Supermicro FatTwin™ with Cool IT Liquid Cooling Technology

Ranks #400 on the Top500 and #41 on the Green500

When architecting a HPC system, power and cooling is a critical aspect. Pursuing ever higher performance is an endless effort for HPC, while power and cooling has been very often a gating factor of that effort. The increases of using higher memory capacity per processor core, the use of GPUs or Xeon Phi Coprocessors, and the system density per rack, result in higher power requirement s and heat concentration. The power consumption per rack commonly exceeds 20 kilowatts (KW), and in some cases, more than 30KW per rack, depending on the system configuration. Most data centers today are limited not just by power, but also by cooling capacity.

The Supermicro Cherry Creek cluster achieved 131.5 TFLOPS in performance while consuming only 74 kW in total power for an efficiency level of 1801 MFLOPS per Watt in 2 racks on the tradeshow floor at SC13. These achievements earned the Supermicro Cherry Creek cluster #400 place on the Top500 list of the world's fastest supercomputers while also positioning it at #41 on the Green 500 list of the world's most efficient supercomputers. The following whitepaper will outline how this was achieved in the tradeshow floor environment of SC13.

Supermicro® 's Cherry Creek at SC`13

Improved Performance, Density, and Efficiency

The Supermicro FatTwin™ SYS-627R2-FT+ is optimized through design and architecture innovations to deliver the highest levels of performance, efficiency, and density in data center, cloud, and HPC environments. Supermicro FatTwin™ efficiency levels are achieved through the use of large-diameter shared redundant fans which displace over 450 cubic feet per minute, motherboard designs with offset CPU and memory modules, optimized cooling airflow, redundant 1620W Platinum Level (94%+) high-efficiency power supplies, and the use of large-diameter shared redundant fans which displace over 450 cubic feet per minute. In addition to these efficiencies that lower total cost of ownership (TCO), the Supermicro FatTwin™ platform offers the most flexible and optimized server architecture for integrating custom third-party solutions such as Cool IT's Rack DCLC™ AHx Liquid cooling with Intel® Xeon Phi 7120P coprocessors

At the SC' 13 tradeshow, Supermicro's energy efficient FatTwin™ enabled with CoolIT liquid cooling technology, raised the operational thermal limits, performance and efficiency of both processors and co-processors in each of the 48 FatTwin™ nodes.

Supermicro® 's FatTwin™ **Cherry Creek Node**

Each of the liquid cooled FatTwin™ nodes support dual Intel® Xeon® E5-2697 v2 processors, 3x Intel® Xeon Phi 7120P coprocessors, 128 GB ECC RDIMM DDR3-1866MHz in 16 DIMMs, 1x Intel® SSD DC-S3700series 400GB and Intel® TruScale QDR 40Gb/s InfiniBand.

Supermicro FatTwin™ SYS-627R2-FT+ with Cool IT Rack DCLC™ AHx for Cherry Creek

CoolIT's Rack DCLC™ AHx™ rack-based liquid cooling solution enables high-performance and high-density clusters anywhere, without the requirement for facility liquid to be plumbed to the rack. The unique AHx configuration consists of a cooling liquid network that is mounted directly onto the Intel E5-2697V2 processors and Intel Xeon Phi 7120p co-processors. This system allows the processor heat to be directly absorbed into circulating liquid which then efficiently transports the heat to a liquid-to-air heat exchanger mounted on the top of the rack. This stand-alone rack solution is modular and is compatible with any rack-computing set-up, enabling ultra-high density clusters to be deployed quickly and easily without the complexity of integrating with facility liquid or chiller systems.

