

2023 – The year to deploy Machine Learning and AI in your data center?

Can you use Machine Learning and AI to balance the impossible – increased data center workloads v energy reduction / ESG targets?

Today's data centers face a challenge that, initially, looks like it's almost impossible to resolve. Operations have never been busier, and analyst projections suggest workload levels are only going to increase over the next five years – with some projecting compound annual growth rates of as much as 21% through 2025. At the same time, critical facilities such as data centers are coming under increased pressure to reduce energy consumption - particularly as the reality of corporate ESG programmes and net zero commitments start to bite.



Jonathan Gould
Vice President NAM
EkkoSense
jonathan.gould@ekkosense.com



So how can data centers resolve what appear to be two potentially conflicting demands: supporting escalating workloads while still cutting their carbon usage and helping to support ESG goals?

In this White Paper EkkoSense argues that traditional data center software toolsets - such as BMS, EPMS, CFD and DCIM - can't provide a credible answer as they don't equip operations teams with a complete view of what's really happening in their data centers.

Instead, the answer requires a fundamentally different and more innovative approach, combining more comprehensive sensing, advanced software optimization tools, and the power of Machine Learning and AI algorithms to provide a true real-time visualization of data center performance.

This White Paper sets out how Machine Learning and an AI-enabled software-based optimization can now play a key role in helping data centers to deliver against both workload demands and energy reduction goals.

It will also detail how this approach is already helping EkkoSense customers to deliver the benefits below.



Reduce average data center cooling energy costs by up to 20%.



Increase data center cooling utilization beyond an industry average of just 40%.



Ensure 100% rack-level ASHRAE thermal compliance – protecting against the cooling issues that still account for almost one third of unplanned data center outages.



Contribute directly, and quickly, to corporate ESG programme targets and Net Zero commitments, with quantifiable carbon reductions that are achievable within weeks.



And all deployed and operational through a light-touch approach that enables ROI within only 24 months.



Moving beyond traditional data center infrastructure tools.

Before unlocking these potential benefits, data center operations need to look beyond their legacy critical infrastructure software toolsets that don't always give them the real-time answers they need.

With data centers focused on supporting the 'essential' technology services that people increasingly demand to support their personal and professional lives, it's hardly surprising that operations have never been busier. Driven by trends that show no sign of slowing down, we're seeing massively increased data usage associated with streaming services, storage, compute demands, smart **IoT integrations**, as well as **5G connectivity** rollouts. However, despite these escalating workloads, the unfortunate reality is that many of today's critical facilities simply aren't running efficiently enough.

Given that the average data center operates for over 20 years, this shouldn't really be a surprise. Efficiency is invariably tied to a facility's original design, and frequently based on expected IT loads that have long been overtaken. At the same time change is a constant factor, with platforms, equipment design, topologies, power density requirements and cooling demands all evolving with the continued drive for new applications. The result is a global data center infrastructure that regularly finds it hard to match current and planned IT loads to their critical network. This will only be exacerbated as data center demands increase.

Traditional data center approaches are struggling to meet these escalating requirements. Prioritizing availability is largely achieved at efficiency's expense, with too much reliance still placed on operator experience and trusting that assumptions are correct.

Unfortunately, the evidence suggests that this model is no longer realistic. **EkkoSense research** reveals an average figure of 15% of IT racks in data centers operating outside of ASHRAE's temperature and humidity guidelines, and that customers strand up to 60% of their cooling capacity due to inefficiencies. And that's a problem, with Uptime Institute estimating that the global value attributed to inefficient cooling and airflow management is around \$18bn. That's equivalent to some 150bn wasted kilowatt hours.

Need to support real-time optimization.

Unfortunately, while traditional data center infrastructure toolsets may have very solid use cases, they're less flexible when it comes to directly supporting real-time optimization activities. BMS is of course a key platform, but it typically doesn't feature any analytics and is usually designed to alert only on hard faults or SLA breaches – which is often too late to prevent outages. Similarly, Electrical Power Management Systems are often treated as a BMS extension.

