LBNL-5128E

ERNEST ORLANDO LAWRENCE
BERKELEY NATIONAL LABORATORY

“Hot” for Warm Water Cooling

Henry Coles, Lawrence Berkeley National Laboratory
Michael Ellsworth, IBM Corporation
David J. Martinez, Sandia National Laboratory
Anna-Maria Bailey, Lawrence Livermore National Laboratory
Farhad Banisadr, Los Alamos National Laboratory
Natalie Bates, Energy Efficient HPC Working Group
Susan Coghlan, Argonne National Laboratory
David E. Cowley, Pacific Northwest National Laboratory
Nicolas Dube, Hewlett Packard Corporation
Parks Fields, Los Alamos National Laboratory
Steve Greenberg, Lawrence Berkeley National Laboratory
Madhusudan Iyengar, IBM Corporation
Peter R. Kulesza, Oak Ridge National Laboratory
Josip Loncaric, Los Alamos National Laboratory
Tim McCann, SGI
Greg Pautsch, Cray Inc.
Michael K. Patterson, Intel Corporation
Richard G. Rivera, Los Alamos National Laboratory
Greg K. Rottman, U.S. Army Engineer Research Development Center
Dale Sartor, Lawrence Berkeley National Laboratory
William Tschudi, Lawrence Berkeley National Laboratory
Wade Vinson, Hewlett Packard Corporation

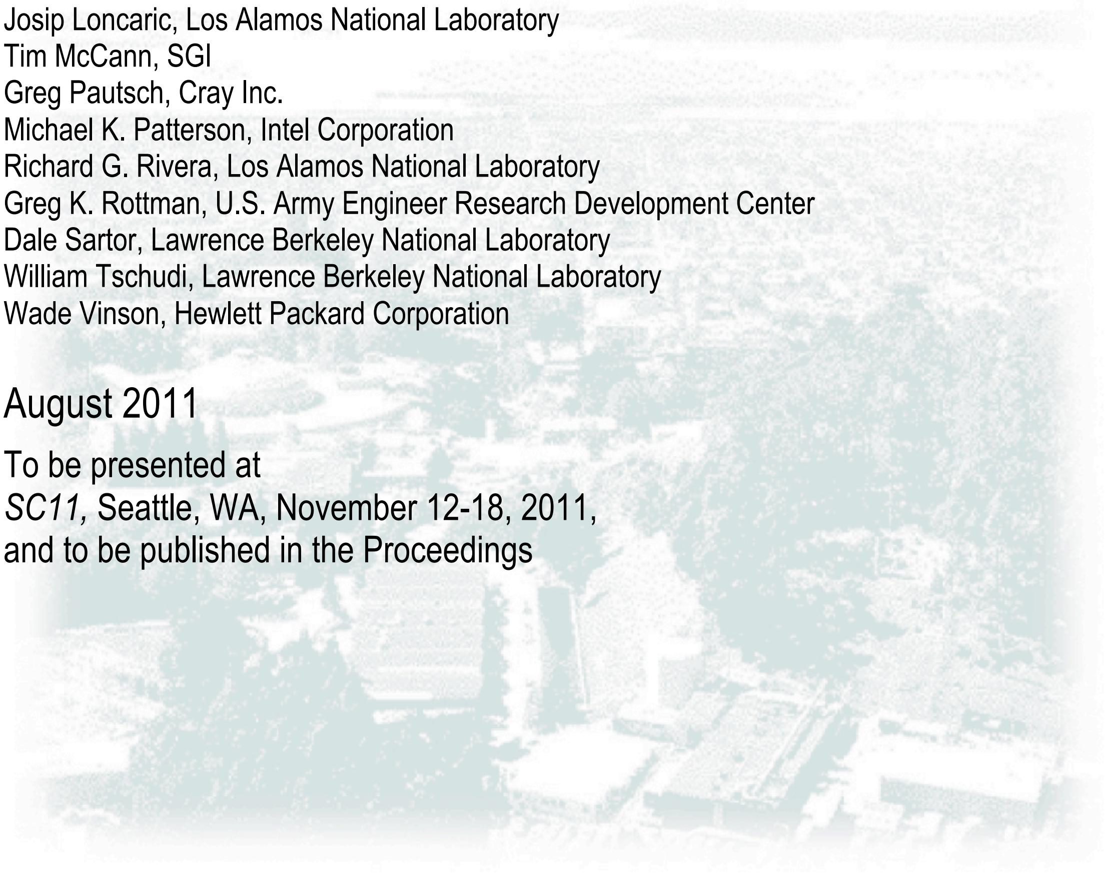

August 2011

To be presented at
SC11, Seattle, WA, November 12-18, 2011,
and to be published in the Proceedings

DISCLAIMER

“Hot” for Warm Water Cooling

Henry Coles
Lawrence Berkeley National Laboratory
One Cyclotron Road
Berkeley CA 94720
hccoles@lbl.gov

Michael Ellsworth
IBM Corporation
2455 South Rd.
Poughkeepsie, NY 12601
mje@us.ibm.com

David J. Martinez
Sandia National Laboratory
P.O. Box 5800
Albuquerque, NM 87185
davmart@sandia.gov

ABSTRACT

Liquid cooling is key to reducing energy consumption for this generation of supercomputers and remains on the roadmap for the foreseeable future. This is because the heat capacity of liquids is orders of magnitude larger than that of air and once heat has been transferred to a liquid, it can be removed from the datacenter efficiently. The transition from air to liquid cooling is an inflection point providing an opportunity to work collectively to set guidelines for facilitating the energy efficiency of liquid-cooled High Performance Computing (HPC) facilities and systems. The vision is to use non-compressor-based cooling, to facilitate heat re-use, and thereby build solutions that are more energy-efficient, less carbon intensive and more cost effective than their air-cooled predecessors. The Energy Efficient HPC Working Group is developing guidelines for warmer liquid-cooling temperatures in order to standardize facility and HPC equipment, and provide more opportunity for reuse of waste heat. This report describes the development of those guidelines.

