Cooling IBM Supercomputers

Birds of a Feather – Dynamic Liquid Cooling, Telemetry, and Controls; Opportunity for Improved TCO?

Chris Marroquin
IBM Systems Packaging Design Engineer
11/17/2015

Recent History:
Blue Gene/Q, SuperMUC, POWER 775

- **All machines used direct-water-cooling of above-dew point water for a majority of the compute rack power**

- **Constant temperature and flow inlet water supply provided by a coolant distribution unit (CDU)**
  - Filtered to not plug quick-connects and fine pitch fins
  - Treated with biocides and corrosion inhibitors

- **POWER 775 had integrated in-rack CDUs, SuperMUC at LRZ had custom 1 MW CDUs, Blue Gene/Q used both commercial rack-sized CDUs and facility level CDUs**
  - For Blue Gene/Q, created fast-acting shutoff valve to protect against leaks. See next slide.

- **All machines had a (correlated) choice of inlet flow rate and inlet water temperature**
  - Some clients would seasonally change inlet water temperature (and flow) to stay within envelope

- **POWER 775 and Blue Gene/Q could measure water temperature and flow.**
  - Periodically stored in a database along with device temperatures
  - Accuracy depended on measurement method

- **IBM HPC offerings trending toward allowing all versions of cooling, but large systems are likely to use facility level CDUs**

Facility Considerations

- **In general, IBM does not recommend fine grained regulation of temperature or flow**
  - Different racks on the facility loop could be experiencing different workloads
  - Individual drawers within the rack could be experiencing different workloads
    - Water flow/temp to each rack must be sized to handle max expected workload
  - Drawer level feedback is possible, but requires feedback and control loops for each drawer to the facility
  - Reducing facility pump flow does not offer the savings in pumping power to justify increased cost of the infrastructure of valves, sensors and feedback loops at the rack or system level

- **IBM recommends following an operating envelope to allow each customer to operate at their most optimal point**
  - Facility is free to change flow rate or temperature to stay within envelope
  - Additionally, inlet water temperatures can always be reduced as dew-point falls
  - To maximize “free cooling”, customer can operate to allow warmest exit water possible

- **Workload redistribution to make more racks even utilization**

Example: Flow Rate vs. T1 (inlet water temp)

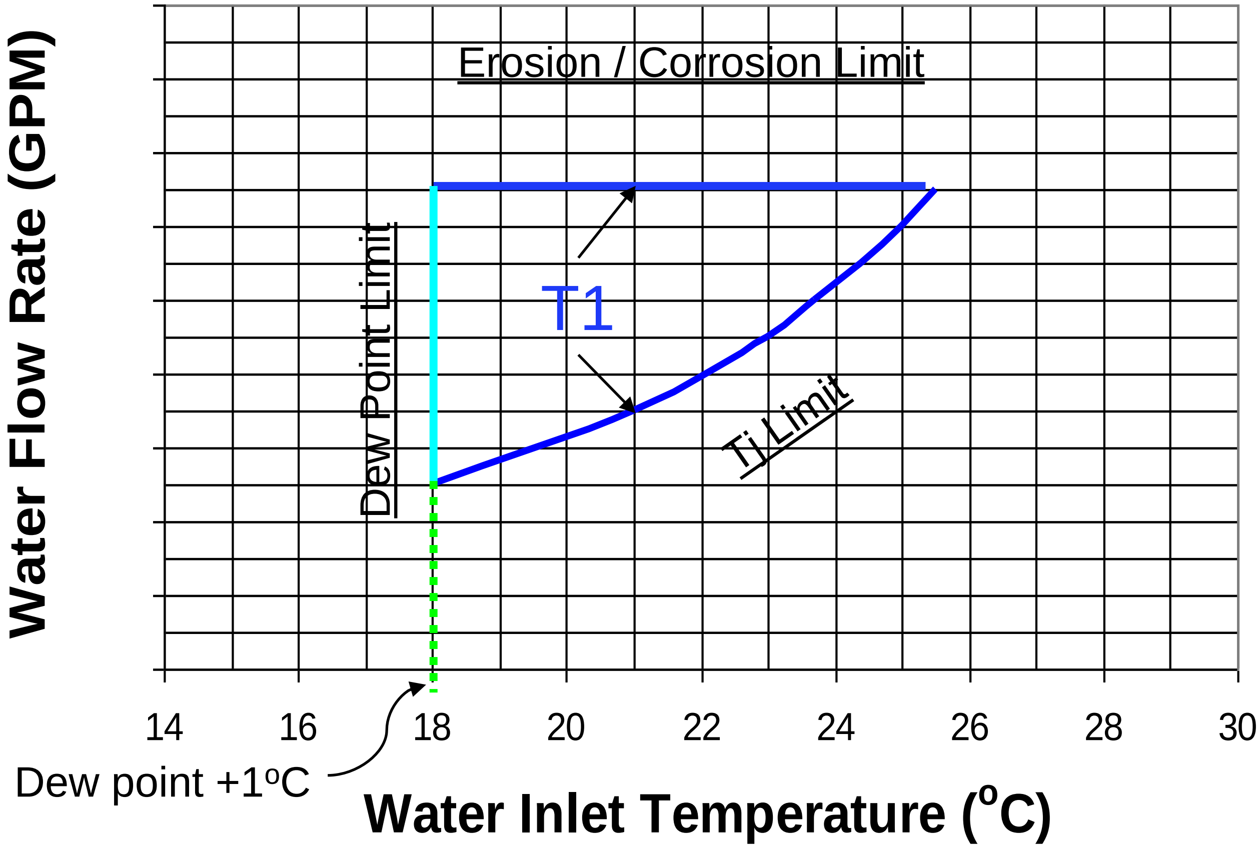

Example: Flow Rate vs. T2 (Outlet Temp)

