance with low-cost dielectric fluid submersion cooling

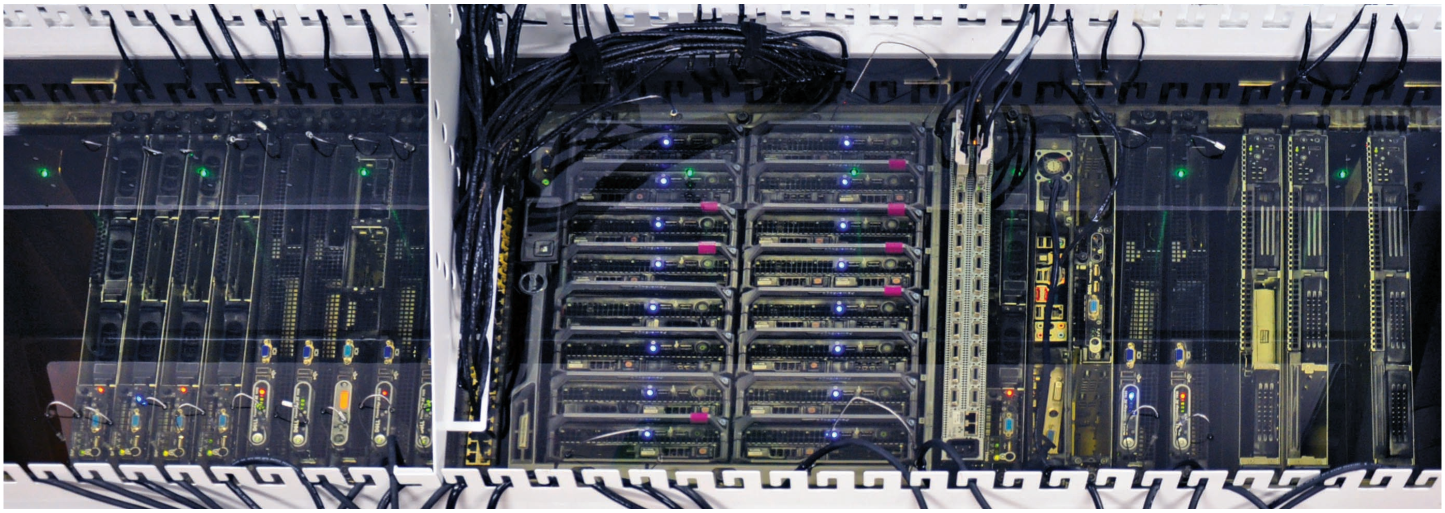

Green Revolution Cooling: Over-Clocking Servers with the CarnotJet™ System

Green Revolution Cooling Headquarters, Austin, TX, May 2010

ABSTRACT: In a series of tests designed to challenge the CarnotJet™ system's ability to cool over-clocked server CPUs, GRC over-clocked current generation Intel Xeon CPUs 35-60% while dissipating 200 Watts per socket without mechanical chilling.

Over-clocking is the practice of running a CPU at a clock frequency higher than the base rate. For example, it's possible to run a 2.6 GHz processor at 3.0+ GHz (15%+ increase in speed), but it will also use more power, which results in a significantly higher temperature and requires more cooling capacity. Anecdotally, gamers have run CPUs up to 100% faster than base using liquid nitrogen for extreme cooling. At the University of Illinois at Urbana-Champaign, IBM's *Blue Waters* supercomputer cools its over-clocked processors with water piped to the CPU. But these extreme methods for cooling demonstrate the fundamental lack of a sustainable cooling method to support over-clocking.

Server over-clocking is limited by air-cooling constraints and reliability concerns. Because server heat sinks are passive and up to several times smaller in size than ones used in desktop computers, they are much more limited in their capacity to dissipate heat. When we ran initial tests using air cooling, we were unable to over-clock our servers more than a few percent. But that all changed when we over-clocked with the CarnotJet™.

Experiment: We wanted to see if the CarnotJet™ fluid-submersion cooling system was up to the challenge of cooling over-clocked server CPUs. Could the CarnotJet™ cool servers with such high CPU temperatures effectively? How far could the CarnotJet™ push these chipsets? What kind of performance increase could it comfortably sustain?

To answer these questions, we first built three 1U servers with Xeon Nehalem CPUs (E5520: 2.26 GHz and X5550: 2.6 GHz). We then enabled over-clocking and installed the servers into a standard CarnotJet™ unit. We proceeded to run ninety (90) individual tests, including BIOS settings like base clock, voltages, and hyper-threading, and a variety of diagnostic tests to ensure that the systems were running properly.

Benchmarks: To measure output (GFLOPs) and residual values, we used Linpack (Intel MLK Xeon 64 bit) set to 90% memory utilization. We used lm-sensors in Linux to measure CPU core voltage and temperature, motherboard temperature, and memory usage. We used *WattsUp* power meters to measure server wattage.

Test results [Figure 1]: GreenDEF™ coolant typically entered the system at 35° C, rose to 37° C near the motherboard, and exited at 40° C. With standard (non over-clocked) systems, GreenDEF temperature of 40° C results in CPU temperature of 60° C (these CPUs are designed to run at ~75 C).