General Terms

Design, Economics, Standardization

Keywords

direct liquid cooling, dry cooler, cooling tower

1. INTRODUCTION

Approximately one third to one half of an air-cooled datacenter's energy consumption is wasted on powering the cooling systems that keep the computer system from overheating. Furthermore, the amount of heat needing to be dissipated by future supercomputers limits the practicality of air cooling. Liquid cooling is key to reducing cooling energy consumption for future supercomputers because the heat capacity and transfer efficiency of liquids is orders of magnitude greater than that of air.

The transition from air-to-liquid cooling is a technology inflection point providing an opportunity to set guidelines for facilitating the energy efficiency of liquid-cooled facilities and systems. Current practice is to use vapor-compression refrigeration systems to provide chilled water or refrigerant solutions for cooling. Substituting cooling towers, hybrid cooling towers, or dry coolers that provide warmer water to supercomputers is the natural progression towards energy efficiency. The U.S. Department of Energy (DOE) supercomputing laboratories are working collaboratively with industry representatives to develop guidelines for warmer liquid-cooling temperatures to guide future supercomputer procurements, and to standardize the design basis for warmer temperature cooling systems.

The vision is to build liquid-cooled solutions that do not require compressors making them more energy-efficient, lower carbon, and more cost-effective than their air-cooled predecessors. The net result will be significant cost savings, reduced capital expenditures, reduced energy bills as well as reliability improvements. Secondary goals are to reduce or eliminate water consumption (i.e. evaporation in cooling towers) and enable more productive use of heat recovered from the supercomputers.

The national laboratories collaborate through the Energy Efficient High Performance Computing Working Group referred to as EEHPCWG, which has approximately two hundred members from supercomputing centers, industry and academia. This working group has prioritized a number of areas to advance such as liquid cooling guidelines with a goal of dramatically improving overall energy performance while maintaining high computational ability. The working group is supported by the DOE Federal Energy Management Program and Sustainability Performance Office.

2. HPC LIQUID COOLING TEMPERATURE GUIDANCE

The EEHPCWG is currently focused on defining liquid cooling guidelines for future use. The goal is to help National Laboratory supercomputer sites by providing procurement guidelines for new supercomputer equipment. These guidelines will specify liquid cooling temperature ranges for liquids cooled by cooling towers or dry coolers. This will establish a common design goal between supercomputer manufacturers and the supercomputer facilities for the definition of liquid cooling temperatures supplied to the supercomputer at the building interface point; see Fig 1. There are a number of attributes necessary to define cooling liquid supplied by the building and provided to the supercomputer.

Liquid cooling guidelines may include:

- Supply temperature minimum and maximum
- Minimum return temperature increase compared to supply
- Quality – chemical and impurity limits
- Percent of total energy removed by the liquid
- Maximum liquid static pressure
- Minimum liquid delta pressure
- Rate of change of supply temperature

Each of these subjects can be treated separately. Initially the EEHPCWG attempted to define more than one specification at a time and decided to focus on the maximum supply temperature along with some investigation on the supply minimum temperature. The maximum supply temperature that can be produced is easily related to environmental conditions given the assumptions listed below with the assumed cooling infrastructure as will be explained in the following sections.

3. ANALYSIS

3.1 Introduction

Liquid-cooling guidelines for IT equipment are defined for cooling liquid supplied by the building to the IT Solution. This guideline does not cover free-air cooling solutions such as the use of outside-air cooling of datacenters. Examples of liquid-cooling solutions are rack coolers or in-the-row coolers that use liquid to remove the heat but cool the electrical components using air, see Figure 1. Other examples use liquid at or near the electrical components to provide the needed cooling using conduction or forced liquid convection and remove the heat by way of a liquid to liquid cooling distribution unit (CDU); see Figure 2. The supplied solution removes a substantial amount of the heat generated by the combination of IT and required cooling equipment. 90 percent or more of this heat is typically removed. The point where the solution connects with the supplied cooling liquid is the building interface point, also referred to as the interface. See Figs. 1 and 2. The goal of the following analysis is to define the water temperature maximum as supplied by the building and accepted by the IT solution. This does not imply that a building needs to supply this maximum temperature - only that the supply temperature should not exceed this value. The maximum supply temperature is measured at the building interface point, see Figure 1. Figure 1 shows air-cooled IT equipment using a solution requiring a CDU, for example a refrigerant to liquid CDU. Figure 2 is a diagram of components that might be found for a direct liquid cooling solution supplied with an internal CDU. The total solution can be supplied in one or more cabinets or modules.

Figure 1: Example of an IT Equipment Cooling Solution Using Air at the Chip for Cooling and Equipped with a Liquid-to-Liquid CDU Connecting with the Building Interface

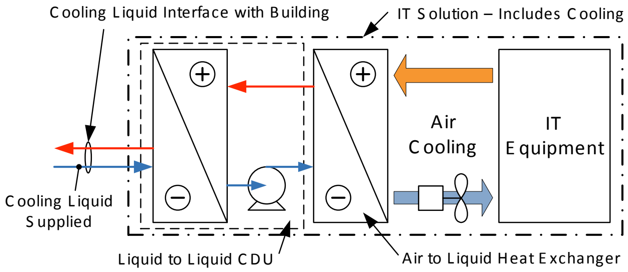

Figure 2: Example of an IT Equipment Solution Using Direct Liquid Cooling at the Chip and Equipped with a Liquid to Liquid CDU that Connects with the Building Interface