While CFD systems can be great to support new build projects or major design changes, they typically only utilize data from a specific point in time under set parameters – making it overly complex and expensive to unlock real-time optimization opportunities. DCIM systems too tend to be largely driven out of IT requirements, with very little focus on the M&E side. Because many DCIM vendors originated from the IT side of the fence and, although they might have claimed comprehensive functionality for 'inside the rack', none of them have yet to properly address the very real M&E needs of data center operators – particularly when it comes to overall energy efficiency and capacity management.

There's no doubt that the size of the optimization prize is considerable and achieving this will require a break from legacy infrastructure tools. A Machine Learning and AI approach doesn't just highlight problems as they occur, but delivers intelligent insights in real-time and recommends solutions before potential issues even have a chance to develop.





Setting out how Machine Learning and AI can really make a difference.

For several years, AI and Machine Learning have been used as buzzwords to signal the vision of an automated data center that's more resilient and costs less to run. But the reality, as we have shown, is that most data center operators are still living in a reactive world – often spending most of their time chasing down problems and putting out fires.

However, rather than consider AI as a kind of miracle data center infrastructure management plugin – one that can suddenly monitor, manage and optimize all of a data center's power, cooling and capacity requirements – it makes more sense to concentrate initially on specific areas where Machine Learning and AI can be applied and can deliver significant results.

Effective AI solutions rely on extensive access to accurate granular sensor data, so that they can learn and adapt when exposed to new findings. Machine Learning systems are capable of analyzing information from very large data sets, and then detecting and extrapolating patterns from that data to apply them to evolving scenarios. And because today's IT platforms clearly have access to massive computing power, they are now able to identify, analyze and act on data. However, the key question remains: what data are they going to process to get the AI insights they require?

EkkoSense believes that the application of Machine Learning and AI creates a new paradigm in terms of how to approach data center operations. Instead of being swamped by too much performance data, operations teams can now take advantage of Machine Learning to gather data at a much more granular level – meaning they can start to access how their data center is performing in real-time.


The key is to make this accessible, and using smart 3D visualizations is a great way of making it easy for data center teams to interpret performance data at a deeper level: for example, by showing changes and highlighting anomalies.

The next stage is to apply Machine Learning and AI analytics to provide actionable insights. By augmenting measured datasets with Machine Learning algorithms, data center teams can immediately benefit from easy-to-understand insights to help support their real-time optimization decisions. The combination of real-time granular data collection every five minutes and AI/Machine Learning analytics allows operations not just to see what is happening across their critical facilities but also find out why – and what exactly they should do about it.

AI and Machine Learning powered analytics can also uncover the insights required to recommend actionable changes across key areas such as optimum set points, floor grille layouts, cooling unit operation and fan speed adjustments. Thermal analysis should also help to indicate optimum rack locations. And because AI enables real-time visualizations, data center teams can quickly gain immediate performance feedback on any actioned changes.


EkkoSense believes that the application of Machine Learning and AI creates a new paradigm in terms of how to approach data center operations.





How Machine Learning and AI can support the five key stages of effective software-based optimization.

Cooling optimization and airflow management are proven areas where the technology can deliver tangible results. At EkkoSense we focus on five core data center optimization stages where Machine Learning and AI can make a real difference:



Gather accurate cooling, power, and space data at a highly granular level.

This means that you can observe how your data center is performing in real-time. Measuring what's happening at the cooling units, rack temperatures, what the UPS is doing, what power is being produced can all quickly add to a significant level of data – potentially over 100,000 new data points each day – and that's before you've worked out how to correlate all that data in real-time. However, with the availability of the latest ultra-low-cost wireless sensors, there's really no excuse not to do this in 2023. Low-cost wireless EkkoSensors, for example allow new levels of high spatial resolution, down to individual rack-level with multiple sensor points.



Make it easy to visualize complex data easily and quickly.

