



# Chiller Optimization

## Where's the data coming from?



**Dr Stuart Redshaw**  
CTIO & Co-Founder, EkkoSense  
[stuart.redshaw@ekkosense.com](mailto:stuart.redshaw@ekkosense.com)

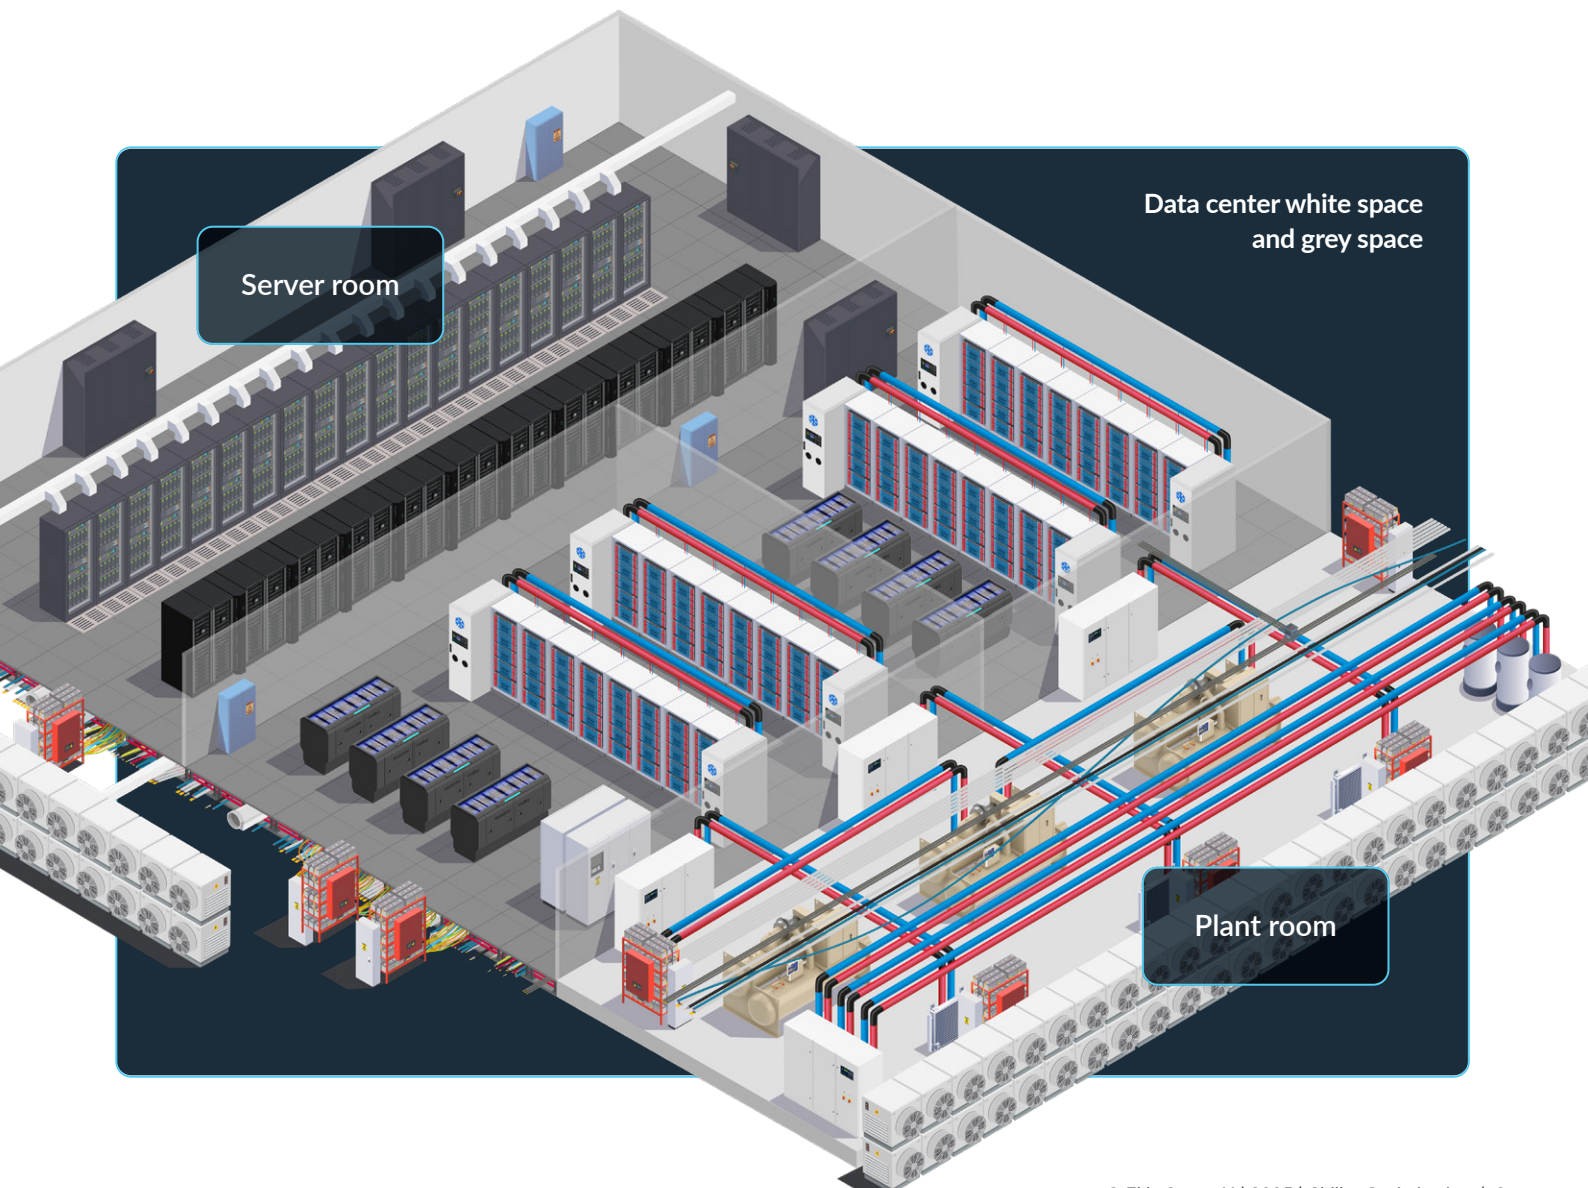




**At EkkoSense we believe it's difficult to unlock the performance improvements needed to handle greater data center workloads and secure energy savings unless you know exactly what's happening across all your multiple rooms in real-time.**

**Realistically this won't be achievable unless data center management commits to bridging the organizational gap between their IT and M&E functions.**

While the latest digital services and AI workloads may run on high density compute platforms with advanced liquid or hybrid cooling solutions, it's still the traditional facilities management teams that manage and maintain the buildings and the critical supporting power and cooling infrastructure within them.







# Bridging the gap between IT and M&E functions

Not surprisingly, many focused IT teams have little interest in the underlying M&E infrastructure that provides the power and cooling that enables their services to keep running.

40%



Average data center cooling utilization

EkkoSense Research

Because of this, it's not unusual to see expensive power and cooling resources being used inefficiently. This can clearly impact negatively on net zero performance, getting in the way of corporate net zero and sustainability initiatives, but also potentially places organizations at risk when critical resources suddenly become depleted or unavailable.

Unfortunately the lack of real-time insight into actual data center management and data center cooling, power and capacity performance means that operations teams often have to over-cool because of this uncertainty and lack of intelligence. Adopting this approach on an ongoing basis effectively prevents data center teams from sensing the true levels of performance optimization that are actually achievable. EkkoSense research certainly confirms this, finding that average data center cooling utilization can often sit at as low as 40% - implying that significant cooling capacity is effectively stranded as teams don't know how to release it and re-apply it elsewhere.

Clearly there's an imperative to optimize and improve cooling performance, but it's important to recognize that the multiple cooling system components – pumps, heat exchangers, compression chillers, dry coolers, heat exchangers and CRAHs – are all interconnected, and can all have a direct impact on actual CPU and GPU server performance.