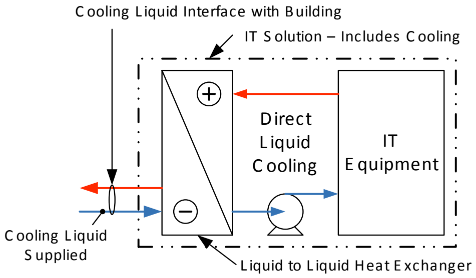

3.2 Methods

This report addresses only the maximum water temperature supplied. The study of additional attributes is planned as time and resources are made available. To clarify the guideline, the building is not required to supply the maximum temperature but if the high temperature liquid is presented to the IT equipment it shall accept that temperature and provide full performance and reliability for an extend period of time if desired by the datacenter owner. If the building supplies a liquid temperature outside the guideline range, the IT equipment can reduce cooling requirements by adjusting performance or shutting down gracefully as needed to protect data or avoid permanent damage.

Defining the upper temperature limit of cooling liquid supplied by the building is required to allow the datacenter owners a way to plan for the most efficient cooling system design that is compatible with their outside environment while providing cooling liquid within the guidelines. The same upper temperature limit is used by the supercomputer manufacturers as part of the thermal design basis for the IT equipment and cooling components required with a purchase.

The processes that produce the cooling fluid provided by the building are constrained starting at the ambient conditions outside the building. Many different cooling equipment constructs can be found installed across the industry for these processes. The water temperature supplied by the building starting with the outside ambient conditions and ending at the interface point is investigated. Systems that don't require compressor cooling are the primary focus of this study. Therefore cooling beginning with cooling towers (performance is constrained by outdoor wet bulb temperature) or dry coolers (performance constrained by outdoor dry bulb temperature) are studied.

The methods and analysis process has four steps plus a recommendation:

- Select U.S. National Laboratory HPC sites
- Obtain the wet bulb and dry bulb temperature environmental cooling design conditions for locations at or near the HPC sites using ASHRAE data.
- Define cooling path constructs used for analysis. Investigate approach temperatures for components used in the defined constructs. (The approach temperature is the delta between the supplied cooling fluid and the leaving temperature of the fluid or surface being cooled.) The constructs and approach temperatures are used to investigate cooling fluid temperatures that can be provided at the building interface and estimate the server thermal design margin.
- Investigate results of forecast cooling liquid temperatures, cooling constructs and environmental conditions that point to natural "break points". Predict thermal design margins using a CPU chip commonly used in HPC IT equipment solutions as a check of IT equipment thermal design feasibility for the constructs analyzed. The electronic device used for this study was the Intel Xeon 5500 series model EC 5545 operating at 85 watts.
- Propose liquid cooling maximum temperatures supplied at the building interface for dry cooler and cooling tower installations and seek ASHRAE adoption of the guidelines.

3.3 Selected National Laboratory HPC Sites

We selected 15 U.S. National Laboratory sites. These locations are spread across the country. The Houston Texas location was added as part of the investigation to see how a city known for high temperatures and high humidity compares to other locations.

Locations used for this study:

- Lawrence Berkeley National Laboratory (LBNL), Berkeley California
- Lawrence Livermore National Laboratory (LLNL), Livermore California
- Pacific Northwest National Laboratory (PNNL), Richland Washington
- Los Alamos National Laboratory (LANL), Los Alamos New Mexico
- Sandia National Laboratory (SNL), Albuquerque New Mexico
- Thomas Jefferson National Accelerator Facility (Jefferson Lab), Newport News Virginia
- Oak Ridge National Laboratory (ORNL), Oak Ridge Tennessee
- National Renewable Energy Laboratory (NREL), Golden Colorado
- Princeton Plasma Physics Laboratory (PPPL), Princeton New Jersey
- Stanford Linear Accelerator Center (SLAC), Menlo Park California
- Argonne National Laboratory (Argonne), Argonne Illinois
- Idaho National Laboratory (INL), Idaho Falls Idaho
- Fermilab, Batavia Illinois
- Ames National Laboratory (AML), Ames Iowa
- Brookhaven National Laboratory (BNL), Upton New York
- Houston Texas (location of interest)

3.4 Wet-Bulb and Dry-Bulb Design Temperatures

The LCWG voted to use the published ASHRAE wet and dry bulb cooling design temperatures corresponding to the 99.6 percent (also known as the 0.4 percent design temperature) of hours design limit to assess environmental conditions at each site. 99.6% corresponds to all but 35 hours per year predicted to be lower than the ASHRAE temperatures. The wet bulb temperatures are used to predict limits when using a cooling tower as the primary source of cooling water. The dry bulb design temperature is used to predict the limits with dry coolers. Dry coolers are air to liquid (usually water or a mixture of water and glycol) heat exchangers. The dry cooler heat exchangers are located outside of the building and typically have fans to provide air flow.

The ASHRAE design data for sites listed above were obtained from the ASHRAE Handbook CD included with the 2009 ASHRAE Handbook of Fundamentals.

The data obtained from the ASHRAE data base is shown in tabular format in Appendix A. There are two tables in Appendix A; one sorted by dry bulb and the other sorted by web bulb temperature. In many cases there is no ASHRAE environmental data for the National Laboratory location city; in those cases a nearby location available in the data base was selected. The locations selected in the data base are listed in the tables in Appendix A. A bar chart type format using the same data as that in the Appendix A tables sorted from high to low for the wet bulb and dry bulb temperatures is presented in Appendix B.

The dry-bulb temperatures across all sites are predicted to be at or below 99.5°F (37.5°C) 99.6% of the time during a typical year. The remaining 0.4% of the time corresponds, as previously mentioned, to 35 hours per year. The 35 hours are not distributed evenly and would be concentrated within the hotter months during the day time. The wet-bulb temperatures across all sites are forecast to be at or below 79.7°F (26.5°C) 99.6% of the time during a typical year. Therefore 37.5°C (99.5°F) will be used to analyze designs using dry coolers as the primary cooling source and 26.5°C (79.7°F) for analyzing design using cooling towers as the primary cooling source.

