# Preventing Catastrophic Cooling Failures in Data Centers

*The Case for Continuous Remote Monitoring and Proactive Water Treatment Architecture*

**PUBLISHED**
August 2026

**FOCUS**
Mission-Critical HVAC / Chiller Plant Integrity

**AUTHOR**
Aaron Kassaw, Advantage Controls

**TARGET AUDIENCE**
CTOs, VP Facilities, Site Reliability Engineers

## 1. Executive Summary

In hyperscale and colocation data centers, thermal management is directly tied to uptime and revenue. High-availability network design, redundant power, and cybersecurity attract sustained capital investment. The physical cooling infrastructure — cooling towers, evaporative condensers, and chillers — receives comparatively little attention despite representing a genuine single point of failure.

Most facilities still manage cooling water chemistry through manual grab sampling on a bi-weekly or monthly cycle. That interval is longer than the timescale on which microbiological fouling develops. Between two sampling visits, an open recirculating loop can move from controlled to severely fouled without any operator becoming aware of it, because nothing in the reactive regime is watching.

The consequences are not hypothetical. Uptime Institute's 2026 Annual Outage Analysis reports that 57% of operators surveyed put the cost of their most recent major outage above $100,000, and for the second consecutive year one in five put it above $1 million.[1] Cooling failures sit alongside software and network faults as recurring contributors to impactful outages, behind power.[1]

This paper examines the mechanics of rapid heat exchanger degradation, situates the financial exposure against published third-party data, and outlines how continuous remote water quality monitoring shifts facilities from reactive response to proactive control.

| 57% | 1 in 5 | ~4× |
|---|---|---|
| OF MAJOR OUTAGES COST MORE THAN $100,000[1] | COST MORE THAN $1 MILLION[1] | BIOFILM INSULATING EFFECT VS. CARBONATE SCALE[3] |

## 2. Illustrative Failure Scenario: Biofouling in a High-Ambient Facility

> *The following is a composite scenario constructed from common failure modes in arid-climate cooling plants. It does not describe a specific facility or incident.*

Consider a colocation facility in the arid Southwest United States, operating open cooling towers under sustained high ambient temperature and high evaporation load. Chemical treatment is managed on a monthly service visit. Biocide is fed on a timer.

Between service visits, the biocide feed pump loses prime. Nothing alarms — the timer continues to cycle, the controller reports normal operation, and no one is measuring oxidant residual. Microbiological populations, released from chemical control during peak thermal season, establish on condenser tube walls.

Two to three weeks later, condenser approach temperature has drifted several degrees from baseline. Head pressure runs high. Operators attribute the change to ambient conditions, which are in fact severe. Under continued load, the unit reaches its high-pressure trip point and locks out. Diagnosis on teardown reveals biological fouling across the heat exchange bundle.

The recovery sequence is where the cost concentrates: emergency chemical cleaning or mechanical tube cleaning, possible rental chiller mobilization, rigging, expedited freight, and premium labor — with hosted workloads degraded or offline throughout.

> **What made this expensive was not the biofouling. It was the two to three weeks during which no one knew it was happening.** A continuous oxidant residual measurement would have flagged the failed feed pump within hours of the loss of prime, when the correction cost was a service call.

## 3. The Kinetics of Degradation

A persistent assumption in facilities management is that water quality degrades gradually, on a scale of weeks to months, and that monthly sampling is therefore adequate. Biological fouling does not behave this way.

### 3.1 Microbiological Fouling

Cooling towers function as large-scale air scrubbers, continuously ingesting dust, organic debris, and airborne microorganisms into a warm, aerated, nutrient-bearing recirculating loop. Operating water temperatures in the 85–105 °F range fall within the growth optimum for most of the relevant organisms, including *Pseudomonas* species and common slime-forming bacteria.