T1: inlet water temperature to Rack
T2: return water temperature to facility

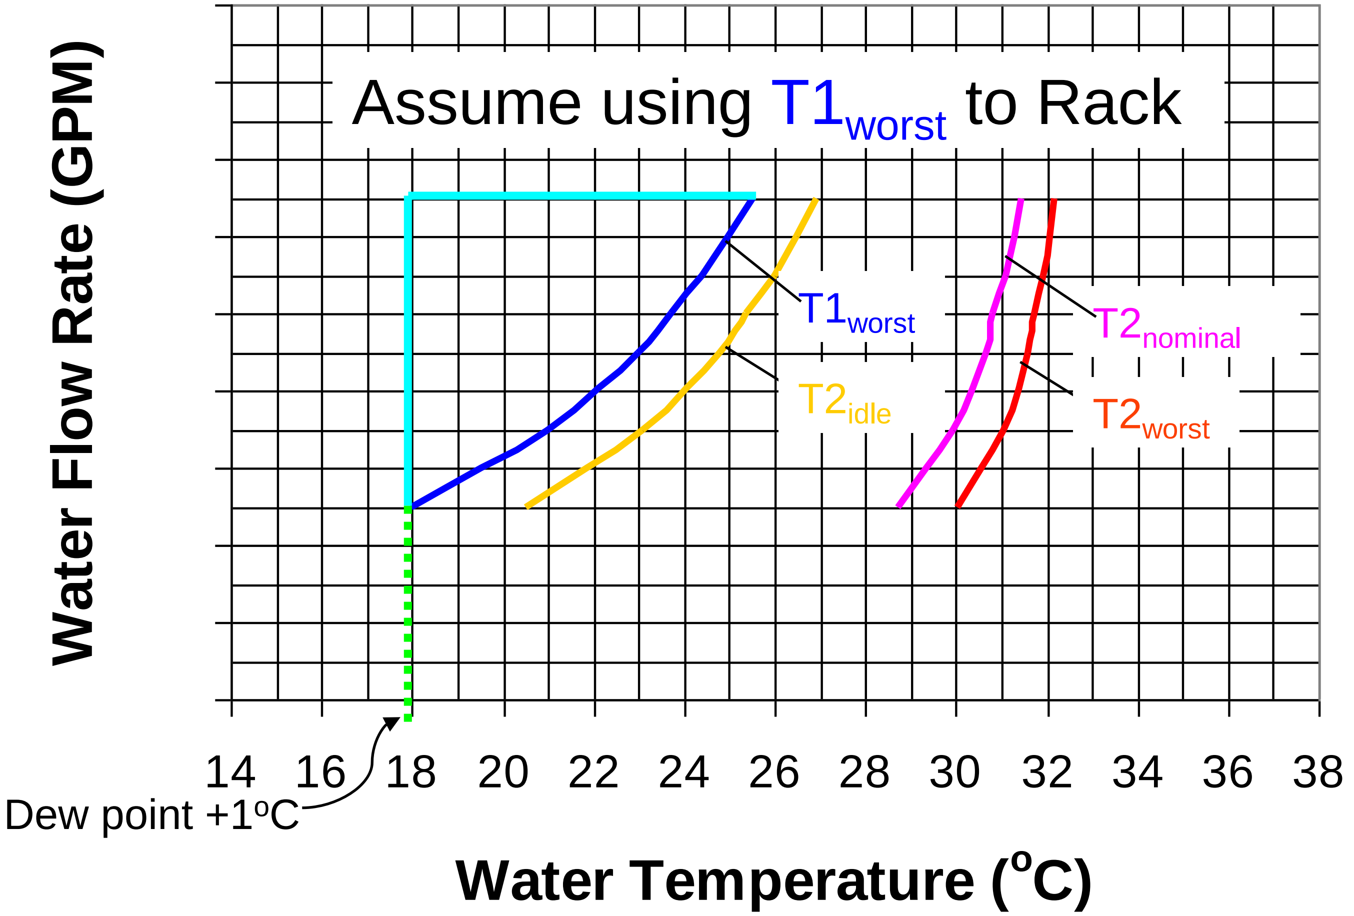

Warm Water Cooling

- **Allowing warmest water to exit the rack allows for maximizing “free cooling”**
  - Dictated by last component in the water stream
  - Improved success if components have uniform power maps (no hot spots)
  - Allows water-to-water heat exchanger to operate at more efficient operating point
  - Can take advantage of air side economizers to remove heat to the ambient

- **Rack level controls to facility CDU might offer opportunity to optimize exit temp**
  - Rack level power feedback to the facility can reduce flow to individual racks – requires valves on each rack

- **Operating envelope is valid**
  - Customer can operate on the Tj limit curve

Water Chemistry Control

- **Copper corrosion inhibitors and biocides (especially in warm water cooling) are critical**

- **Monitoring is crucial to ensure long term stability**
  - Clogging fine pitch cold plate fins
  - Pitting of copper lines -> leaks

- **Equipment exists to continuously monitor water chemistry, conductivity, pH, corrosion rates, and turbidity**
  - Can make on-the-fly modifications to chemistry to keep within optimal range