Operations teams often have to over-cool because of uncertainty and lack of intelligence.





# AI places an increased premium on real-time white space visibility

With the processing of GPU-intensive workloads such as AI clearly set to generate even more heat within data centers, operations teams need to think hard about their current infrastructure and how it will need to evolve.

At EkkoSense we're already seeing AI-driven heat loads exhibiting significant dynamic variability. It's all very different from the reassuring certainty of traditional enterprise workloads. AI applications load up differently, obviously with very high heat loads, but it's the sheer rate of change that's more significant.

Until now there has been a familiar cadence to how traditional data centers are run, but with AI compute workloads increasing it's much more dynamic. As AI compute loads kick in things heat up within seconds. Suddenly the heat is sitting across the rack tops as extra megawatts hit the room. It's exciting, but it also signals that you've absolutely got to be on top of exactly what's going on within your data center rooms before you start making any optimization changes!

That's why at EkkoSense everything we do is informed by engineering first principles. Core heat transfer and thermodynamics disciplines are at the heart of our proposition. This ensures we are able to provide a deep understanding of energy efficiency and heat transfer across all aspects of today's data center operations from HVAC, building services and facilities management through to the latest revolutionary clean tech and AI compute GPU platforms. With this kind of insight, operations teams can be confident of what's likely to happen from an engineering perspective when they think about optimizing different elements of their cooling infrastructure.

It's vital that facilities teams don't set out to optimize individual cooling systems and components without considering the entirety of their activity's impact on overall data center performance. Take chiller optimization as an example. We believe that it's essential to understand both the complexities of chiller optimization within the grey space as well as its direct relationship with real-time white space performance before even considering any standalone AI-powered optimization approaches.

Core heat transfer and thermodynamics disciplines are at the heart of our proposition. This ensures we are able to provide a deep understanding of energy efficiency and heat transfer across all aspects of today's data center operations.

AI

$$\begin{aligned} x_1 &= \frac{m_d}{m_p + m_d} \text{ but } \frac{m_d}{m_p + m_d} + \frac{m_p}{m_p + m_d} = 1 \text{ (100\%)} \\ \frac{m_d}{m_p + m_d} &= 1 - \frac{m_p}{m_p + m_d} = 1 - x_1 \\ \frac{m_p}{m_p + m_d} T_1 &= \frac{m_d}{m_p + m_d} T_2 + T_2 \rightarrow (1 - x_1) T_1 = x_1 T_2 + T_2 \\ x_1 &= \frac{T_2 - T_1}{T_3 - T_1} \text{ and } \frac{m_p}{m_p + m_d} = 1 - x_1 = 1 - \frac{T_2 - T_1}{T_3 - T_1} \end{aligned}$$



