SuperMUC-NG

Dr. Herbert Huber, Head of High Performance Systems Department

SuperMUC-NG

TOP 500 SUPERCOMPUTER SITES

19.5 PFlop/s HPL Performance

Rank 8 (November 2018)

CPU	INTEL thin	INTEL fat
Processor type	Intel Skylake	Intel Skylake
Total number of nodes of this type	6336	144
Number of islands with this node type	8	1
Number of processors (sockets) per node	2	2
Number of cores per processor	24	24
Memory per node (GByte)	96	768
Nominal frequency (GHz)	3,10	3,10
Peak frequency (GHz)	3,90	3,90
Measured AVX-512 frequency (all cores active)	2,30	2,30
Floating point operations per core per clock (Fused MulAdd counts as 2)	32	32
Peak floating point performance of one core @ nominal frequency (GFlop/s)	99,20	99,20
Peak AVX-512 floating point performance of one node @ measured AVX-512 frequency (all cores active)	3533	3533
Total	INTEL thin	INTEL fat
Total CPU Cores	304128	6912
Total Peak PFlop/s @ measured AVX-512 frequency (all cores active)	22,38	0,51
SUM PEAK PFlop/s @ measured AVX-512 frequency (all cores active)	22,89	
Total Nodes	6480	

SuperMUC-NG

- Compute Node ThinkSystem SD650
  - ✓ 2 Intel Xeon 8174 processors
  - ✓ 12x DDR4 (R/LR) 2667MHz DIMM
  - ✓ 1x Intel OPA100 adapter
  - ✓ 1x onboard 1 GbE adapter

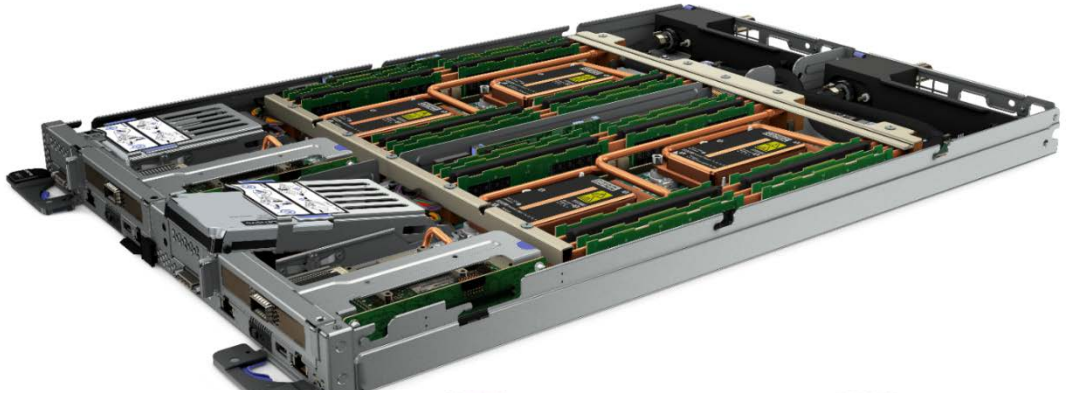

- Compute Node Chassis
  - ✓ Lenovo ThinkSystem SD650 1U full wide tray double-node in 6U12N chassis
  - ✓ 6x 2000W Platinum power supplies

- Intel OmniPath Interconnect Network
  - ✓ 8x 788 thin node islands fat tree network topology
  - ✓ 1x fat and management node island
  - ✓ 1x storage island
  - ✓ Blocking factor of 4 between islands

