

# Dynamic Liquid Cooling, Telemetry and Controls

Wade Doll  
Principal Infrastructure Architect, Cray Inc.  
11/17/2015

# Can We Reduce Operating Cost's and Energy Consumption?

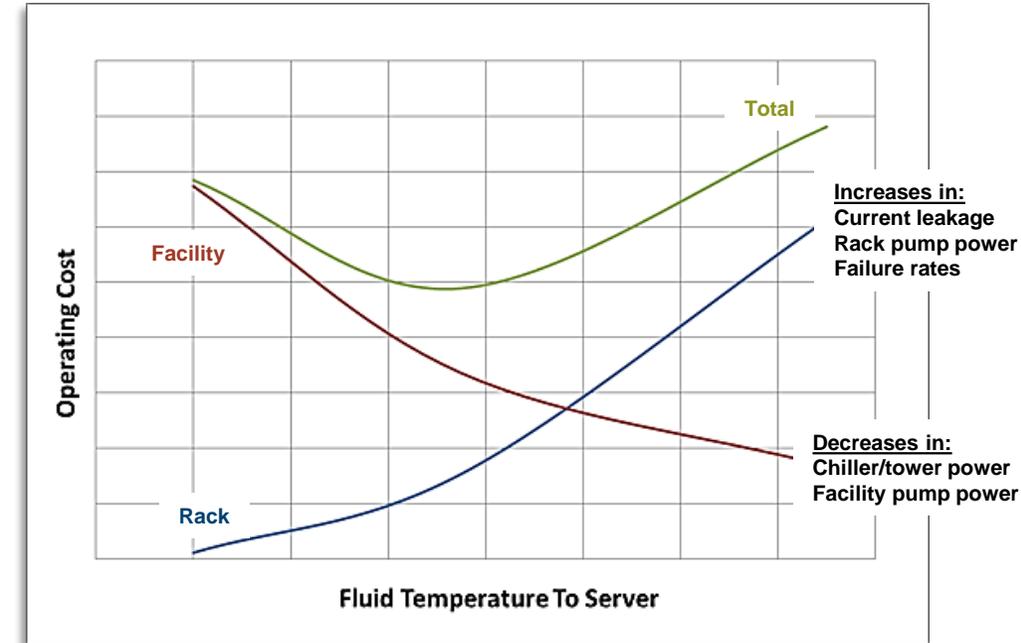
- **Yes, but by how much and at what cost is the real question.**
  - Need to think beyond operating cost (→TCO)
  - Deeper situational analysis
  - Case studies
  - Generating a “Tool” based off analysis and studies would be useful
  
- **Answer will be datacenter and server specific.**
  - Facility dedicated to the system
  - Granularity of control
  - System complexity and flexibility
  - Power variation
  - .....etc
  
- **Integrators and datacenters need to make a conscious effort**
  - Incorporating all the “hooks”
  - Right sizing and proper use of redundancy
  - Procurements need to put a value on this for adoption





# Where Are Controls Best Implemented – Facility or Rack?

- For the most “optimized” solution, needs to work as one harmonious system. But....
  - How do they talk with each other?
  - Who is the master and who is the slave?
  - Should there be control boundaries?
- One system effects the other. There is an optimum point where the total energy consumption is lowest.
  - Non-linear dependencies
  - This may not be at the highest temperature
  - Potentially a third axis “Node Power”
- A practical solution may be both, but with boundaries
  - Share communication and telemetry, but not control
  - Keeps within set “near optimum” conditions
  - Control algorithms specific to facility, server and their interactions



# Thoughts and Experiences

- Getting “all” the telemetry may be expensive
- Idea of control response times; 10 min, 1 min, 1 sec
- Need feed-forward mechanisms for control algorithms
- Changing a variable can have un-attended consequences. May be a good application for Analytics (machine learning/neural networks)
- *HPC procurements need to value this and be the drivers*





**Thank You.**