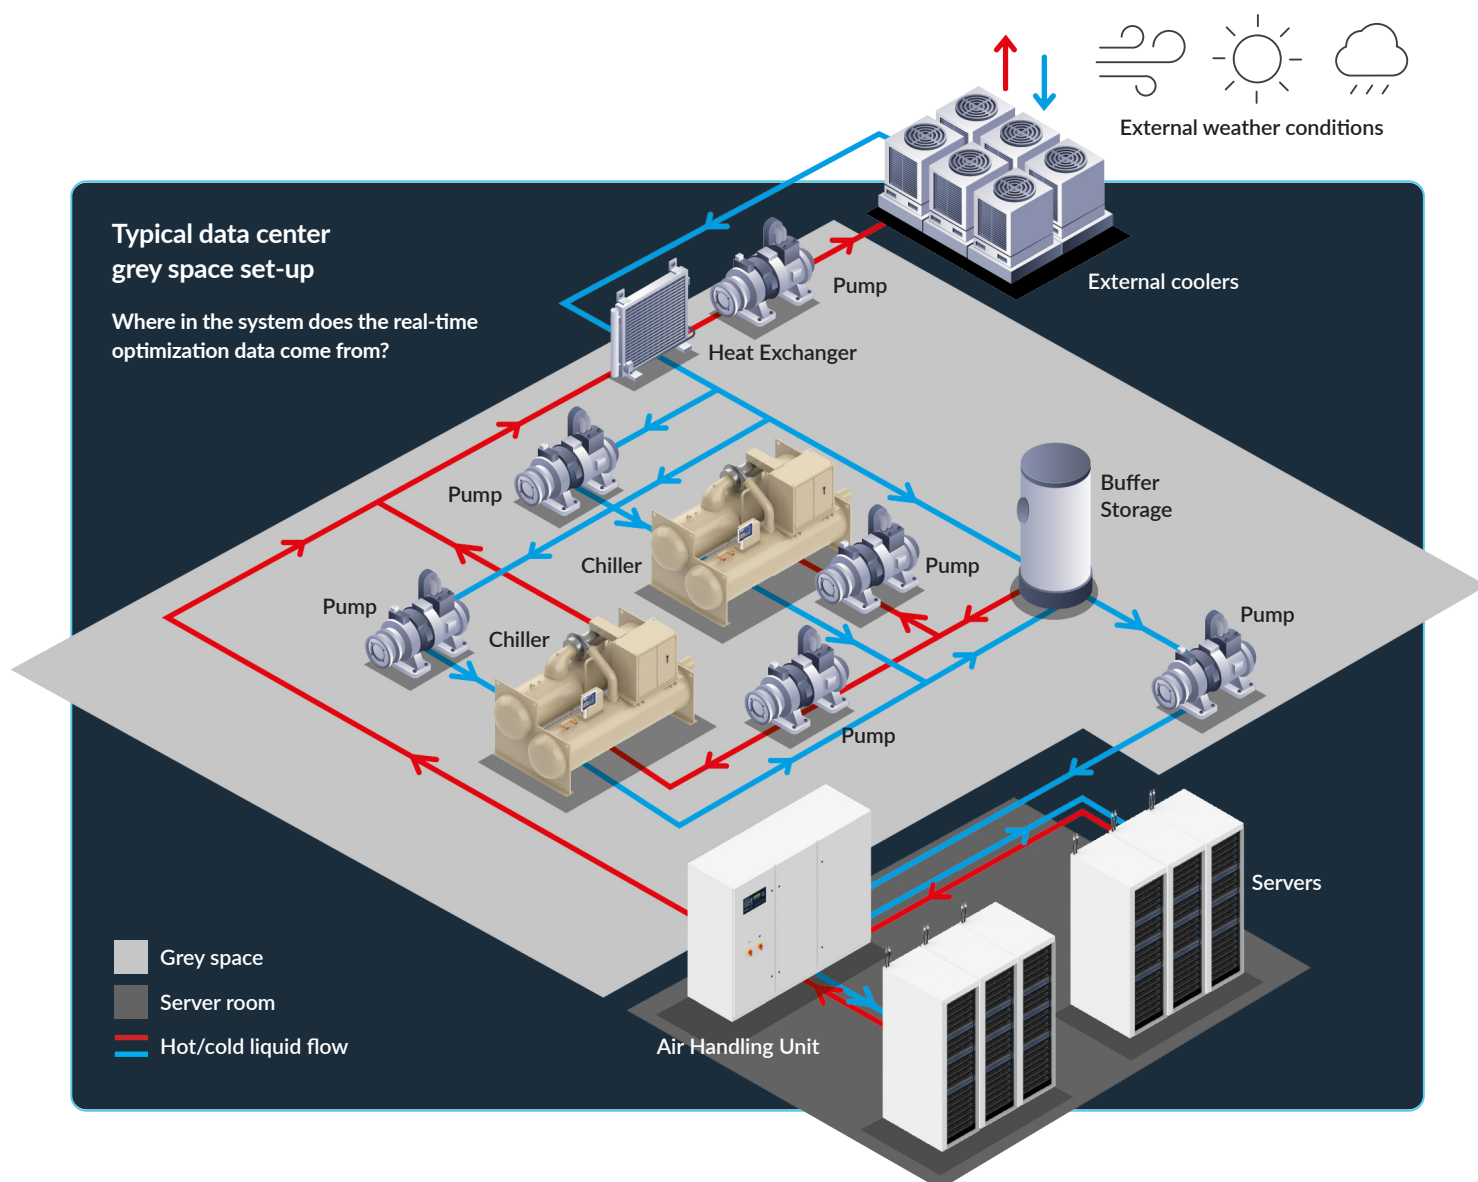
# There's more to chiller optimization than just the chillers