Lenovo ThinkSystem SD 650

Top-Down View

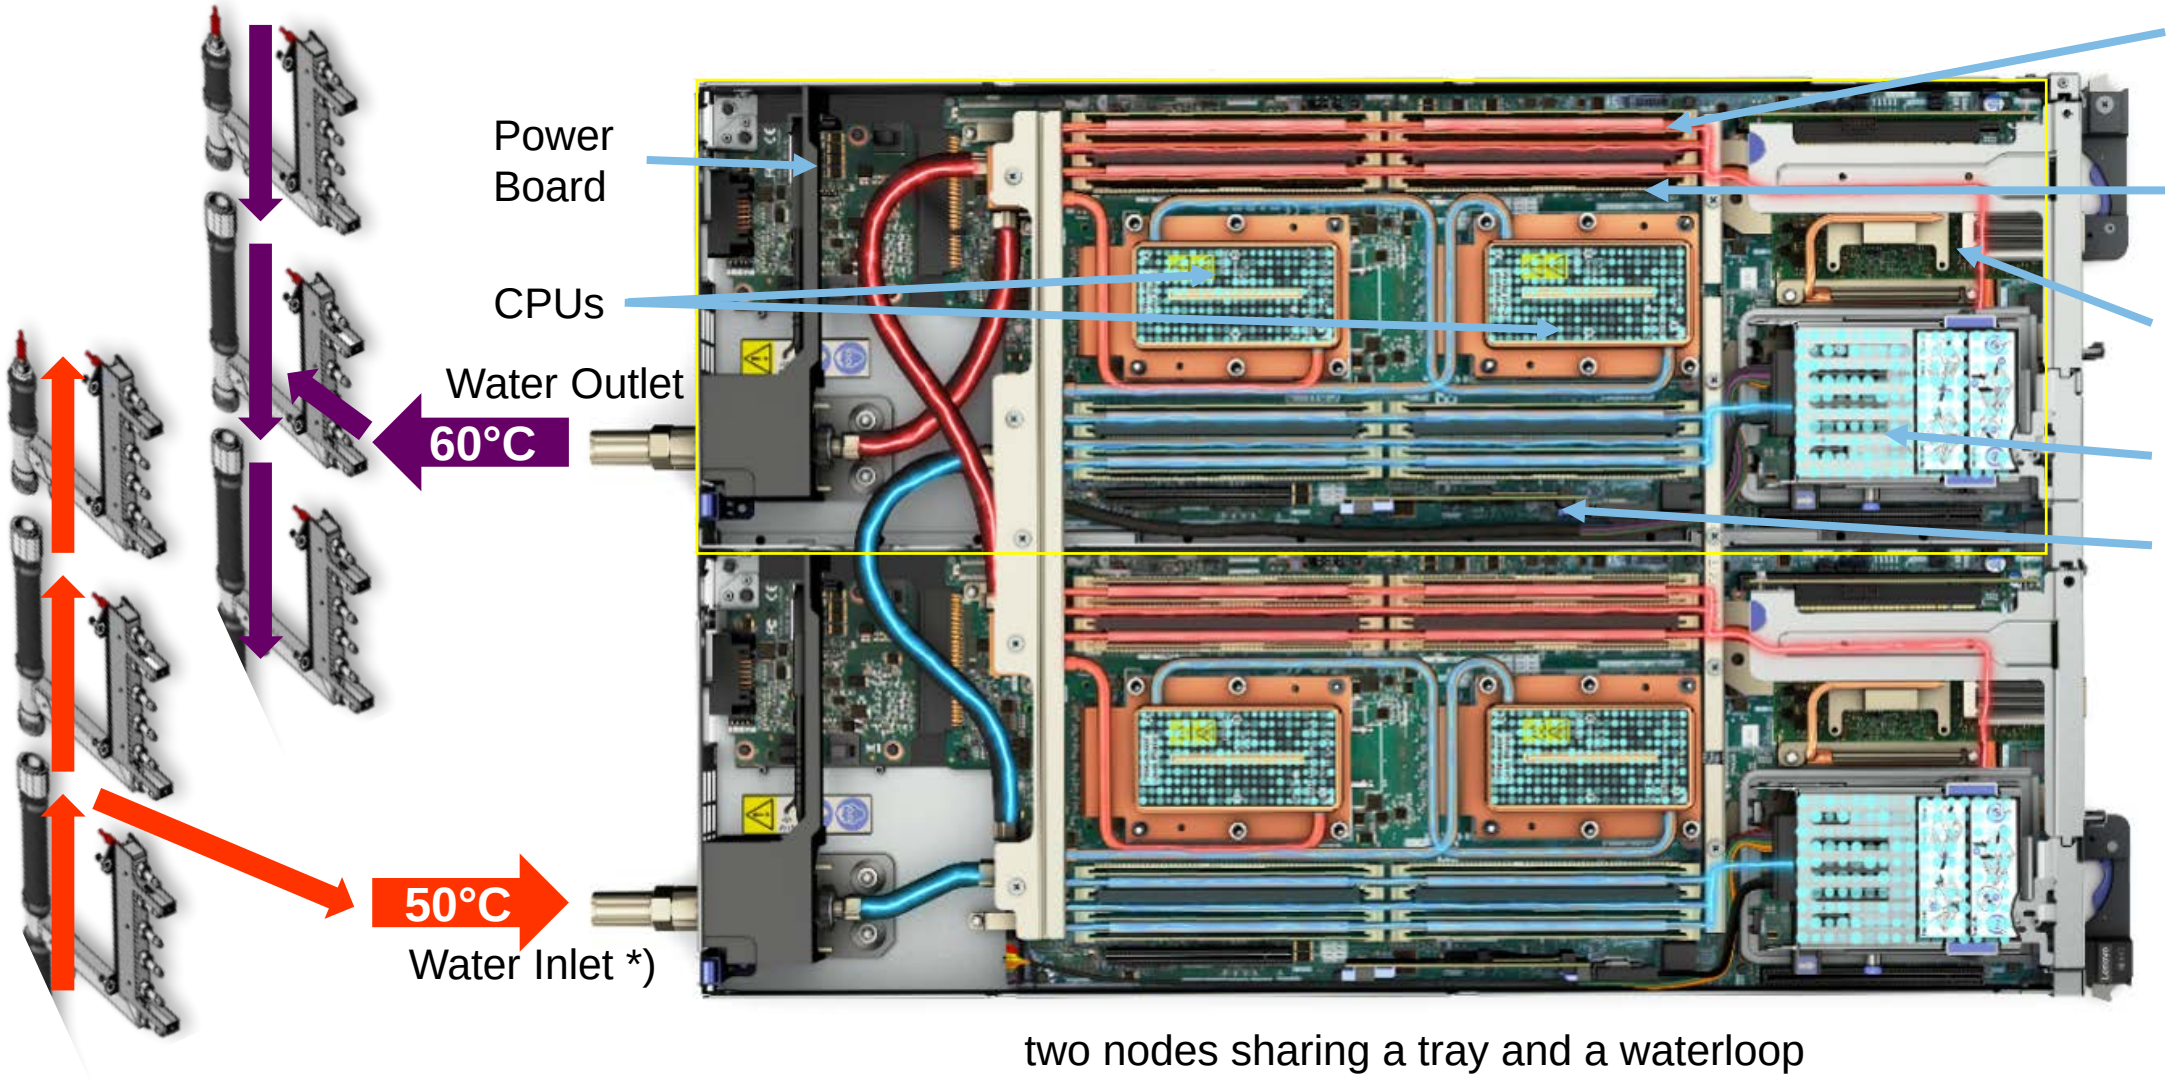

two nodes sharing a tray and a waterloop

2018 Lenovo - All rights reserved.

SuperMUC-NG

- Parallel File System with
  - ✓ 50 PByte capacity and
  - ✓ 500 GByte/s I/O bandwidth

- Data Science Storage for Long Term Data Storage with
  - ✓ 20 PByte capacity and
  - ✓ 70 GByte/s I/O bandwidth