As humans we're simply not good at staring at complex spreadsheets and trying to extrapolate what they are telling us. What we need is a new way to understand the relationships between all those data points, understanding the Zones of Influence, and learning what impacts what. Thermally you want to see those greens and light blues, not the dark blues and reds. By creating comprehensive **3D digital twin** visualizations, it's a lot easier for data center teams to interpret this information – making it much easier to see what's happening across your estate, highlighting potential anomalies and displaying suggested airflow and cooling improvements.



Apply Machine Learning and AI analytics to provide actionable insights.

Augmenting measured datasets with Machine Learning algorithms to provide data center teams with easy-to-understand insights to support real-time optimization decisions. Rather than treat Machine Learning and AI as a universal solution, we expect to see data center teams focusing on those areas – such as cooling optimization and airflow management – where these technologies

can be applied and deliver significant results. AI and Machine Learning changes the game, taking all those complex data sets and crunching the numbers, doing in seconds what people would take weeks to attempt and most likely fail to achieve. With AI in place, you benefit from fully correlated real-time data that's presented in a distinctive, actionable way.



Ensure delivery with actionable recommendations for human auditability.

Rather than rely on unwieldy automation solutions, EkkoSense advocates a light-touch approach to data center infrastructure management, with operations teams supported by AI-powered, actionable recommendations for greater human auditability. Operators still maintain control over decisions, recognizing that often they will have insights that simply weren't available to the data sets. And because optimization recommendations are presented each time for human auditability, data center teams are always on hand to ensure that any changes are delivering the expected results.



Deploy an ongoing continuous optimization approach.

Giving data center staff the capability for continuous optimization, while supporting them in keeping pace with their ever-changing critical facilities. You cannot optimize periodically and expect to stay in sync. Drift is natural, because regardless of how static you believe your environment to be, change is always occurring. The Machine Learning algorithm is persistently monitoring and analyzing your environment. It is always looking for ways to improve the operation of the environment, allowing you to continually improve the way you operate the facility.

AI and Machine Learning changes the game, taking all those complex data sets and crunching the numbers, doing in seconds what people would take weeks to attempt and most likely fail to achieve.





Taking things to the next level with 3D digital twins that help manage the Machine Learning process.

The only reliable way for data center teams to troubleshoot and optimize data center 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 potential challenge in terms of the sheer volume of real-time data that is being collected.

Combining the power of Machine Learning with real-time data from a 'fully-sensed' room enables the creation of a digital twin of an organization's data center layout – one that not only visually represents current cooling, power, and thermal conditions, but also provides tangible recommendations for optimization. This level of decision support can help operations teams to take things to the next level, as the software will recognise when changes have been made and even refresh the digital twin to reflect updates.

At EkkoSense we address this through the powerful 3D visualization capabilities offered by our EkkoSoft Critical software, enabling the building, population and editing of customer data center rooms for presentation via first person, plan, 3D views as well as power schematics. Combining the power of Machine Learning with real-time data from a 'fully-sensed' room enables the creation of a fully-realised Digital Twin of your data center layout – one that not only visually represents current thermal conditions, but also provides tangible recommendations for thermal, power, and capacity optimization.

So instead of simply automating systems and trusting AI to get on with managing the sensitive security and controls needed for critical data center cooling duty performance, EkkoSense offers a far more productive approach. Gathering cooling, power and space data at a granular level, visualizing that complex data using the latest video gaming technology to make it much easier to compare changes, highlight trends and anomalies, and then using Machine Learning techniques and AI-powered recommendations to provide actionable insights to your data center team.

Overview

The Power of EkkoSense AI

EkkoSense AI and Machine Learning are changing the game for data center operators

Typical BMS View

Most data center teams see only the cooling unit temperatures. Rack inlet temperatures are not properly monitored and their true status is invisible.

EkkoSense 3D Interactive Digital Twin

Comprehensive thermal data capture and monitoring gives full visibility and uncovers risk, predicts failures and provides opportunities for improvement.

EkkoSense AI and Machine Learning

AI powered software with Machine Learning allows operators to quickly and easily fine tune the data center to ensure maximum efficiency.

