An Evolution in Liquid Cooling: LBNL 50B-1275 as a Case Study

Steve Greenberg

Executive Summary

Building 50B, Room 1275 at the Lawrence Berkeley National Laboratory is a 5600-square foot, raised-floor data center that was rebuilt in the mid-1990s to house the National Energy Research Scientific Computing center (NERSC). NERSC outgrew the room in space, power, and cooling needs and moved first to a leased building in nearby Oakland in 2000, and then back to its own building on the main LBNL site. 1275 became an enterprise data center for in-house needs (email, accounting, etc.) and scientific computing. Email and other business applications were moved to the cloud or other spaces, and from 2010 the center began hosting clusters of HPC (high-performance computing) clusters owned by scientific groups at LBNL. The original cooling scheme used CRACs (Computer-Room Air Conditioners) that discharged their cold air into the underfloor plenum and rejected their heat to water cooled by the building's cooling tower. Poor air management meant that there was an apparent shortage of cooling even though the CRACs had significant excess rated capacity relative to the cooling load.

Improved air management got the cooling needs and capacity in balance, and as load grew, the center began testing alternative cooling schemes to meet load and improve energy efficiency. Passive rear doors, rejecting to the cooling tower water and/or chilled water were used as early as 2008. Bringing liquid cooling to the chip was demonstrated in 2014.

The current evolution seeks to provide as much HPC hosting as possible using the available power and meeting the cooling needs as efficiently as possible. This includes moving batteries and spinning disks (with tighter temperature requirements) to another space, operating 1275 as an ASHRAE A3 area (64°-81°F per the Recommended temperature range, but allowable up to 95°F IT inlet), using active (fan-assisted) rear doors on tower water only and with higher cooling water temperature difference to allow more cooling with the same water flow, direct-to-chip water cooling, and removing the worn-out, inefficient CRACs. The ultimate build-out is projected to be 1.5 MW total power and the PUE decreasing from today's 1.4 to 1.1.

Lessons learned from the 1275 evolution and discussed further in this case study:

- Liquid cooling is effective and reliable
- Rear doors are a good bridge technology from air to warm liquid cooling (compressor-free)
- Distribution of heat within passive rear doors is important
- There are tradeoffs using control valves and/or flow balancing for liquid cooling
- Active rear doors can unload server fans, making PUE worse but resulting in net energy savings.

Introduction and background

1990's: Supercomputer Center in Transition

Building 50B, Room 1275 is a 5600-square foot, raised-floor data center that was rebuilt in the mid-1990s to house NERSC, the National Energy Research Scientific Computing center. NERSC outgrew the room in space, power, and cooling needs and moved first to a leased building in nearby Oakland in 2000, and then back to its own building on the main LBNL site. Room 1275 became an enterprise data center for in-house needs (email, accounting, etc.) and scientific computing used by LBNL researchers. The original cooling scheme used CRACs (Computer-Room Air Conditioners) that used vapor-compression refrigeration cycles and discharged their cold air into the underfloor plenum and rejected their heat to water cooled by the building's cooling tower. Poor air management (see Figure 1) meant that there was an apparent shortage of cooling even though the CRACs had significant excess rated capacity relative to the cooling load imposed by the IT and power distribution equipment in the space.

Figure 1. Air management issues in the early days. At left, portable fans were used to address hot spots; note CRAC with inlet from the space rather than above-ceiling plenum. At right, floor tiles with large openings addressed inadequate flow in this aisle but dropped the under-floor pressure, causing air deficiencies elsewhere.

2000's: Science-and-House Computing

Improved air management got the cooling needs and capacity in balance (Bell, 2010); see Figure 2.

Figure 2. Improved air management. At left, CRAC unit with return duct from above-ceiling plenum. Today all but one of the seven original CRACs have been removed. At right, floor tiles with appropriate opening size to properly balance airflow.

As load grew, the center began testing alternative cooling schemes to meet load and improve energy efficiency. For example, water-cooled rear doors on 6 server racks were installed (Mathew et al, 2010). These doors were "passive", i.e. the server fans provided the necessary pressure to force air through the heat exchangers in the doors. See Figure 3. This first set of doors used closed-loop cooling water cooled in a heat exchanger by the cooling tower water. A booster heat exchanger could cool the water further (rejecting to building chilled water) but was not used until a recent retro-commissioning effort. Because they used the relatively warm closed-loop cooling water, these doors removed only about half of the server heat; the remaining heat was removed by the CRACs. Following this successful demonstration, additional passive doors were added, with no booster heat exchanger; again, the CRACs removed residual heat.

Figure 3. Passive rear-door heat exchanger (RDHX). At left, RDHX with door opened 90°. Hoses (at right) supply and return water from an under-floor distribution system.