- Software
  - ✓ Operating system: SLES12 SP3
  - ✓ Parallel file system: IBM Spectrum Scale (GPFS)
  - ✓ Batch scheduling system: Slurm
  - ✓ Programming environment: Linux + Intel Parallel Studio, OpenHPC

- OpenStack Compute Cloud with
  - ✓ 32 nodes with 2x Intel Xeon 6148 processors, 192 GByte memory and 1x Mellanox ConnectX-5 100 GigE adapter
  - ✓ 32 with nodes with 2x Intel Xeon 6148 processors, 2x Nvidia Volta 100 GPUs, 768 GByte memory and 1x Mellanox ConnectX-5 100 GigE adapter
  - ✓ 1 huge memory node with 8x Intel Xeon 8160 processors, 6144 GByte memory and 2x Mellanox ConnectX-5 100 GigE adapter

Cloud Computing

LRZ
A Pioneer of “Direct Warm Water Cooling”

Up to 50°C Hot Water Cooling

LRZ procurement requirement: 97% Heat Removal Efficiency!

2012

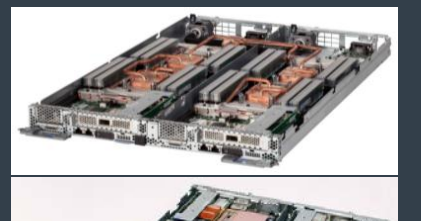

2014

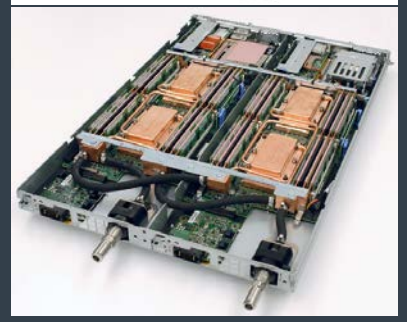

2018

Value of direct warm water cooling

- Reduced noise level in the Data Center
- Maximum TDP CPU Choice
- Reduced server power consumption
  - No fan per node (~ 10%)
- Reduce cooling power consumption
  - At 40°C free cooling all year long (> 15%)
- Energy Aware Scheduling
  - Only CPU bound jobs get max frequency (~ 5%)

> 30%
Total Saving

- Less conventional chillers needed to cool the Computing System
  - But conventional chillers are still needed
    - ✓ to cool the system waste heat of air cooled components such as network switches, power supplies and storage
    - ✓ as well as to cool the heat flux of the system into the room air (temperature of cooling devices in the nodes is more than 40°C → high surface temperature of nodes and compute node racks).

Re-Use of Waste Heat

How to reuse waste heat during summer?

Reuse of Waste Heat

- New buildings in Germany are very good thermally isolated: Standard heat requirement of only 50 W/m^2

  → SuperMUC-NGs waste heat would be sufficient to heat 40.000 m^2 of office space (~10 x space of LRZ office and visualization buildings)

- What to do with the waste heat during summer?

  ✓ Heat a large pool or market garden?
  ✓ Heat asparagus fields?
  ✓ Drive a Stirling motor?
  ✓ Drive an adsorption chiller?

Adsorption Chiller

The method of using solid materials for cooling via evaporation.