3.5 Cooling Architecture Constructs

In order to predict the liquid temperatures that the building could supply, critical temperatures for the processors were assumed. The processor temperature limit is defined as the maximum case temperature or Tcase max. The ability to cool below these critical temperatures can be forecast by adding the approach temperatures (change of temperature through various cooling components) in series starting from the outside environment up to the processor case. There are many possible designs to produce cooling water but by making a few assumptions we arrive at a reduced set of combinations to consider.

We made the following assumptions:

- We used an Intel CPU 5500 processor as the critical electronic component in terms of cooling. The Intel 5550 processor model EC5545 is commonly used in HPC IT equipment. The critical temperature is referred to as Tcase max and is assumed to be 77.5°C (172°F), consistent with the March 2010 Intel Thermal/Mechanical Design Guide. The power level is assumed to be 85 watts.
- The approach temperatures of thermal components, materials or devices in the heat transfer path from the ambient temperature at the cooling tower or dry cooler to the processor case, are added to forecast the margin of chip case allowable temperature. For heat transfer components where no fluids are involved, the approach temperature term is conduction and/or forced convection delta temperature and is accounted for in the same manner as adding approach temperatures.
- Constructs relying primarily on a cooling tower to generate cooling liquid are assumed to have one liquid to liquid heat exchanger; see Fig 3. A plate and frame type heat exchanger is commonly used. Plate and frame heat exchangers are very efficient at transferring heat from one liquid to another without mixing the two fluids.
- Constructs relying primarily on a dry cooler to generate cooling liquid are assumed to have one liquid to liquid heat exchanger. This heat exchanger is referred to as a CDU see Fig 4. CDU devices commonly use plate and frame type heat exchangers.
- In some cases the facility owner or solution supplier requires an additional heat exchanger component to control temperature, reduce leak risk, manage condensation or better control cooling liquid quality. The constructs described below contain the minimum or close to the minimum number of heat exchangers. We show that there is sufficient thermal cooling margin using direct liquid cooling for additional heat exchangers if needed in all cases.

- Two methods of cooling at the chip level were assumed in order to find the approach temperatures:
    - Air cooling – typical constructions use finned heat sinks held to the top surface of critically cooled components and a number of fans inside the server chassis provide air flow across the heat sink.
    - Direct Liquid Cooling (a.k.a. Direct Cooling) – this technology is new and designs vary. One concept uses conduction to transfer heat from the top of the chip to the bottom of a cold plate cooled by refrigerant.
- A concept of "Pre-Heat" is used. Pre-Heat is defined as heat absorbed by the cooling medium (air or liquid) as it passes from one component to another in the heat transfer path. Pre-Heat is added in the same manner as approach temperatures of other components.
- There are two Pre-Heat types:
    - Air Cooling at the Chip Level – After entering the front bezel, the air is heated by components such as disk drives and memory modules prior to entering the chip-cooling heat sink. This temperature delta is Pre-Heat.
    - Direct Cooling – Liquid flowing in serial paths cooling multiple components or absorbing heat from surroundings increases in temperature, this reduces the available cooling. This reduced cooling is accounted for with the addition of pre-heat.

Combinations yielding four constructs were investigated. Two heat transfer processes between the system and environment (primary cooling process) and two types of cooling at the chip level were selected to yield four cooling combinations listed in Table 1.

Table 1. Combinations of Primary Cooling Process and Cooling at the Chip Level Considered for this Study

Primary Cooling Process	**Cooling at the Chip Level**	**Figure Number**
Cooling Tower	Direct Liquid Cooling	3
Dry Cooler	Direct Liquid Cooling	4
Cooling Tower	Air Cooling	5
Dry Cooler	Air Cooling	6

Figure 3 shows the components used to forecast Tcase max. for the Intel 5500 processor chip when cooled with direct cooling at the chip and using a cooling tower as the primary heat transfer to the environment.

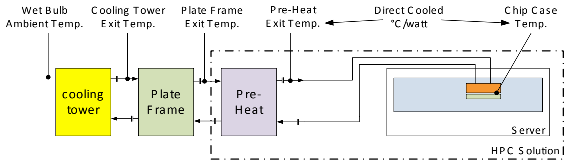

Figure 3: Cooling Tower with Direct Liquid Cooling at the Chip

Figure 4 shows the components used to forecast Tcase max. for the Intel 5500 processor chip when cooled with direct cooling at the chip and using a dry cooler as the primary heat transfer to the environment.

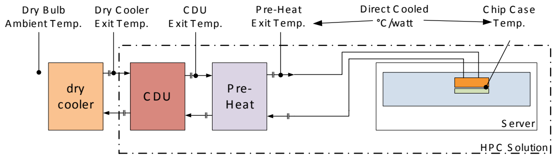

Figure 4: Dry Cooler with Direct Liquid Cooling at the Chip

Figure 5 shows the components used to forecast Tcase max. for the Intel 5500 processor chip when cooled with air cooling at the chip and using a cooling tower as the primary heat transfer to the environment.

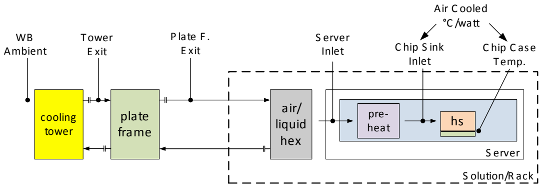

Figure 5: Cooling Tower with Air Cooling at the Chip

Figure 6 shows the components used to forecast Tcase max. for the Intel 5500 processor chip when cooled with air cooling at the chip and using a dry cooler as the primary heat transfer to the environment.

