QUANTIFYING THE COST OF DATA CENTER COOLING APPROACHES INCLUDING SERVICE-LEVEL AGREEMENTS

This thesis developed a comprehensive cost model for evaluating and comparing the financial viability of different data center cooling systems. As data center demand and energy consumption rise, driven by AI and cloud computing, optimizing cooling systems is critical for reducing operational expenditures. Existing cost models overlook the financial penalties associated with downtime defined in Service-Level Agreements (SLAs). This research addresses this gap by creating a model that integrates capital and operational costs with a stochastic discrete-event simulator for maintenance, failures, and repairs. By quantifying downtime penalties in addition to energy and capital expenditures, the model provides a more accurate total cost of ownership. The model calculates the Return on Investment (ROI) and Internal Rate of Return (IRR) to assess the economic justification for deciding when investing in advanced or redundant cooling technologies is economically justifiable. A case study on the CEETHERM lab data center at Georgia Tech shows that the specific method of calculating penalties based on either cumulative downtime, continuous downtime, or interruptions can significantly shift the projected Return on Investment and breakeven timeline. The results demonstrate that interruption penalties had the largest penalties which prevented the system from ever reaching a breakeven point, demonstrating that Service-Level Agreements based on the frequency of service interruptions pose a significantly higher financial risk than those based on outage duration.