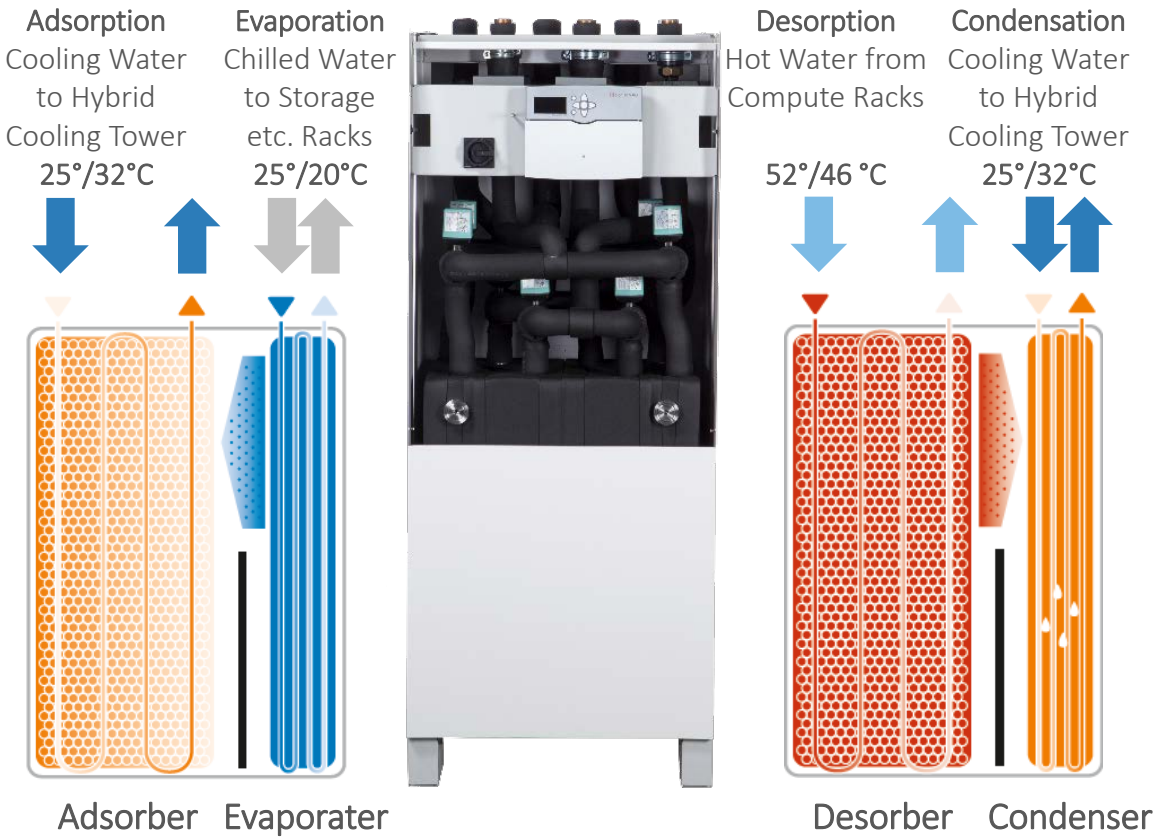

Module 1

Module 2

- Adsorption chiller consists of two identical vacuum containers, each containing two heat exchangers – and water.
  - **Adsorber (Desorber)**
    Coated with the adsorbent (e.g. silica gel)
  - **Evaporator (Condenser)**
    Evaporation and condensation of water
- Adsorption process has 2 phases
  - in the **adsorption phase** the water on the evaporator is taken in by the coated material in the adsorber. Through that evaporation the evaporater and the water flowing through it does cool down while the adsorber fills with water vapor and heats up the water flowing through it. When the adsorber is saturated the process is reversed.
  - in the **desorption phase** hot water is passed through the adsorber acting as a desorber rather as its desorbing the water vapor and dispensing it to the evaporator which is acting as condenser at that point condensing the vapor back to water. Again the process is reversed when the adsorber is emptied.

Value of Warm Water Cooling and Heat Reuse with Adsorption Chiller

- Reduced noise level in the Data Center

Energy Savings

- Maximum TDP CPU Choice
- Reduced server power consumption
  - No fan per node (~ 10%)
- Reduce cooling power consumption
  - At 50°C free cooling all year long (> 15%)
  - Heat reuse by adsorption chiller to generate 600 kW chilled water cooling capacity (~ 6%)
- Energy Aware Scheduling
  - Only CPU bound jobs get max frequency (~ 5%)
- No conventional chillers needed to cool the Computing System

> 35%
Total Saving

Lenovo
lrz