




reducing carbon emissions locally and globally



Data Centers and Mission Critical Facilities

Data Centers & Mission Critical Facilities

(Data Centers • Research Labs • High Availability Compute)

Chilled Water & Glycol Based Cooling Systems Only

Why This Matters

In data centers and mission critical environments, **cooling energy and compute performance are tightly coupled**, but the dominant inefficiency is often misunderstood.

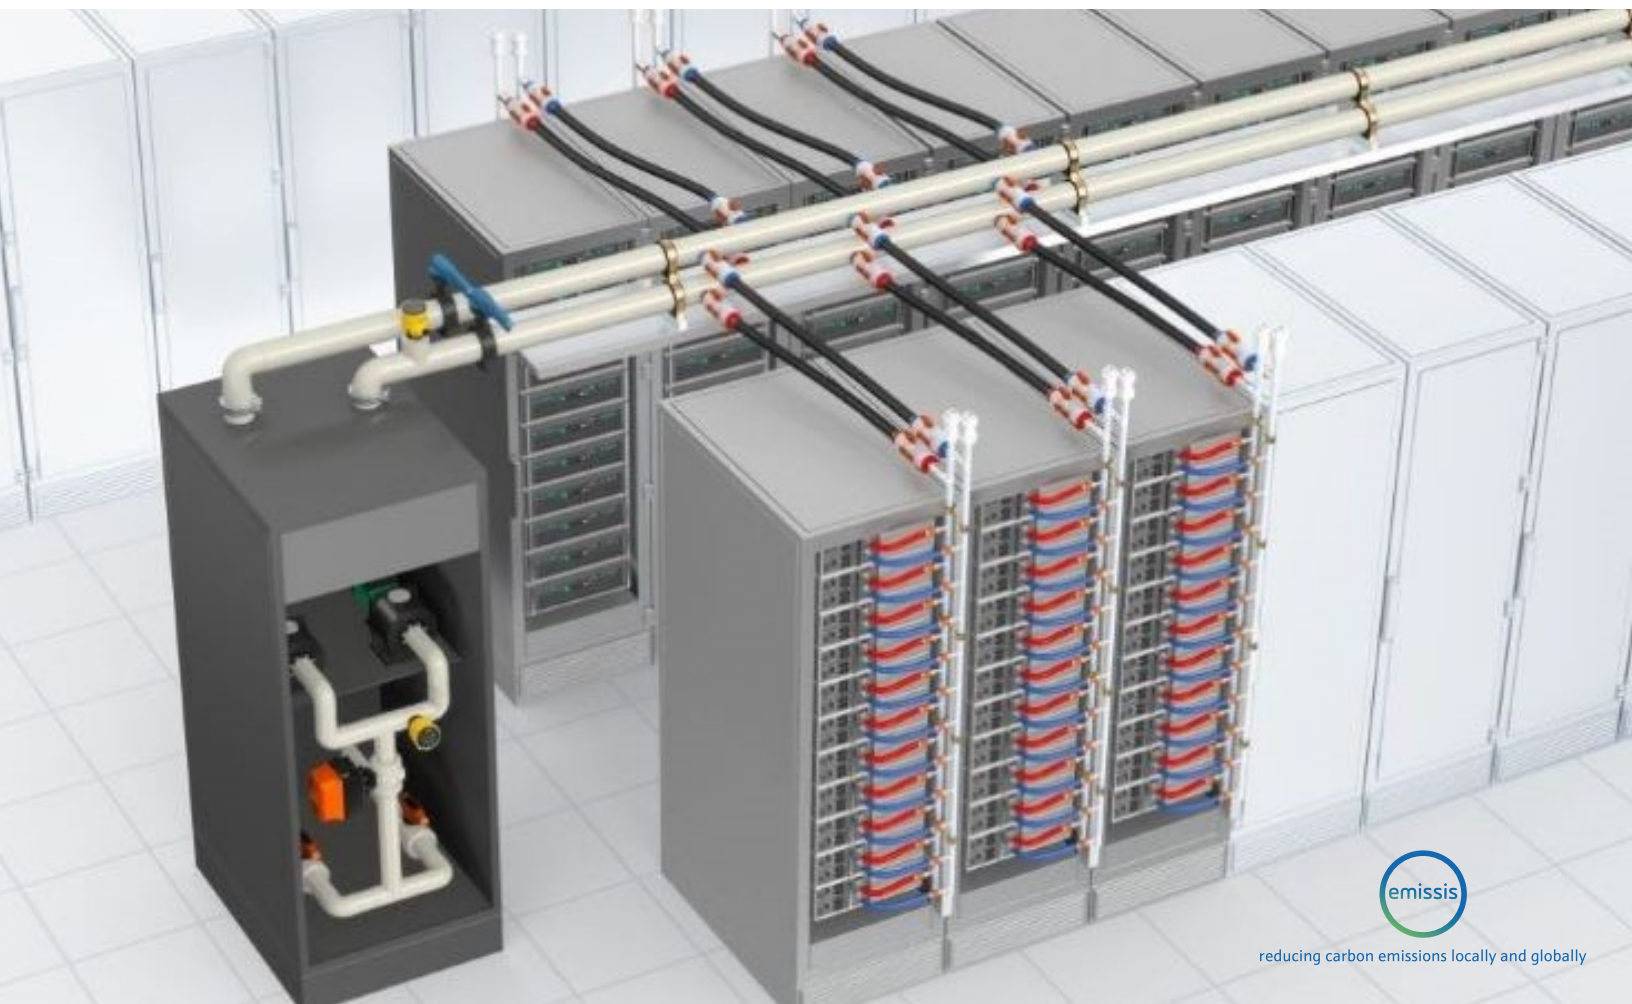
Operators face a persistent trade off at the **thermal conversion layer**:

- Increase cooling energy to protect performance and reliability, **or**
- Constrain energy use and accept higher throttling risk, performance variability, or reduced resilience

This tension is **not caused by poor controls or setpoints**, but by the **electrical energy required to convert power into stable cooling** within hydronic plants.

Modern workloads—AI, HPC, and burst compute—intensify this challenge by creating **rapid, transient load changes** that expose weaknesses in thermal responsiveness, not just steady state capacity.

This makes data center cooling a prime **Layer 2 optimization opportunity** within a stacked energy strategy.



Position in the Layered Model

- **Layer 1 (Electrical Baseline)**: Reduces upstream electrical waste across all loads
- **Layer 2 (This Use Case)**: Reduces the electricity required to maintain existing thermal setpoints
- **Layer 3 (Operational Stability)**: Protects performance under real world load variability

This use case addresses **thermal conversion efficiency only**, ensuring savings stack cleanly without overlap.



The Emissis Advantage

Emissis deploys **enPact Chiller™**, powered by **Maxwell™ patented heat transfer technology**, into **closed loop chilled water and water glycol systems** serving data center cooling plants.

By improving heat transfer at the **fluid physics level**, enPact Chiller™ reduces the **electrical energy required to maintain existing optimal thermal setpoints**—without:

- Raising temperatures
- Altering control logic
- Impacting redundancy or resilience design

Important boundary: This use case applies **only to hydronic based cooling architectures**. Maxwell™ is not applied to DX only systems.



What Changes (Explicit Layer 2 Logic Chain)

1. Lower Energy to Hold the Same Temperature

Improved heat transfer allows the cooling plant to meet identical inlet and supply temperature targets with **less compressor and pump work**.

2. Greater Thermal Stability (Not Just Lower Averages)

Improved thermal responsiveness reduces **temperature variance and transient excursions**, particularly during burst loads and rapid workload shifts.

3. Reduced Thermal Throttling Risk

Because throttling is triggered by **local or transient exceedances**, tighter thermal control results in:

- Fewer throttling events
- Shorter throttling duration
- Reduced internal server fan ramping

4. Higher Realized Processing Output

While nameplate hardware capability is unchanged, **delivered performance over time improves**:

- Fewer cycles lost to involuntary under clocking
- More consistent clock frequencies under sustained load
- Improved job completion times and performance predictability

This reflects **recovery of suppressed performance**, not overclocking or added hardware stress.



Proven Impact (Layer Specific, Cooling & Energy)

- ~12–13% **chiller level COP improvement** demonstrated in a hydronic plant serving a data center environment
- ~9–17% **system level COP improvement** observed across comparable large scale hydronic deployments
- Energy savings achieved **without increasing operating temperatures**


(All results measured within defined plant boundaries using normalized, M&V aligned methods.)

Why It Works

Data centers operate within **tight thermal margins** to protect both performance and reliability. enPact Chiller™ improves the **efficiency and responsiveness of the hydronic loop itself**, allowing operators to decouple **performance protection from energy escalation**.

The outcome is a rare but defensible result in data center operations:

- **Lower cooling energy** (Layer 2 gain)
- **Unchanged thermal setpoints** (no compliance or warranty risk)
- **Improved compute stability** (feeds Layer 3 benefits)



Where It Fits Best

- Enterprise and colocation data centers with central plants
- Research labs and HPC facilities
- AI and burst compute environments
- Mission critical facilities using **chilled water or glycol cooling**

Not applicable:

DX only packaged or split system architectures without hydronic loops.