Lately the chiller optimization process has become synonymous with AI, with algorithms developed to adjust setpoints and save energy by increasing temperatures.

But does chiller optimization really require AI to do the things that it's doing? Firstly, there's a lot more to it than just the chillers! It's actually everything that's working in the grey space to deliver the cooling that needs optimizing.

There's pumps of various types and functions that don't always sit within the chiller's AI algorithm – so you always need to be asking where exactly does the chiller start and stop? The pumps clearly have an impact on the process, but what about the effects of external weather conditions or secondary load demands? More importantly, where's the actual raw data coming from that can help you to optimize chiller performance? If you don't have real-time visibility of data center thermal and power performance, how can you expect an AI to make the correct decisions? We believe it's actually where the data is coming from that's important – not the algorithm, not how you process it – it's always about the raw data.

Clearly there are risks associated with optimizing a data center's chillers and pumps without having complete visibility over what's actually going on in the room. Chiller optimization shouldn't be entirely reliant on AI and has to consider the end-to-end cooling system – including pumps and room performance. That's why we believe it's important to optimize a room's thermal, power and capacity performance before moving up the stack. Of course any optimization and performance improvement of cooling within the room will already have a positive impact on the energy consumption of the chillers before any focused optimization within the grey space.







# Finding the sweet spot where data collection works best


**The only truly reliable way for data center teams to troubleshoot and optimize white and grey space performance is to gather massive amounts of data from right across the facility – with no sampling.**

This removes the risk of visibility gaps but also introduces a new challenge in terms of the sheer volume of real-time data that is being collected. Early on EkkoSense established that each single rack in a data center is a contained thermal environment, and as such, each has its own thermal reality. This was at a time when people were only capturing thermal data at the room level, or every 10th or 20th rack.

The EkkoSense team determined that the granularity of data really matters. However, it's not just data quantity that matters, the quality is also vital. ASHRAE determined many years ago that the specific location of a sensor makes a huge significant difference in the quality of data gathered. Data center operations teams simply can't assume that their historic sensor locations have kept up with the IT systems deployed, and that they continue to provide the best indicators of performance and risk for their sites.

And what matters for the white space also matters for the engineering grey space. This very specific and detailed consideration of the quantity and precise location of sensors is totally at odds with the AI algorithm-only approach. To get optimization right you have to be incredibly structured, very precise, and you need to measure specific things – but the good news is that you don't need to measure everything everywhere. EkkoSense research has identified the key specific locations to measure that all data centers share. But these locations are outside of the traditional places that BMS systems have concentrated on, since their focus has always been on the performance of individual system components – rather than the performance of the system as a whole.

EkkoSense research has identified the key specific locations to measure that all data centers share.

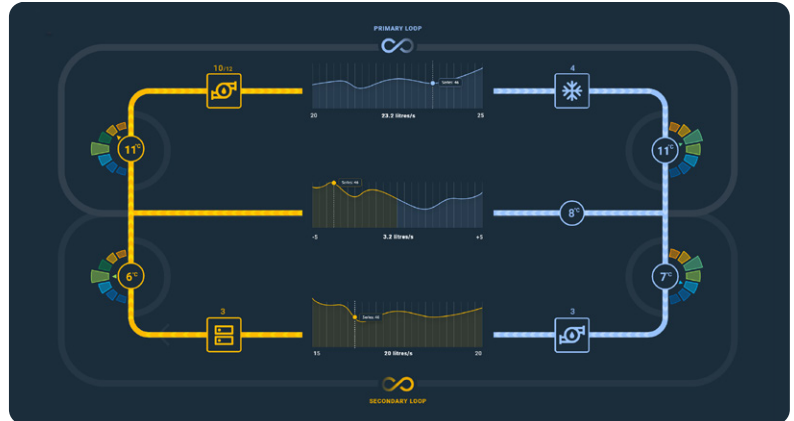




# Chiller optimization must be all about getting the right balance

When it comes to effective chiller optimization it's all about successfully balancing chiller staging, flow rates and temperatures with the requirements from IT loads.