Population growth under favorable conditions is exponential:

$$N(t) = N_0 \times 2^{(t / g)}$$

*where t = elapsed time and g = generation time*

The generation times cited in general microbiology — on the order of 20 minutes — describe laboratory conditions with unlimited nutrient availability and are not representative of field cooling systems, where doubling times are substantially longer and vary with nutrient loading, temperature, and residual oxidant. The relevant point is not the specific value of *g* but the shape of the curve: growth is exponential, not linear, so the interval between "controlled" and "fouled" is short, and a sampling regime that assumes linear degradation will consistently detect problems after they have matured.

Planktonic organisms attach to metal surfaces and secrete extracellular polymeric substance (EPS) — a polysaccharide and protein matrix commonly called biofilm. Once established, biofilm resists oxidant penetration far more effectively than free-floating populations, meaning that the chemical dose sufficient for prevention is well below the dose required for remediation.

Biofilm is a disproportionately effective insulator. Published treatment industry guidance holds that biofilm deposits are up to four times more insulating than calcium carbonate scale of comparable thickness, and that slime will therefore drive up head pressure and shut down a chiller considerably faster than mineral scale.[3]

### 3.2 Mineral Scaling and the Design Fouling Allowance

In high-evaporation climates, dissolved solids concentrate rapidly in the recirculating loop. When cycles of concentration exceed the saturation limits for the makeup water chemistry — typically because blowdown control is delayed or conductivity setpoints are wrong for the season — calcium carbonate precipitates onto the hottest available surfaces, which are the condenser tubes:

$$Ca^{2+} + 2HCO_3^- \rightleftharpoons CaCO_3(s)\downarrow + CO_2\uparrow + H_2O$$

The efficiency consequence is best understood through the rating standard rather than through deposit thickness. Chiller performance is not published for clean tubes: AHRI Standard 550/590 requires that standard ratings incorporate a water-side fouling factor allowance, specified as 0.000250 $h\cdot ft^2\cdot °F/Btu$ for tower heat exchangers and 0.000100 $h\cdot ft^2\cdot °F/Btu$ for the cooling-mode evaporator.[2] That allowance *is* the design margin — the thermal resistance the manufacturer has already assumed will accumulate on the tube surface over normal service.

The practical implication is direct. A chiller carrying condenser-side deposits beyond the rating allowance is no longer delivering its published capacity or efficiency, whatever the nameplate says. The deposit does not need to be thick or visible; it needs only to exceed a margin measured in ten-thousandths of an hour·square-foot·degree per Btu. AHRI treats field variation in fouling as significant enough to warrant a dedicated guidance document on how these allowances are applied across the industry.[4]

> This framing also changes the operational question usefully. **"How much scale is too much?"** has no answerable field version, because no one measures deposit thickness on a running chiller. **"Has condenser approach temperature drifted from its commissioned baseline?"** is answerable continuously — and it is the same question.

The failure risk is downstream of the efficiency loss. Sustained operation at elevated head pressure imposes thermal and mechanical stress on the compressor, and the eventual high-pressure trip is a symptom of a condition that has been developing for months.

### 3.3 Regulatory and Risk Management Context

Open cooling towers are a recognized amplification and dissemination route for *Legionella pneumophila*. ASHRAE Standard 188 establishes risk management requirements for building water systems,[5] with tower-specific practice addressed in CTI Guideline WTB-148[7] and NFPA 214.[8] Many jurisdictions and enterprise insurers now expect a documented water management program with verifiable control measures.

The professional qualification infrastructure has caught up with this. ASSE/IAPMO/ANSI/CAN Series 12000-2024 defines certification tracks for personnel serving on building water management teams — Standards 12060 through 12063 for employers, plumbers, pipefitters and HVAC technicians, and sprinkler fitters, and Standard 12080 for the Legionella Water Safety and Management Specialist.[6] Notably, the 12080 competency list explicitly requires the certified specialist to differentiate between proactive and reactive approaches to prevention. The distinction this paper draws is a named competency in an ANSI- and SCC-approved consensus standard, not a vendor's framing.