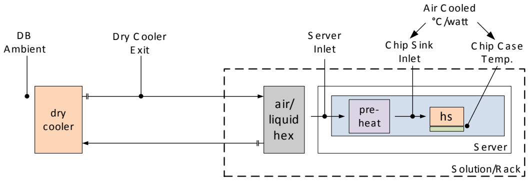

Figure 6: Dry Cooler with Air Cooling at the Chip

4. Approach Temperature Assumptions

A forecast for each component approach temperature value in the previous section is listed in Table 2. A few comments are outlined below. Reference information or assumption explanation for each component approach value is listed in Section 8, to locate the reference information use the superscript number found in Table 2. The approach temperatures assumed are thought to be what can be achieved with good design and are not a best case. For example the approach temperature for a dry cooler is listed at 10°F. Better performance for dry coolers is practical by increasing the size, there will be an increase in purchase cost and additional space is needed, of course a diminishing returns analysis is needed for each case considering capital and recurring costs.

Chip air cooling using a heat sink performance of approximately 0.225 degrees Celsius per watt is thought to be easily achievable in current designs. If more space is allowed for a larger heat sink or improved air flow better performance may be achieved.

The value used for direct cooling pre-heating may turn out to be generous considering the technology is relatively new. This pre-heating value comes from a company with some experience with recent prototype designs. One can imagine different designs with lower pre-heating values. Future market forces will determine the range of this value and is likely to change over time. The pre-heating associated with air cooling at the chip level is well known for current designs, reconfiguring server layouts at reduced densities may enable improvements.

Cooling distribution units (CDUs) are assumed to have larger approach temperatures compared to plate and frame heat exchangers used to separate cooling tower water from the building chilled water supply. The space available for a CDU heat exchanger is more constrained due to CDUs being commonly found inside or near the datacenter room. The smaller allowable space results in heat exchangers with larger approach temperatures.

Table 2. Heat Transfer Component Approach Temperatures

Heat Transfer Component	Description	Temperature Or Delta Temperature
Open Cooling Tower	wet bulb temperature to water temp. leaving	3.8°C (7°F)[1]
Dry Fin Cooler	dry bulb temperature to liquid temp. leaving	5.5°C (10°F)[2]
Plate and Frame Heat Exchanger	cooling liquid temp. entering to cooled liquid temp. leaving	1.67°C (3°F)[3]
Cooling Distribution Unit (CDU)	cooling liquid entering to cooled liquid leaving	2.77°C (5°F)[4]
Air to Liquid Heat Exchanger	cooling air entering to cooled fluid leaving	9°C (16.2°F)[5]
Server Bezel Pre-Heat	server air entering temp. to chip heat sink entrance temp.	3°C (5.4°F)[6]
Direct Liquid Cooling Pre-Heat	allowance for heat transferred to cooling fluid and non-parallel circuits	5°C (9°F)[7]
Chip Air Cooling Heat Sink	cooling air temp. entering heat sink to component case temp.	0.225 °C/watt (0.405 °F/watt)[8]
Chip Direct Liquid Cooler	cooling fluid temp. entering device to component case temp.	0.175 °C/watt (0.315 °F/watt)[9]

4.1 Building Supply Temperature and Server Thermal Margin Estimates

The forecast for each construct is presented. Each construct has a unique but similar set of thermal approaches and temperature values.

For example: Table 3 contains the calculation detail of the construct using a cooling tower and direct liquid cooling. Note the Tcase forecast, building to IT solution hand off point and server thermal design margin. The Tcase maximum for the Intel Xeon EC5545 at 85 watts is included as 77.6°C (171.6°F) and is compared to the forecast to estimate the server thermal design margin.

Table 3: Detailed Forecast Example Using a Cooling Tower and Direct Liquid Cooling at the Chip

Heat Transfer Reference or Component	Temp. Value	Approach or Delta Temp.
Environment Ambient	26.5°C (79.9°F) Wet bulb	
Cooling Tower Approach		3.9°C (7°F)
Plate Frame Heat Exchanger Approach		1.7°C (3°F)
Building Cooling Liquid to Solution Interface Temp.	32°C (90°F)	
Pre-Heat		5°C (9°F)
Liquid Cooling Device Approach		14.8°C (26.6°F)
Tcase Forecast	51.9°C (125.4°F)	
Tcase Maximum	77.6°C (171.6°F)	
Thermal Design Margin		26°C (46°F)

The results for all four constructs including from Table 3 above are contained in the following Table 4. Some values are rounded to the nearest integer.

Table 4: Summary of Key Temperatures and Estimated Server Thermal Design Margins for Constructs Studied

Infrastructure and Chip Cooling Design	Ambient Temp.	Building Cooling Liquid to Solution Interface Temp.	Tcase Forecast Temp. 78°C (172°F) Allowed	Thermal Design Margin Temp.
Cooling Tower Direct Cooling	Wet Bulb 26.5°C (79.7°F)	32°C (90°F)	52°C (125°F)	26°C (46°F)
Dry Cooler Direct Cooling	Dry Bulb 37.5°C (99.5°F)	43°C (110°F)	65.7°C (150°F)	12°C (21°F)
Cooling Tower Air Cooling	Wet Bulb 26.5°C (79.7°F)	32°C (90°F)	63°C (146°F)	14°C (26°F)
Dry Cooler Air Cooling	Dry Bulb 37.5°C (99.5°F)	43°C (110°F)	74°C (166°F)	3.4°C (6°F)

5. CONCLUSIONS AND RECOMMENDATIONS

The study found two “natural-break” points for maximum cooling liquid supply temperatures that can be recommended as part of HPC liquid cooling design guidelines.