	CPU Readings		Outputs		Temperatures	
Chip	GHz	Core Volts	Watts	GFLOPs	CPU	Motherboard
E5520	2.26	1	175	35.9	54	32
E5520	2.82	1.1	218	43.3	66	32
E5520	3.2	1.1	254	51.0	64	27
E5520	3.4	1.1	272	53.35	74	32
E5520	3.66	1.2	320	54	83	32
X5550	2.66		207	43.5	65	37
X5550	3.0	1.18	233	48.5	70	37
X5550	3.2	1.24	254	51.4	77	37
X5550	3.4	1.24	270	55.2	79	37
X5550	3.5	1.3	296	56.7	82	37
X5550	3.6	1.34	323	57	84	37

Figure 1: Temperatures, Voltage, Server Total Watts consumption, and Output

Stable CPU over-clocking was achieved at speeds 35-60% faster than rated. We stopped tests when the CPU core voltage hit 1.35V in order to avoid potential CPU damage. These gains required relatively modest net temperature increases for CPUs, as the coolant proved adept at rejecting most of the additional heat.

Previously, we believed that over-clocking would produce too much heat to dissipate without adding electricity-intensive mechanical cooling. However, we were able to cool all over-clocked CPUs in this test using our low-power system.

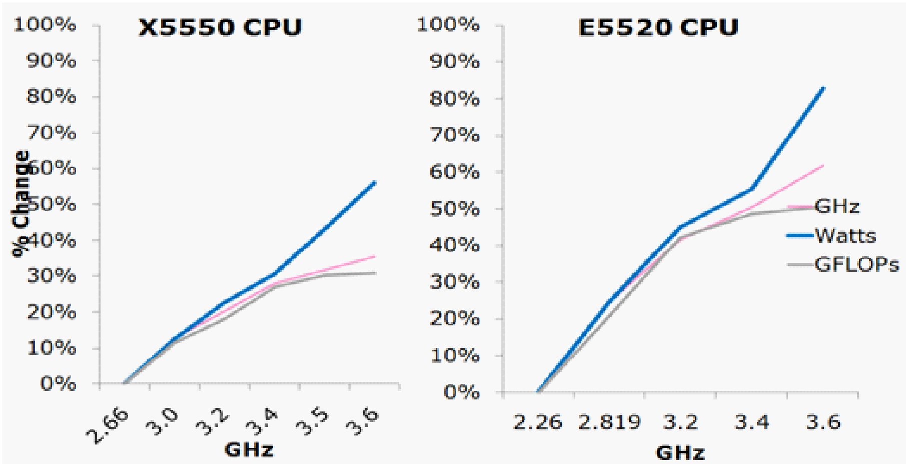

Figure 2: [LEFT]: X5550 over-clock to 3.6 GHz vs. its 2.66 GHz base (+35%) [RIGHT] E5520 over-clock to 3.66 GHz vs. its 2.26 GHz base (+62%)

Figure 2 shows the relationship between CPU processing power and energy use. Note that power per FLOP does not significantly increase until ~3.4 GHz.

Additional Findings: While we didn't measure power from the CPU itself, we did measure server power and can derive directional power usage for the CPU. We found power dissipation per socket can be over 200 Watts, compared to 80-95 Watts for the typical server CPU. This leads us to believe that CPU/GPU manufacturers would increase CPU speed if they had significantly better cooling capabilities (as evidenced by IBM's use of 3.5 - 4 GHz CPUs after adding liquid cooling to the nodes in *Blue Waters*).

More output from the same CPU/server can lower the cost per calculation because modest increases in speed are often accompanied by large increases in price. For example, a 2.4 GHz Intel CPU (E5620) costs $387 at the time of writing, while a 2.66 GHz CPU (E5640) costs $774. The increase in price is not commensurate with the gains in speed: 18% more frequency and 21% higher Linpack output result in 100% higher cost. By comparison, minimal cost is required to over-clock a server in the CarnotJet™ system.

	Frequency	CPU Watts	Non CPU W	Total W
Stock	2.66	95 (S)	80 (D)	175 (M)
Over-clocked	3.6	220 (E)	100 (D)	320 (M)

Figure 3: Power usage comparison. [M = Meter measured, S = Intel Specs, E = Estimate, D = Total watts – CPU watts]

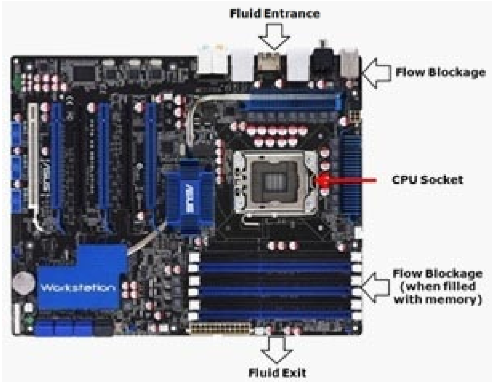

Figure 4: Terrain features of the ATX motherboard impeded flow of liquid coolant

Future Considerations: Our ATX motherboard used for over-clock testing [Figure 4], was not optimal for CPU heat dissipation in a server because the memory sticks disrupt coolant flow. Server motherboards do not have the same issue.

Conclusion: Servers comfortably achieved 35-60% gains in CPU output while over-clocking in the CarnotJet™ system. Although over-clocked chips required significantly more energy, there was minimal additional cooling energy needed to cool the system.

While we found little ability to over-clock servers using air-cooling, we found that fluid-submersion cooling can dissipate over 200 Watts per CPU, which has implications for centers contemplating GP GPUs or other high wattage (130W) CPUs. Over-clocking with fluid-submersion may be a cost-effective way to maximize longevity and utility of servers while providing stability and keeping cooling costs low.