The Series 12000 risk framework is also instructive about where a data center sits. Its Appendix E classifies building water systems by type and by occupancy risk group. A system that includes a cooling tower falls into the highest system category, and above the lowest occupancy risk group that combination lands in the top class of control measures — which call for documented chemical control of corrosion, scaling, and microbiological fouling on the towers. Appendix E is informative rather than normative, and it was written with healthcare occupancies in view. The engineering logic transfers regardless: it is the presence of the tower, not the sensitivity of the tenant, that drives the classification.

Continuous monitoring supports this documentation burden directly. It produces a time-stamped, auditable record of oxidant residual, pH, and conductivity rather than twelve annual data points. This is a compliance and evidentiary benefit. **No monitoring or treatment program eliminates Legionella risk, and this document should not be read as claiming otherwise** — the benefit is that control measures are verified continuously rather than assumed between visits.

## 4. Reactive vs. Proactive Treatment Architecture

| VECTOR | REACTIVE (LEGACY) | PROACTIVE (CONTINUOUS REMOTE SENSING) |
|---|---|---|
| **Monitoring frequency** | Manual grab samples, bi-weekly to monthly | Continuous telemetry, 1–5 minute intervals |
| **Biocide control** | Timer-based slug dosing or manual addition | Demand-based dosing against ORP/pH feedback with trend tracking |
| **Fouling detection** | Discovered at high-pressure trip, capacity loss, or tube failure | Early flagging of delta-T and delta-P drift from baseline |
| **Equipment failure detection** | Detected at next service visit | Feed pump loss of prime, sensor drift, and valve failure detected within hours |
| **Operational risk** | Elevated — undetected biological excursions between visits | Substantially reduced — chemical drift surfaces while correction is inexpensive |
| **Compliance evidence** | Twelve data points per year; control assumed between visits | Continuous auditable record supporting a documented water management program |
| **Lifecycle cost** | Higher chemical waste, reactive maintenance, shortened asset life | Optimized chemical consumption, extended equipment service life |

Continuous monitoring does not eliminate cooling system risk. It changes what an operator knows and when, which changes what a correction costs.

## 5. Quantifying Exposure

**DIRECT COSTS**

*Tangible financial loss*

- **Downtime.** Uptime Institute reports 57% of major outages exceeding $100,000 and one in five exceeding $1 million.[1] Facility-specific hourly figures vary widely with tenant mix, SLA structure, and workload criticality; operators should model against their own contracts rather than industry averages.
- **Emergency equipment and mobilization.** Rental chiller placement, rigging, expedited freight, and premium labor for an unplanned multi-hundred-ton replacement.
- **SLA penalty credits.** Contractual breach payouts, frequently structured as multiples of monthly rack rate.
- **Hardware stress.** Thermal excursion shortens component service life and elevates drive failure rates, with costs that appear months after the incident.

**INDIRECT & DEFERRED COSTS**

*Long-term capital drag*

- **Tenant confidence.** Mission-critical workloads relocated after an extended outage rarely return.
- **Commercial impact.** Public incident tracking affects prospective tenant evaluation and lengthens enterprise sales cycles.
- **Engineering opportunity cost.** Senior operations staff absorbed by emergency response are not doing capacity or efficiency work.
- **Insurance.** Loss history influences property and business interruption premiums at renewal.

## 6. Strategic Recommendations

Water treatment in a mission-critical facility should be managed as an instrumented engineering discipline, not a scheduled maintenance line item. Three pillars:

**Pillar 1 — Deploy continuous real-time telemetry.**

Install multi-parameter sensor arrays measuring oxidation-reduction potential (ORP), conductivity, pH, and temperature at minimum. Connect to cloud analytics with 24/7 alert thresholds on both absolute values and rate-of-change. The rate-of-change alarm is what catches equipment failures; the absolute threshold alarm is what catches chemistry excursions.