BEFORE EkkoSense optimization

AFTER EkkoSense optimization

100% ASHRAE compliant operating environment

Safer IT and M&E equipment
More efficient with less energy consumed

ekkosense.com

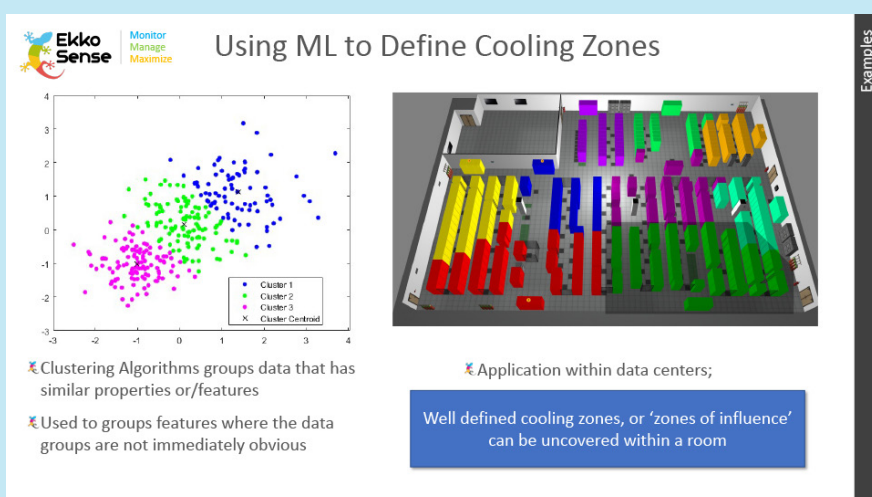


Driving software-based optimization with EkkoSense Machine Learning algorithms.

Instead of simply automating systems and trusting AI to get on with managing the sensitive security and controls needed for critical data center cooling duty performance, EkkoSense believes in a more productive approach.

Gather cooling, power and space data at a granular level, visualize that complex data to make it easier to compare changes and highlight anomalies, and then use Machine Learning algorithms to identify specific areas where actions could be needed. Key EkkoSense examples here include:

Using machine learning to define cooling Zones of Influence



By capturing entirely new levels of data center cooling data (going beyond basic temperature measurements to also include energy usage and airflow distribution) EkkoSense can map zone-by-zone cooling analytics within the data center to help with resiliency and capacity planning decisions. This creates live airflow 'zones of influence' that can group racks into clusters that specifically match the cooling unit that's provisioning them. By then sharing these insights with Cooling Advisor – the industry's first embedded AI-enabled advisory solution – EkkoSense is then able to show precisely which critical equipment would be impacted by changes to cooling configuration, and advise on the most efficient operational changes.

With this kind of real-time thermal monitoring in place, data center teams can track cooling outputs and identify any poorly performing cooling systems in advance so

timely improvements can be made. EkkoSense supports both granular rack and cooling unit level monitoring to find the hidden – but easy to fix – cooling and airflow problems that typical cooling PPMs and BMS systems fail to find or diagnose.

Taking advantage of a powerful clustering algorithm, the Zones of Influence approach works particularly well for data centers that are air-cooled, either under floor or directly into the space – effectively supporting most data centers worldwide from legacy systems to new builds. By knowing 'which' equipment is reliant on 'which' cooling helps determine a risk profile for the operation of the data center. Because most data centers have sizeable cooling redundancy still, single inactive units may not be such an issue – but a faulty unit or an inefficient one may cause recirculation problems that have a direct effect on the equipment within their zone.



Visualizing critical equipment performance within data centers

Visualizing the Data


Visual representation of temperature
Good for alerting and managing SLAs, but the cooling relationships remain hidden.