A building supplied liquid temperature of 32°C (90°F) was found associated with using cooling towers (evaporative cooling) as the primary cooling liquid process. This temperature can be produced using only a cooling tower 99.6% of the time at 100% of the National Laboratory HPC sites studied. Cooling liquid supplied by the building at 32°C (90°F) is estimated to provide adequate server design margin for both direct liquid cooling and air cooling at the chip. The margins are 26°C (46°F) and 14°C (26°F) for direct liquid cooling and air cooling respectively. A building supplied liquid temperature of 43°C (110°F) was found as a natural-break point associated with a cooling process that starts with a dry cooler. The estimate using the assumptions stated above indicate that a server design using air as the cooling medium at the chip may be problematic because of only a 3.4°C (6°F) margin. However by using direct liquid cooling an estimated margin of 12°C(21°F) can be obtained and should allow a thermal design operating under the Tcase maximum for 99.6% of the time for the HPC sites studied. In this study the chip power was assumed to be 85 watts.

Some server designs use power levels close to 130 watts. For these cases, assuming a similar heat sink performance, the thermal design margins will be considerably less. For example the margin is reduced by 10°C (18°F) for air cooling and 8°C (14°F) for direct liquid cooling. A reduced thermal margin for direct liquid cooling starting with a dry cooler, approaches the Tcase maximum limit, hopefully in these cases a more efficient heat transfer device can be provided, if necessary, compared to the 0.225 degrees Celsius per watt performance assumed for this study. As noted earlier other approach temperatures can be reduced providing increased thermal margin.

When ambient temperatures are above the maximum associated with the 99.6% design limit for either cooling tower or dry cooler based cooling a number of alternatives may be used. The required cooling can be reduced during these exceptional conditions by considering operational adjustments or features including: automatically reducing performance as chip temperatures reach internal limits (requires servers incorporating this feature), turning off servers to reduce heat load, spraying dry coolers, incorporating thermal storage or using a chiller for limited period. Table 5 has the final recommendations proposed for HPC maximum liquid supply temperature for liquid cooling guidelines.

This study does not include a total cost of ownership type financial comparison of current commodity servers to future HPC server solutions that may incorporate direct liquid cooling or advanced air cooling. The possible increased purchase cost or reduced density of a future HPC solution that can use warm liquid cooling should be offset by reduced capital costs for cooling infrastructure and savings from reduced ongoing energy consumption costs. These guideline recommendations have handed off to the ASHRAE liquid cooling sub-committee.

A graphical presentation of thermal component additions and thermal design margins is provided in Appendix C.

Table 5. Proposed HPC Building Supplied Cooling Liquid Maximum Temperatures

Liquid Cooling Class	Main Cooling Equipment	Supplemental Cooling Equipment	Building Supplied Cooling Liquid Maximum Temperature
L1	Cooling Tower and Chiller	Not Needed	17°C (63°F)
L2	Cooling Tower	Chiller	32°C (89°F)
L3	Dry Cooler	Spray Dry Cooler or Chiller	43°C (110°F)

6. ACKNOWLEDGMENTS

This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

7. ADDITIONAL AUTHORS

Due to formatting constraints all authors could not be listed at the top of the paper. The following people, in alphabetical order, were also part of the study effort.

Bailey, Anna-Maria, Lawrence Livermore National Laboratory
Banisadr, Farhad, Los Alamos National Laboratory
Bates, Natalie, Energy Efficient HPC Working Group
Coghlan, Susan, Argonne National Laboratory
Cowley, David E, Pacific Northwest National Laboratory
Dube, Nicolas, Hewlett Packard Corporation
Fields, Parks, Los Alamos National Laboratory
Greenberg, Steve, Lawrence Berkeley National Laboratory
Iyengar, Madhusudan, IBM Corporation
Kulesza, Peter R., Oak Ridge National Laboratory
Loncaric, Josip Dr., Los Alamos National Laboratory
McCann, Tim, SGI
Pautsch, Greg, Cray Inc.
Patterson, Michael K., Intel Corporation
Rivera, Richard G., Los Alamos National Laboratory
Rottman, Greg K, US Army Engineer Research Development Center
Dale Sartor, Lawrence Berkeley National Laboratory
Tschudi, William, Lawrence Berkeley National Laboratory
Vinson, Wade, Hewlett Packard Corporation
Wescott, Ralph, Pacific Northwest National Laboratory

8. REFERENCES

The following references , assumptions or explanations refer to the information presented in Table 2. Due to the proprietary and nascent nature of some cooling performance attributes listed, the thermal approach design values may not be easily found or confirmed in the public domain. We appreciate the candor of companies that provided information. The actual performance of a system may contain approach values considerably above or below the values used. The values used are considered to be achievable using good design practice.

(1) 1996 ASHRAE Systems and Equipment Handbook page 36.2

(2) www.drycoolers.com

(3) Taylor Engineering, Plate Frame BGP47-90-TMTL4 – 8/23/2010 LBNL ALS USB Server Rm. Drawing M0.2

(4) Additional approach added to plate frame heat exchanger due to smaller space likely for a liquid to liquid heat exchanger contained in a CDU.

(5) Coolcentric (formally Vette) air to water passive rear door heat exchanger performance.

(6) Discussion with Electronic Cooling Solutions (B. Maltz) – assumes no upstream high heat components and well designed internal recirculation control.