Front IO of Supermicro®'s FatTwin™ SYS-627R2-FT+ with Cool IT Liquid Cooling

Cool IT Rack DCLC™ AHx System

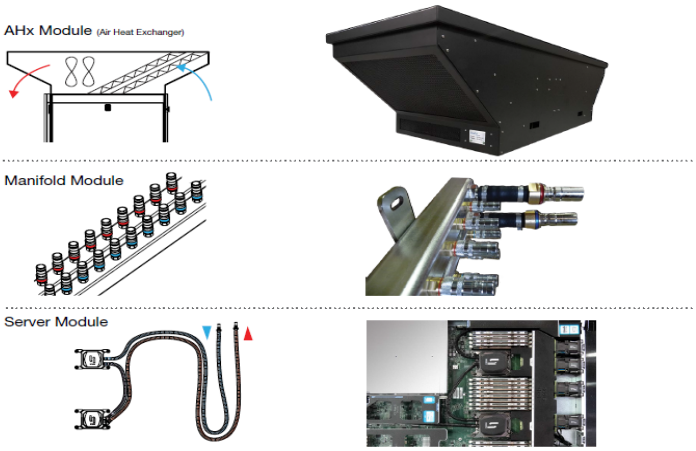

Manifold Module

Utilizing a stainless steel body and Staubli dry-break 'dripless' quick connects, the backbone to the entire system is incredibly robust. No plastics are used to ensure worry free operation for many years while the parts are color coded for easy installation and service.

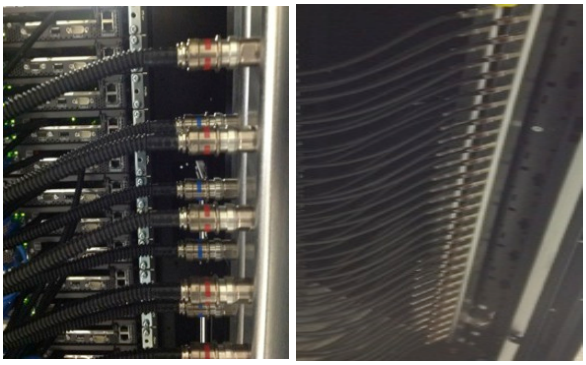

Cool IT Manifold Module

The manifold racks like a standard PDU, are flexible in size and orientation allowing servers to be hot swappable.

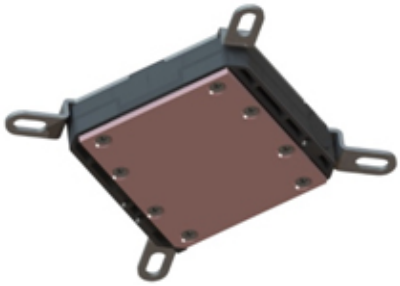

Utilizing a state of the art copper cold plate the server module's direct contact with the heat producing components in the server removes the heat via a liquid loop. The CoolIT Systems reference cold plate is a solid copper component that leverages several CoolIT design optimizations including a unique V-Groove micro-channel construct that maximizes performance for today's high thermal density applications (see Table 1.1).

Mechanical	
Fin Pitch	0.26 mm
Fin Density	100 Fins Per Inch
Surface Flatness	0.15 mm
Operational Range	5 °C ~ 80 °C
Sealing Methodology	O-Ring
Coldplate Material	Copper
Housing Plastic	PPS 40% GF
Thermal Resistance	0.046 °C/W @ 1L/min *
Pressure Drop	0.88 psi @ 1L/min *
Dimensions	Height 15.68 mm, Width 63.0 Length 63.0 mm
Socket Support	Intel 155X, 2011, 2011 narrow

TABLE 1.1

Connecting the system together is a network of ultra-high-reliability Fluorinated Ethylene Propylene (FEP) tubing.

Performance	
Operational Range	0 °C ~ 200 °C
Melting Point	255 °C
Bend Radius	< 20 mm
WVTR	2 g/m/year @ 45 °C, 5 g/m/year @ 60 °C
Tensile Strength	> 4,500 psi
MIT Flex life ASTM D 2176	250,000 cycles
Maximum Pressure	380 psi @ 25°C, 180 psi @ 8

TABLE 1.2

The cold insertion joint of a tube and barb is simply that of expanding a smaller polymeric tube over a larger barbed end feature. The seal and retention of this joint is created by the pressure or stress that is built up in the tube as the assembly is made.

A properly made joint is necessary to ensure full performance of the system. Force required to make the insertion has been measured to be between 300 and 400 Newtons (67 to 90 lb/f). This force is highly dependent upon proper alignment (see Table 1.2).

