Electronic Design

RUPESH MAINALI, Senior Member Technical Staff,
Reliability Engine, www.reliabilityengine.com

The Hidden Signals Inside a Liquid-Cooling System

Stable component temperatures could conceal a liquid-cooling system about to fail. But closely monitoring the coolant can reveal clogged filters and shrinking pump reserves earlier on.

Take two snapshots of the same liquid-cooled server rack, and the thermal conditions can look identical. Rack power, coolant supply temperature, and processor temperature are nearly unchanged.

But beneath these stable operating conditions, the cooling loop may be struggling to keep up with the heat. In the later snapshot, the coolant distribution unit (CDU) drives the pumps at a higher speed while the filter consumes more of the available pressure budget, and the flow of coolant in one branch has started to slow down compared to its peers. The processors are still being cooled to the temperature required for the data center, but the cooling loop is working harder to prevent overheating.

By opening valves and increasing pump speed in real-time, a liquid-cooling system can conceal many of these issues. However, these subtle changes can be early indicators of increasing flow resistance, imbalances, sensor degradation, or other forms of deterioration that increase the risk of eventual overheating and failure.

These risks may remain hidden if performance is judged by the temperature and load alone. As a result, the control system needs to monitor not only temperature, but also the effort required to achieve it.

Coolant data becomes a reliability input when it changes a decision. High component temperature, loss of flow, and other signals are used by fast protection logic to reduce performance impacts or limit damage from coolant leaks and other failures. Particle counts, inhibitor concentration, and dissolved-metal results develop more slowly and belong in maintenance planning. Put them on the same asset timeline, but not in the same alarm window.

How to Spot Early Signs of Liquid-Cooling Failure

Coolant conditions ebb and flow as the processors and other electronics in the rack throttle up and down in performance. But not every change is a sign of a potential failure. To tell the difference between what's normal and what's not, engineers need to ask themselves some fundamental questions:

How much heat is the liquid carrying?

Compare rack electrical power with the heat carried by the coolant, estimated from density, volumetric flow, specific heat, and the return-to-supply temperature rise. The values will not match exactly because of sensor uncertainty, heat rejected to air, and differences in measurement boundaries. Track the gap from the commissioned reference. A sudden shift can point to a biased flow sensor, reversed temperature inputs, an undocumented air-cooling contribution, or the wrong power boundary.

How hard is the loop working?

View flow beside differential pressure, pump command or speed, valve position, and filter differential pressure. More pump speed at higher rack power may be normal. More pump speed at the same flow, valve state, and coolant temperature means the path now needs more pressure. Look for filter loading, a valve or quick-disconnect problem, trapped gas, a branch imbalance, or a bad measurement.

Is the fluid still within its intended condition?

Start with the specified fluid, treatment plan, and wetted material list. Conductivity, pH, inhibitor or glycol concentration, particles, turbidity, microbiological activity, dissolved oxygen, and dissolved metals answer different questions. Choose tests tied to a known failure mechanism and a practical response. No single set of limits fits every fluid and material system.

What changed around the same time?

A fill, top-off, filter change, hose replacement, cold-plate service, firmware update, control-sequence change, or workload migration can explain a step in the data. Put service and operating events on the same timeline before deciding that the change is equipment degradation.

The Importance of Sensor Placement in a Liquid Cooling System

Liquid-cooling systems use a network of embedded sensors to track pressure, temperature, flow, fluid levels, and other parameters in real-time. As a result, strategic sensor placement is critical to accurate monitoring and early problem detection.

- **At the CDU:** Capture the temperature of the coolant flowing out to the electronics and the coolant flowing back from the hot spots in the system, also known as supply and return temperatures of the technology cooling system (TCS).
  - Total flow and pump command or speed must also be monitored. Record the facility-side valve or heat-exchanger state that controls TCS supply temperature. For condition-based filter service, measure differential pressure within the CDU. When available, engineers can track pump electrical power, too, which provides an independent check on control effort.
- **At the rack boundary:** Capture flow and differential pressure between the supply and return headers. If parallel branches share a manifold, branch or tray flow tells more than rack total flow alone. The total can remain steady while a restriction, valve change, or balancing error moves coolant away from one path and into another.
- **At the electronics boundary:** Retain processor or package temperature, reported power, throttling or power-capping state, and enough workload context for a fair comparison. Rack averages can hide the hottest branch, a degraded interface, or the component with the least thermal margin.
- **At the sample point:** Choose a location that's representative, accessible, and repeatable. Record the loop, port, fluid temperature, collection method, and time. A sample drawn from a stagnant leg or an unflushed valve may look credible while saying little about the circulating fluid.

The Berkeley Lab liquid-cooling controls guidance organizes liquid-cooling inputs by measurement frequency, accuracy, and ownership. The framework also works for reliability monitoring. Write down the decision each signal supports. Control inputs, historian trends, and laboratory results need different collection methods, but they still require a shared time base and asset identifier.

Build a Baseline to Monitor the Performance of a Liquid-Cooling System

Commissioning should leave operations with more than a pass-or-fail test sheet. Capture a reference after the loop has been cleaned, filled, vented, filtered, balanced, and verified. Record fluid identity and lot, concentration where applicable, wetted-material configuration, and filter type. Include sensor calibration, control mode, rack power, temperatures, flow, pressure loss, pump command, and the values required by the treatment plan.

Once the loop is stable, capture repeatable low, typical, and high load points, including valve and pump control modes. These form a reference curve for checks after a filter change, rack expansion, hardware refresh, or controls update. If the configuration changes materially, create a new dated baseline. An obsolete reference may send a diagnosis in the wrong direction.