(7) Discussion with IBM and Cray – conservative estimates averaged

(8) Discussions with IBM, SGI and Cray – conservative estimates averaged

(9) Discussions with IBM and Cray – conservative estimates averaged

9. APPENDIX A: NATIONAL LABORATORY ENVIRONMENTAL DATA – TABLE FORMAT

Locations sorted by Dry Bulb Temperature

HPC Data Center Owner	U.S. State	City Name	Close Location Available in ASHRAE Database	Dry Cooling Dry Bulb 0.4% Design Temperature		Evap. Cooling Wet Bulb 0.4% Design Temperature	
				°F	°C	°F	°C
Pacific Northwest National Laboratory	Washington	Richland	Pasco	**99.5**	**37.5**	72.1	22.3
Lawrence Livermore National Laboratory	California	Livermore	Livermore Municipal Apt	**98.8**	**37.1**	70.8	21.6
Houston Texas	Texas	Houston	Bush Intl. Apt.	**96.8**	**36.0**	80.1	26.7
Los Alamos National Laboratory	New Mexico	Los Alamos	Albuquerque Intl Apt	**95.2**	**35.1**	65.3	18.5
Sandia National Laboratory	New Mexico	Albuquerque	Albuquerque Intl Apt	**95.2**	**35.1**	65.3	18.5
Jefferson Laboratory	Virgina	Newport News	New Port News	**94.5**	**34.7**	79.7	26.5
Oak Ridge National Laboratory	Tennessee	Oak Ridge	Nashville Intl Apt	**94.4**	**34.7**	78.2	25.7
National Renewable Energy Laboratory	Colorado	Golden	Denver Stapleton Intl Apt	**93.5**	**34.2**	64.4	18.0
Princeton Plasma Phisics Laboratory	Princeton	New Jersey	Mcguire AFB	**92.9**	**33.8**	78.8	26.0
Stanford Linear Accelerator Center	California	Menlo Park	San Jose Intl Apt	**92.3**	**33.5**	69.5	20.8
Argonne National Laboratory	Illinois	Argonne	Chicago Midway Apt	**92.1**	**33.4**	78.0	25.6
Idaho National Laboratory	Idaho Falls	Idaho	Fanning Field Apt	**91.7**	**33.2**	64.9	18.3
Fermilab	Illinois	Batavia	Aurora Municipal Apt	**90.8**	**32.7**	77.7	25.4
Ames National Laboratory	Iowa	Ames	Ames Muni Apt	**90.5**	**32.5**	79.2	26.2
Brookhaven National Laboratory	New York	Upton	Long Island Macarthur Apt	**88.4**	**31.3**	76.7	24.8
Lawrence Berkeley National Laboratory	California	Berkeley	Oakland	**81.8**	**27.7**	67.6	19.8

Locations Sorted by Wet Bulb Temperature

HPC Data Center Owner	U.S. State	City Name	Close Location Available in ASHRAE Database	Dry Cooling Dry Bulb 0.4% Design Temperature		Evap. Cooling Wet Bulb 0.4% Design Temperature	
				°F	°C	°F	°C
Houston Texas	Texas	Houston	Bush Intl. Apt.	96.8	36.0	**80.1**	**26.7**
Jefferson Laboratory	Virgina	Newport News	New Port News	94.5	34.7	**79.7**	**26.5**
Ames National Laboratory	Iowa	Ames	Ames Muni Apt	90.5	32.5	**79.2**	**26.2**
Princeton Plasma Phisics Laboratory	Princeton	New Jersey	Mcguire AFB	92.9	33.8	**78.8**	**26.0**
Oak Ridge National Laboratory	Tennessee	Oak Ridge	Nashville Intl Apt	94.4	34.7	**78.2**	**25.7**
Argonne National Laboratory	Illinois	Argonne	Chicago Midway Apt	92.1	33.4	**78.0**	**25.6**
Fermilab	Illinois	Batavia	Aurora Municipal Apt	90.8	32.7	**77.7**	**25.4**
Brookhaven National Laboratory	New York	Upton	Long Island Macarthur Apt	88.4	31.3	**76.7**	**24.8**
Pacific Northwest National Laboratory	Washington	Richland	Pasco	99.5	37.5	**72.1**	**22.3**
Lawrence Livermore National Laboratory	California	Livermore	Livermore Municipal Apt	98.8	37.1	**70.8**	**21.6**
Stanford Linear Accelerator Center	California	Menlo Park	San Jose Intl Apt	92.3	33.5	**69.5**	**20.8**
Lawrence Berkeley National Laboratory	California	Berkeley	Oakland	81.8	27.7	**67.6**	**19.8**
Los Alamos National Laboratory	New Mexico	Los Alamos	Albuquerque Intl Apt	95.2	35.1	**65.3**	**18.5**
Sandia National Laboratory	New Mexico	Albuquerque	Albuquerque Intl Apt	95.2	35.1	**65.3**	**18.5**
Idaho National Laboratory	Idaho Falls	Idaho	Fanning Field Apt	91.7	33.2	**64.9**	**18.3**
National Renewable Energy Laboratory	Colorado	Golden	Denver Stapleton Intl Apt	93.5	34.2	**64.4**	**18.0**