The corrugated tube allows for application requiring tight bends and more flexible applications. The design of this tube consists of both CVT (Corrugated Tube) and straight wall sections.

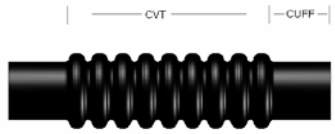

Liquid Cooled FatTwin™ Node for Cherry Creek

Supermicro's FatTwin™ Node used in the Cherry Creek deployment was unique in that processing power is highly condensed into a single FatTwin™ node, while having very few outlets for tubing to reach the rack level manifold. Utilizing an internal manifold, custom Xeon Phi coldplates and Intel Xeon coldplates CoolIT removes over 80% of the heat produced by the system.

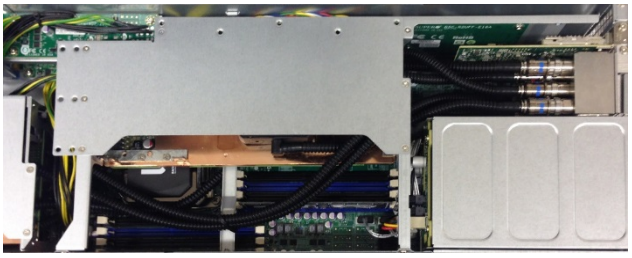

Supermicro® 's FatTwin™ Cherry Creek Node

Increasing Efficiency with Liquid

Even greater efficiencies are realized when a facility water source is used to transport the heat directly out of the building.

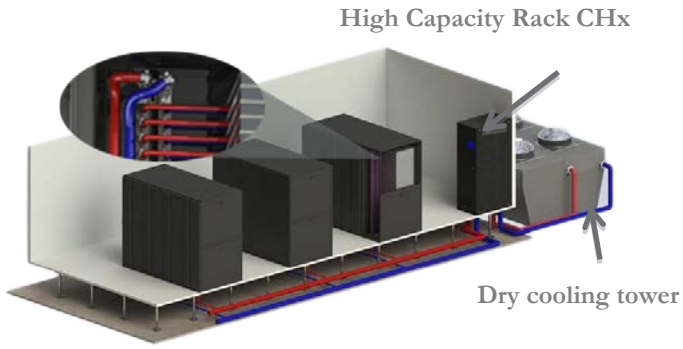

CHx-HC layout with Dry Cooling Tower

The High Capacity CHx Module (CHx-HC) is capable of managing the distribution of clean, treated coolant to and from a network of IT cabinets. The number of cabinets that can be serviced with a one CHx-HC is dependent on the organization and power load of IT equipment within a rack. Typically a single CHx-HC can manage 10-20 cabinets.

	CHx – HC
kW Capacity (@30C facility water)	750 kW
Racks per system	10-20
Power Consumption (Max)	4.2 kW
Centralized Pump Capacity (Max)	360 lpm

TABLE 1.3

Liquid cooling not only increases the efficiency of the Data Center, it reduces the overall footprint of the cluster with a minimal CAPEX increase while achieving a substantial OPEX decrease. Using a typical 240 rack deployment we see incredible savings while utilizing the same compute power.

	Air Cooled	Rack DCLC	
Racks Required for Same Compute	240	120	Half the space required
Total Capital Cost	$103.68M	104.46M	ROI < 1 year
Total Operating Expense ($/year)	$6.40M	$4.99M	24% reduction in OPEX
Total Data Center Power Use	11.23 MW	8.51 MW	2.72 MW power savings

TABLE 1.4

Summary

Supermicro's FatTwin™ technology combined with Intel technologies (CPUs, Xeon Phi Co-processors, and IB add-on cards), and CoolIT's Liquid Cooling were able to achieve a Top 500 (#400) and Green 500 (#41) Supercomputer in just 2 racks on the tradeshow floor of SC'13. On a per-rack basis, with fully populated hardware the total potential energy consumption per rack is 40-45 kW. 85%+ of the IT load is cooled by the Rack DCLC solution. This layout allows for a 50% reduction in floors pace and 2.72 MW power savings to manage the same IT load The overall power savings of Cherry Creek enables a 24% decrease in annual operating expenses and a return on investment of less than 1 year.