First though it's important to consider the characteristics of the different types of chiller systems:



## Chilled water systems with a primary-only system

Here primary chiller water flow is pumped continuously at a constant flow rate and temperature to the loads, regardless of the variations in cooling demand. Three-way valves located at the loads are then used to control the temperature by diverting a portion of the chilled water during part-load conditions to maintain the required thermal balance. This approach can provide reliable cooling performance against various scales, though at the expense of reduced efficiency. If the return water temperature is too low, it can also lead to low Delta-T syndrome.



## Chilled water systems with a primary/secondary arrangement

This approach is mostly found in larger chilled water systems such as data centers with variable cooling loads. The primary circuit consists of low-head pumps that maintain a constant flow rate through the chillers, ensuring stable operation. The secondary circuit contains high-head pumps that ensure sufficient water flow through the cooling loads. This approach is more complex due to the additional pumps and controls involved. However, by dynamically adjusting flow rates in response to varying loads, the primary/secondary systems ensures consistent cooling while reducing energy consumption in larger data centers.



## Variable Primary Flow Systems

State-of-the-art data centers with variable cooling loads need advanced cooling systems that optimize controls and efficiency of the chillers by adopting variable speed drive pumps with integrated controls. Secondary variable frequency drive pumps are sized to overcome the friction through the load and distribution piping connecting it to the decoupler, while primary pumps need only overcome the friction through the primary piping and the chillers evaporator. Control of the primary pumps is based on the temperature at each side of the decoupler – and should be set up so there is no flow through the decoupler when there is sufficient cooling load. The chiller itself should be protected against low flow by monitoring chiller differential pressure or using a flow meter to set the primary flow pump to the minimum speed until the demand for chilled water increases. Where and how this low flow monitoring is located will be critical to the performance of the overall grey space and the white space facilities it serves.





# Re-imagining grey space optimization with EkkoSense

As set out above, it's essential for operations teams to draw on the physics and flow dynamics of the cooling system if they're to effectively re-imagine managing and optimizing chiller performance.

This is where EkkoSense has been focusing its activities, leveraging its expertise to identify precisely the optimum location for monitoring grey space performance.

What's critical is to focus in on the right place – the nexus of grey space information – that will tell us everything we need to know about temperatures and flows at that precise point. We can't just expect to access this data directly – we have to collect it ourselves, accessing information from the critical locations that will allow the system to provide the required chilled water without wasting energy.

Accessing the correct data is just the start for EkkoSense. Using this data allows us to write the most powerful algorithms that can then not only support grey space optimization, but also tie it directly to what's happening within your data space rooms in real-time. This approach recognizes that it's never just about the chillers. If you're serious about tracking and balancing cooling loads across all your rooms then you need to be looking beyond just chillers to include pumps, secondary pumps, primary pumps, cooling towers, heat exchangers – effectively the end-to-end cooling system – it's the whole piece.



**If you're serious about tracking and balancing cooling loads across all your rooms then you need to be looking beyond just chillers to include pumps, secondary pumps, primary pumps, cooling towers, heat exchangers – effectively the end-to-end cooling system – it's the whole piece.**





# Contact EkkoSense to continue the conversation and discuss your optimization challenges



Click or scan

## Have you read...



Click or scan



UK Headquarters: +44 (0) 115 678 1234  
North America: 1-833-921-3335  
Germany: +49 89262025276  
Australia: +61 2 8358 0031  
[info@ekkosense.com](mailto:info@ekkosense.com) [www.ekkosense.com](http://www.ekkosense.com)



WINNER

