

TITLE:
CHALLENGES WITH LIQUID COOLING

ABSTRACT:

Liquid cooling is key to dealing with the heat density, reducing energy consumption, and increasing the performance of this generation of supercomputers and becomes even more predominant on the roadmap for the foreseeable future. The transition to liquid cooling, however, comes with challenges.

One challenge is that each system and each site comes with its specific issues regarding liquid cooling. These complicate procurements, design, installation, operations, and maintenance.

This panel will review the immediate past history from a lessons learned perspective as well as discuss what's needed for liquid cooling to be implemented more readily in the future. This includes challenges such as data and metrics to assess the relative efficiencies of the different cooling technologies, water quality, heat recovery, the disparity between building life and cluster life, and other issues.

How can we as a community address these challenges?

MODERATORS:

Torsten Wilde, LRZ
TBD – Josip Loncaric, LANL

PANELISTS:

Thomas Blum, Megware
Michael Patterson, Intel
Nic Dube, HP
Laurent Cargemel, Bull
Giovannibattista Mattiussi, Eurotech
TBD – Ingmar Meijer, IBM

Format: Moderators are from key supercomputing centers in Europe and the United States. Panelists are representatives from key vendors, also from Europe, the United States and Canada.

The panel will start with the moderators giving a brief presentation on the issues and challenges that they face at their data centers. This will set the stage. These moderators will then promote dialogue and discussion about potential solutions to these issues and challenges.

Below is an example of a common 'liquid cooling' issue.

Background:

Each platform comes with its specifics regarding liquid cooling, which have to be addressed at installation. When the platform is decommissioned, the facility water systems stay, but platform-specific piping is gone.

Facilities support multiple platforms at a time and multiple generations of platforms over time. Facility water systems therefore need to provide enough flexibility to do this.

We need to separate facility water systems, which are stable but flexible; from platform water systems, which are specific and transient.

One way to make progress is to realize that platform vendors have already dealt with this problem: They need to pre-ship test equipment they sell, and that means connecting liquid cooling. Normally, they do this in sensible chunks, e.g. a rack or multiple racks with their CDU. So, if the facility is built using commonly used piping sizes, it should not be too hard to design and build platform water connections to those headers.

So, here is an idea: Ask platform vendors how they supply water to platforms in pre-ship testing. Discuss, and select some pipe sizes that seem to be common. Build facilities with these kinds of headers; and ask platform vendors to design piping to connect to them -- which should not be too hard because they have to do this anyway at their factory before platform delivery.

One example: Design facility to provide N connections of 4-8" each (some common diameter in that range). Design platform 1 to connect to K1 such headers, platform 2 to connect to K2 headers, etc. Pressures, water flows, temperatures, dynamic responses have to be considered; but one has to chunk up the aggregate flow to each platform anyway, at the CDU or rack level -- so that could be done without too much disruption if we follow common industry practices.