**Pillar 2 — Implement demand-based chemical dosing.**

Replace fixed-timer biocide addition with dosing tied to measured ORP targets. Oxidant residual targets should be established with your treatment provider against the specific system and biocide chemistry; ORP in the 650–700 mV range is a common target band for halogen-based programs, but the correct setpoint is system-specific and depends on pH, temperature, and interfering species.

**Pillar 3 — Track heat exchanger efficiency continuously.**

Calculate cleanliness factor from real-time approach temperature and flow, referenced to the commissioned baseline rather than to an absolute threshold. This is the field-measurable proxy for the design fouling allowance discussed in Section 3.2. Algorithmic drift from commissioned baseline flags developing fouling well before a high-pressure alarm, and — critically — distinguishes fouling from ambient-driven performance changes, which is the distinction operators most often get wrong.

**WHERE THIS FITS WITH YOUR EXISTING TREATMENT PROGRAM**

Continuous monitoring is not a replacement for a water treatment service provider. The chemistry program, the treatment selection, and the system-specific control targets remain the provider's domain — and under the ASSE Series 12000 framework, program team membership carries its own qualification requirements. What instrumentation changes is the interval at which anyone knows whether the program is working: from monthly to continuous, and from inferred to measured. Facilities operating under a service contract should evaluate this alongside their provider, not around them.

**A NOTE ON LIQUID COOLING**

Direct-to-chip and rear-door heat exchanger deployments do not remove this exposure — they relocate it. Closed loops face different chemistry risks (corrosion, galvanic issues in mixed metallurgy, microbiological growth in low-flow branches, glycol degradation) at far tighter tolerances, since fouling in a cold plate microchannel has consequences at thicknesses that would be unremarkable in a shell-and-tube condenser. Most high-density facilities will run hybrid architectures for the foreseeable future, with a facility water loop rejecting heat through open towers regardless of what happens at the rack. Both loops require monitoring, on different parameters.

## 7. Conclusion

In high-density computing environments, water chemistry is thermal management. A treatment regime built on periodic manual sampling assumes a rate of degradation that biological fouling does not obey, and it leaves an interval during which equipment failures and chemistry excursions go undetected by design.

Continuous monitoring does not make cooling systems immune to failure. It compresses the time between a problem starting and someone knowing about it — which is, in nearly every expensive cooling incident, the variable that determined the cost.

## References


1. Uptime Institute, *Annual Outage Analysis 2026* (8th Annual Outage Trends Report), May 2026. Cost figures derive from Uptime's 2025 Annual Survey as reported in the 2026 analysis. Executive summary: uptimeinstitute.com/resources/research-and-reports/annual-outages-analysis-2026
2. AHRI Standard 550/590 (I-P)-2020 with Addendum 1, *Performance Rating of Water-Chilling and Heat Pump Water-Heating Packages Using the Vapor Compression Cycle*. Rating fouling factor allowances: 0.000250 h·ft²·°F/Btu (tower heat exchangers); 0.000100 h·ft²·°F/Btu (cooling-mode evaporator). Published by AHRI.
3. ChemAqua Technical Bulletin TB2-043, *Monitoring and Troubleshooting Chiller Efficiency*.
4. AHRI Guideline E-1997, *Fouling Factors: A Survey of Their Application in Today's Air-Conditioning and Refrigeration Industry*. Published by AHRI.
5. ASHRAE Standard 188, *Legionellosis: Risk Management for Building Water Systems*.
6. ASSE/IAPMO/ANSI/CAN Series 12000-2024, *Professional Qualifications Standard for Water Management and Infection Control Risk Assessment for Building Systems*. ASSE International, August 2024. Standards 12060–12063 and 12080; risk framework in Appendix E (informative).
7. CTI Guideline WTB-148, *Best Practices for Control of Legionella*. Cooling Technology Institute.
8. NFPA 214, *Standard on Water-Cooling Towers*.



**ABOUT THE AUTHOR**

Aaron Kassaw, Advantage Controls.