10. Appendix B: National Laboratory Environmental Data – Chart Format

Locations sorted by Dry Bulb Temperature

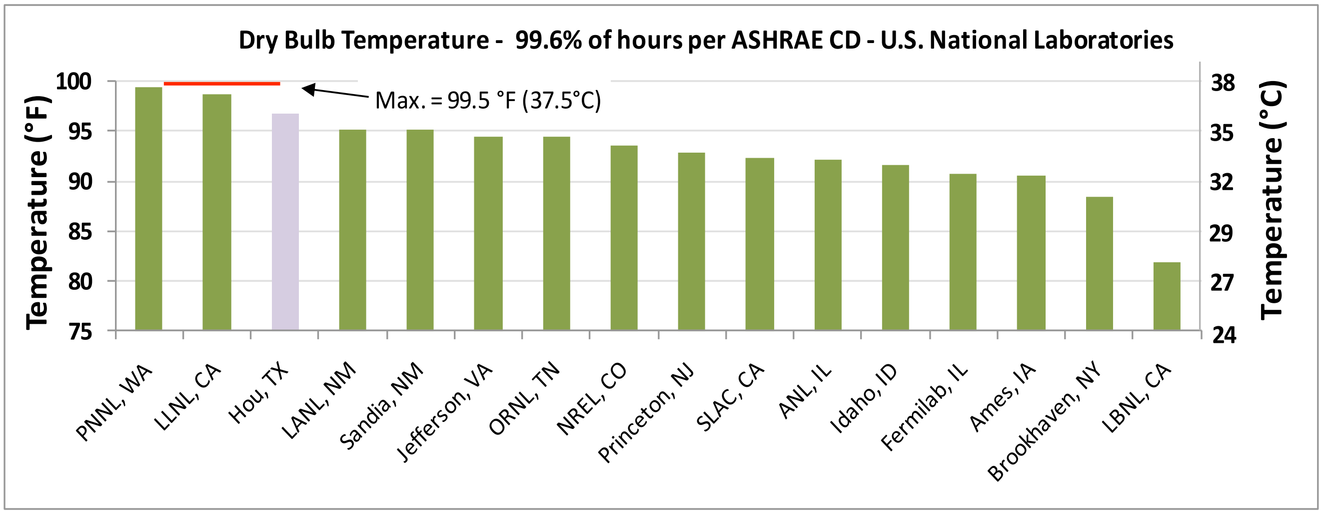

Locations sorted by Wet Bulb Temperature

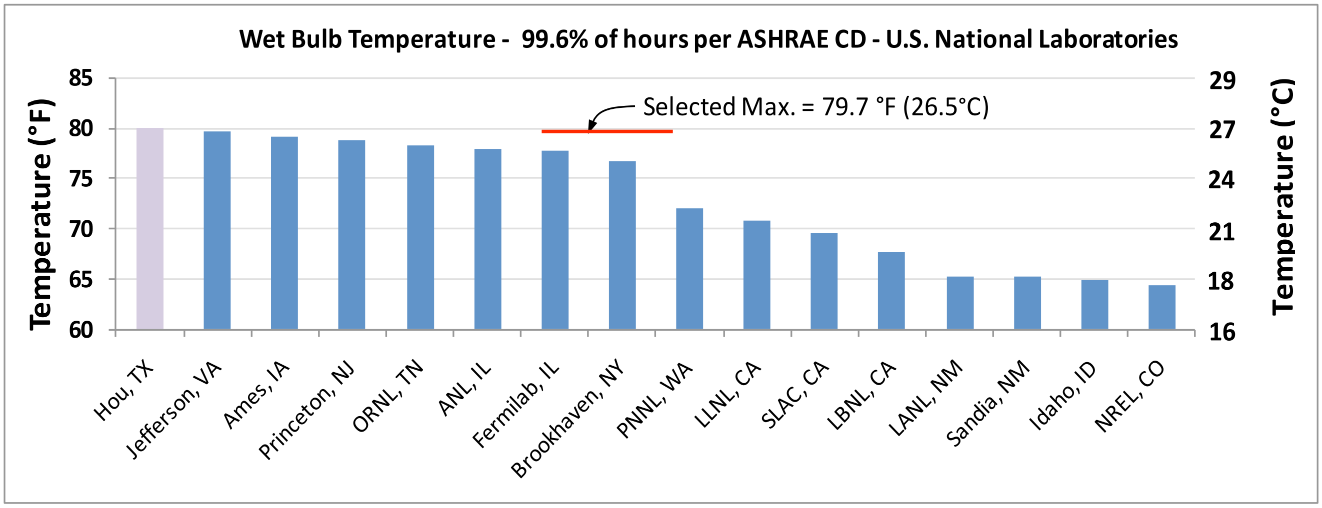

11. APPENDIX C: DIRECT LIQUID AND AIR COOLING COMPONENT TCASE FORECAST GRAPHICS

Direct Liquid Cooling Thermal Components
and Resulting Thermal Design Margins

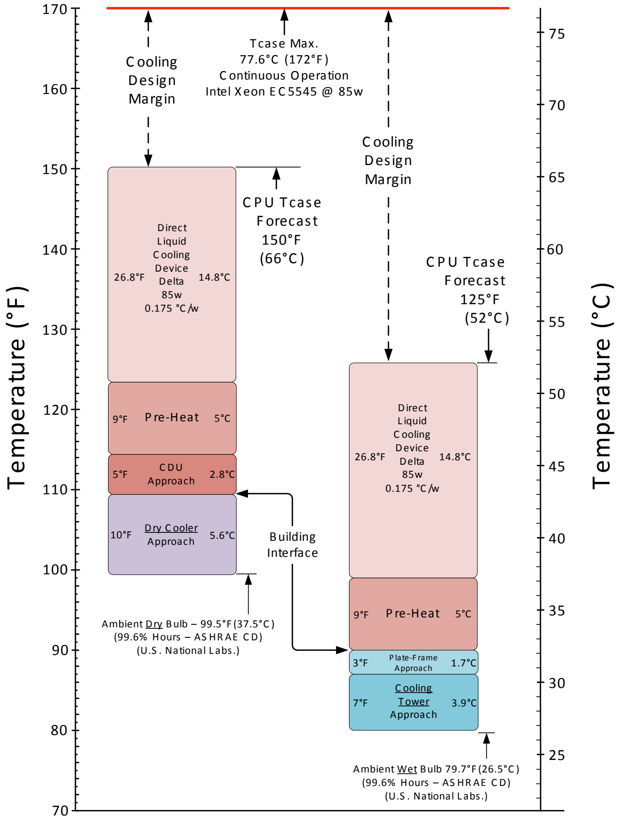

Air Cooling Thermal Components
and Resulting Thermal Design Margins