The Open Compute Project guideline for water-based transfer fluids calls for a representative sample from each TCS loop at startup and for routine testing against the treatment plan. Keep those analytical records with the operating baseline. A result can remain within specification while its rate of change, or its disagreement with the hydraulic data, still warrants a closer look.

Comparing Liquid-Cooling Performance on Equal Terms

Cooling trends are easy to misread. Pressure loss rises with flow. Viscosity changes with fluid temperature and concentration. Conductivity changes with temperature. Package temperature responds to power, firmware, workload placement, coolant supply temperature, and interface condition. Before labeling a trend as degradation, compare similar operating points.

For a fixed path in turbulent flow, differential pressure divided by flow squared can serve as a hydraulic index. It's a comparison tool, not a universal resistance coefficient. Valve position, fluid properties, geometry, sensor calibration, and flow regime must stay comparable. Within those limits, the index helps separate a change in the path from a routine change in flow demand.

Another useful comparison is the component temperature rise above coolant supply divided by component power. Trend this apparent thermal-resistance proxy only at steady, comparable operating points. It's not a qualification metric. Power reporting, sensor location, thermal-interface behavior, and heat shared with air introduce error. Its value appears when nominally identical branches or trays begin to separate.

Keep the time scales straight. Electronics and control telemetry can show the onset of a change. A slower fluid analysis may explain why. Tie each result to its interval, service event, and equipment state. A laboratory sample describes a place and period, not a second-by-second condition.

Liquid-Cooling Trends That Signal a Problem

1. The filter is consuming more of the pressure budget in the system. At comparable flow and coolant temperature, filter differential pressure rises while the controller increases pump command to hold rack flow. Component temperatures remain normal. Verify the pressure taps and flow measurement, then check the service history. If the evidence agrees, inspect or replace the filter under the approved procedure before the pump reaches its control limit.

Don't interrupt the rack at the first indication, but don't wait for temperature to confirm a restriction. A filter alarm must account for flow or it will produce false alerts after ordinary load changes.

2. The coolant flow in one branch diverges from the others around it. Rack total flow looks acceptable, but one branch loses flow and its component-to-coolant temperature rise widens at comparable power. Check branch valves, quick disconnects, trapped gas, sensor agreement, and local pressure data before blaming the cold plate. If hydraulic and thermal evidence move together, a localized restriction becomes the leading explanation.

The OCP cold-plate development guidance treats thermal performance, pressure drop, materials, and reliability as connected qualification concerns for good reason.

3. Fluid condition changes after service. Conductivity, concentration, pH, or particle results shift after a top-off or component replacement, while temperatures and hydraulics remain steady. Confirm the fluid identity, quantity added, sample location, temperature compensation, instrument calibration, and handling method. Then collect a confirmation sample using the defined procedure. This is a reason for changing the investigation plan, not the protection limits.

A Three-Level Alarm Strategy for Liquid-Cooling Systems

An alarm earns its place only when the receiving engineer knows what it may mean, how to confirm it, and which action is authorized. A three-level response keeps early warnings separate from conditions that demand immediate protection.

An advisory indicates that a normalized trend has moved outside its established envelope for a defined persistence period. The first steps are to check data quality, operating mode, and recent events. An advisory should not interrupt compute.

An investigation is warranted when an independent measurement or a second signal tied to the same mechanism supports the trend. The response may include a controlled sample, filter inspection, sensor check, branch inspection, or planned maintenance window.

Protection is required when a validated condition can threaten hardware faster than an engineer is able to investigate it. Confirmed critical loss of flow, a confirmed leak, or a thermal limit may justify derating or shutdown. The safe state and reset conditions should be designed and tested before an event occurs.

For each level, name the owner, confirmation method, response time, and closure record. Set persistence periods and limits from the equipment design, approved treatment plan, commissioning data, and consequence analysis. Copying one threshold across unlike loops may simplify configuration, but it weakens the engineering basis for the alarm.

Questions to Resolve Before Liquid-Cooling Design Release

- **Measurement boundary:** Can the team reconcile rack electrical power with fluid-side heat removal closely enough to detect sensor or boundary errors?
- **Hydraulic visibility:** Can a rack-total measurement distinguish a CDU restriction from a branch-level problem, or is more instrumentation required?
- **Clean baseline:** Will commissioning deliver repeatable operating points, retained fluid records, calibration evidence, and known control states?
- **Data context:** Are workload, firmware, valve state, top-offs, filter changes, and component service written to the same timeline as the cooling data?
- **Alarm intent:** Does every threshold name a failure mechanism, an independent confirmation, an owner, and an action?
- **Protection boundary:** Which conditions justify immediate automated action, and which require confirmation before maintenance is scheduled?

The current ASHRAE AI data center energy and thermal framework addresses variable-speed pumping, differential-pressure and flow control, redundant sensing, water-quality management, and continuous commissioning within the same TCS framework. That combination matters. When the cooling path can limit electronics availability, fluid-side instrumentation and commissioning become part of the reliability case for the system.

Look at More Than Temperature to Improve Liquid-Cooling Reliability

Temperature remains a primary protection variable, but it can't describe the entire cooling path. Ask whether the same temperature now requires more flow, more pressure, more control effort, or a different fluid condition. That comparison turns coolant telemetry into useful engineering evidence without treating every measurement as a fast control input.

In high-density electronics, the least disruptive maintenance window often opens before there's movement in temperature. Clear measurement boundaries, a reproducible baseline, comparable operating points, and alarms tied to specific decisions make that window visible while the team still has room to act.

Rupesh Mainali is a Senior Member of Technical Staff (SMTS) at Reliability Engine, where he works on reliability for liquid-cooled AI infrastructure. His technical writing focuses on coolant condition, thermal operations, and the engineering evidence needed to keep high-density systems dependable.