2010's onward: HPC Clusters Owned by Scientific Groups

In this period, email and most other business applications were moved to the cloud or other spaces, and in the 2010s the center began hosting HPC (high-performance computing) clusters owned by scientific groups at LBNL. Density continued to increase. An IBM idataPlex cluster from the NERSC site with passive rear doors (Coles and Greenberg, 2012a) was moved to the center. These rear doors were operated on chilled water, which allowed them to not only remove all the heat from the IBM cluster, but provide net cooling to the data center. There was some condensation due to the low water temperature used, but due to the modest levels of humidity in Berkeley, the dripping was readily evaporated by the high under-floor air flows and did not cause operational problems.

In addition to the above water cooling of production IT equipment, Room 1275 hosted additional liquid-cooling tests. For example, an APC prototype in-rack cooler with dual coils (chilled water and closed-loop cooling water) was demonstrated (Coles and Greenberg, 2012b), as was an Asetek direct-to-chip water cooling system (Coles and Greenberg, 2014). There was also an HP cooling system with an in-row cooler in between two equipment cabinets, all enclosed and room neutral in normal operation, which used chilled water.

Current implementation: maximizing cooling value

Current practice in the data center is incrementally increasing load (and density) as more researchers locate their HPC equipment in the space. In 2018, a cooling master plan was developed to guide the process, driven by the mandate to serve the growing load given the power and cooling constraints of the existing infrastructure, especially the main transformers and cooling tower. Key elements of the plan include:

1. Moving UPS batteries and spinning disks, with relatively tight temperature requirements, to another space.
2. Enabled by a., the room will be operated as an ASHRAE A3 (air-cooled, level 3) area. This means that the *Recommended* temperature range is still the same as A1-A4, i.e. 64.4° to 80.6°F, but the *Allowable* IT inlet temperature can range up to 95°F (ASHRAE, 2015). This temperature range enables compressor-free cooling year-round in the Berkeley climate. It also requires that IT equipment purchased for use in the space is specified for A3 allowable conditions.
3. In order to achieve the Recommended IT inlet temperature nearly all of the year with closed-loop cooling water, active rear doors will be used. Active doors have larger heat exchangers than passive doors in order to get a closer approach temperature (the difference between the entering cooling water temperature and the leaving air temperature), and a set of fans to provide the necessary air flow through the exchangers. See Figure 4.

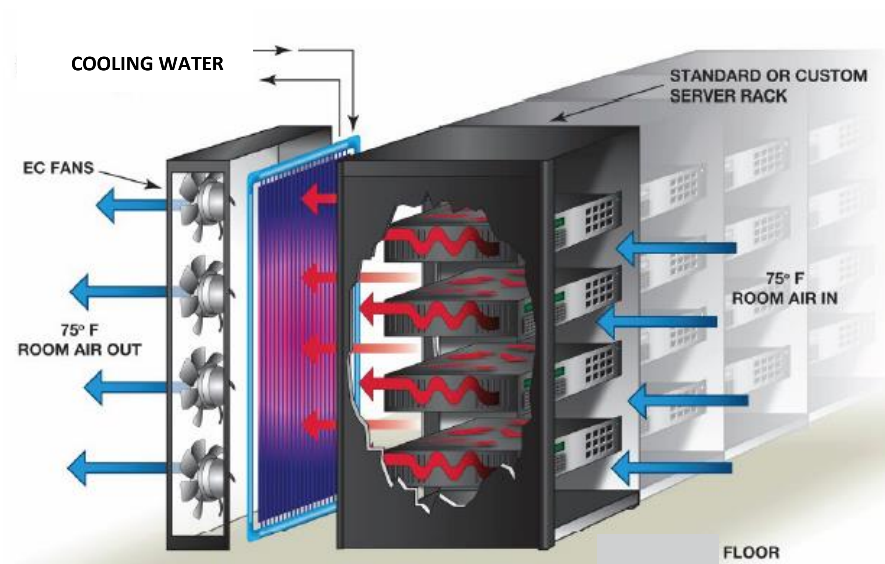

Figure 4. Example of active rear door heat exchanger. The air temperatures at Room 1275 will vary from those shown. Courtesy MotiveAir.

4. Using the larger heat exchangers in the active rear doors to realize higher temperature rise between the entering and leaving cooling water. This enables more cooling for the same water flow.
5. As density climbs further, utilize direct-to-chip water cooling for more effective and efficient heat removal. This cooling system is now available directly from server manufacturers.
6. As the above are phased in, remove six of the seven CRACs, freeing up space in the computer room as well as cooling capacity in the closed-loop water system (since much less fan and compressor heat will be added to the water) and power capacity (since much less power will be used by the doors than on the fans and compressors in the CRAC units).
7. The ultimate build-out of the space will use all of the 1.5 MW total available capacity of the two main transformers serving the center.
8. Additional efficiency and operational improvements include addition variable-frequency drives (VFDs) to the tower water and closed-loop cooling water pumps, adding a second closed loop cooling water-to-tower water heat exchanger to allow maintenance without shutdown, and adding a water filtration system to meet rear-door manufacturer specifications.
9. The PUE is projected to decrease to about 1.1 from its current value of about 1.4, resulting in savings of about 400kW or 3.5 MWh/yr.

Summary and Lessons Learned

1. Liquid cooling is effective and reliable. Various liquid-cooling technologies have been in use at the 50B-1275 data center since 2008, some operating continuously through the whole period.
2. The more closely coupled the cooling is to the heat source, the more efficient the cooling is and the warmer the cooling medium can be (Bell, 2012). As such, rear door heat exchangers, because they effectively eliminate bypass and recirculation air outside of the rack, are a good intermediate technology in the range from cold air to warm liquid cooling (compressor-free).
3. While rear doors are an effective means for efficient IT cooling, their performance is reduced if the IT loads are too concentrated in the rack. This challenge was identified at an Infosys installation.
4. Control strategies for the cooling water to rear doors can range from manually setting a throttling valve, to using calibrated flow balancing valves to set specific design flows, to automatic control valves with control based on discharge air or water temperature or on rack load. The automated scheme makes best use of the cooling water resource as loads vary, but the extra sophistication requires additional capital and maintenance expense.
5. There are interactions between rear doors and the IT equipment fans. Passive doors force the fans to work slightly harder to maintain cooling airflow against the back-

pressure imposed by the heat exchanger. The fans of active doors, because they boost the airflow through the exchanger, unload the IT fans. Since IT fans are counted in PUE denominator and the active door fans are only counted in the numerator, PUE will increase with active doors. Since the fans in the rear doors are more efficient than the IT fans, there should be a net decrease in energy usage.

6. The Center of Expertise for Energy Efficiency in Data Centers at LBNL (Center of Expertise, 2020) is tasked by FEMP with assisting federal data center operators with their energy efficiency efforts. The various consultations, demonstration projects, and case studies noted here, as well as others in the center (e.g. data center infrastructure management, variable-speed CRAC fans, server energy use) over the last 2 decades are all examples of this technical assistance.
7. Data centers always have room for improvement. With focus on the space, power, and cooling constraints of each center, in combination with the computing needs of the center's users, an optimal utilization of resources can be achieved.

Resources and Annotated References

ASHRAE, *Thermal Guidelines for Data Processing Environments*, 4th Edition, 2015. American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Atlanta.

Bell, G.C., 2010. "Data Center Airflow Management Retrofit". As data center energy densities increase, energy savings for cooling can be realized by optimizing airflow pathways within the data center. https://datacenters.lbl.gov/sites/default/files/airflow-doe-femp.pdf

Bell, G.C., 2012. "Improving Data Center Efficiency with Rack or Row Cooling Devices". In new data center designs, capacity provisioning for ever-higher power densities puts into question whether conventional room-conditioning systems can manage future information technology (IT) loads. https://datacenters.lbl.gov/sites/default/files/dc_chilloff2.pdf

Center of Expertise for Energy Efficiency in Data Centers, 2020. A Department of Energy facility located at the Lawrence Berkeley National Laboratory offering technical support, tools, best practices, analyses and technologies to help federal government agencies and other organizations implement data center energy-efficiency projects. https://datacenters.lbl.gov/

Coles. H.C. and S. Greenberg, 2012a. "Energy Efficiency Investigation for the Magellan Super Computer at the National Energy Research Scientific Computing Center" The Magellan supercomputer consists of an IBM iDataPlex system with 400 compute nodes. This system is cooled with passive rear doors. https://datacenters.lbl.gov/resources/energy-efficiency-investigation-magellan-super-computer-national-energy-research

Coles, H.C, and S. Greenberg, 2012b. "Data Center Economizer Cooling with Tower Water". A prototype computer equipment rack-level cooling device with two heat exchangers was demonstrated to illustrate an energy efficient cooling capability. https://datacenters.lbl.gov/sites/default/files/PIER_Economizer%20Cooling%20Demo_2012.pdf

Coles, H.C., and S. Greenberg, 2014. "Direct Liquid Cooling For Electronic Equipment". Report on the demonstration of direct liquid cooling for electronic equipment. Cisco C200 M3 servers were retrofitted with the Asetek direct cooling technology. The energy-savings potential for this technology was quantified using measurements and data center energy use models. https://datacenters.lbl.gov/resources/direct-liquid-cooling-electronic-equipment

Mathew, P., W.F. Tschudi, and S. Greenberg, 2010. "Data Center Rack Cooling with Rear-door Heat Exchanger". Liquid-cooling devices were installed on server racks in a data center at Lawrence Berkeley National Laboratory (LBNL). The passive-technology device removes heat generated by the servers from the airflow leaving the server rack. https://datacenters.lbl.gov/sites/default/files/rdhx-doe-femp.pdf

1. Contacts

Primary Author:

Steve Greenberg
Lawrence Berkeley National Laboratory
segreenberg@lbl.gov
510-486-6971

For more information about FEMP:

Rachel Shepherd
Rachel.Shepherd@ee.doe.gov