Clustering algorithms are great at uncovering the hidden cooling configurations within the data

Examples

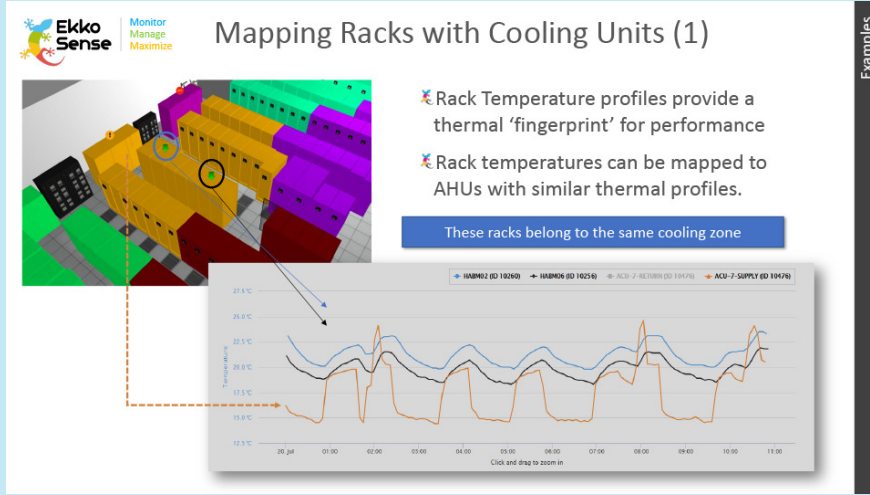
Viewing data center rooms via real-time 3D visualizations provides a great visual representation of temperatures thanks to monitoring points on all the critical equipment. However, while it provides an immediate way of alerting and managing SLAs, the underlying cooling relationship remains hidden. Thanks to the Zones of Influence Machine Learning cooling algorithm, a different view

can show the same section of the room broken into cooling zones (as shown on the right side of the image). You can see that the cooling unit at the top of this image is off, meaning that the orange cooling unit next to it is having to provide cooling to four rows of racks. This then requires further insight as to how data is captured on individual racks.





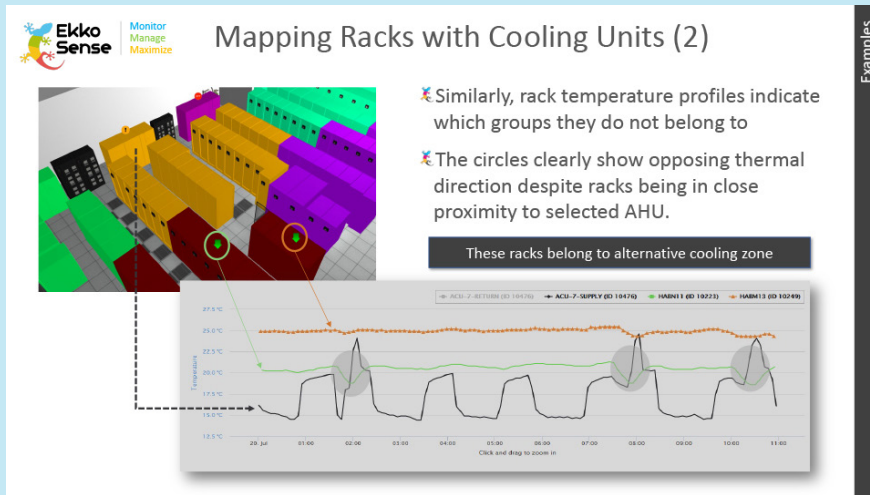
Mapping racks with cooling units



Examples

Selecting two of the racks within the same zone as the cooling unit shows a pattern that is very similar to the cooling unit supply temperature pattern. Rack temperatures go down when the unit is cooling, and go

up when one of the compressors is off. Adjusting for latency within the system, the rack is being 'cooled' by the selected cooling unit.




Examples

By then taking two racks that are in the neighbouring cluster, and bringing up some of the data we can see that the thermal profile for the rack inlet temperatures doesn't match the cooling unit supply temperature well. The circles marked on the graph below clearly show areas where the rack inlet temperature is falling despite the cooling backing off. While this example is quite dramatic in terms of temperature swings, with current

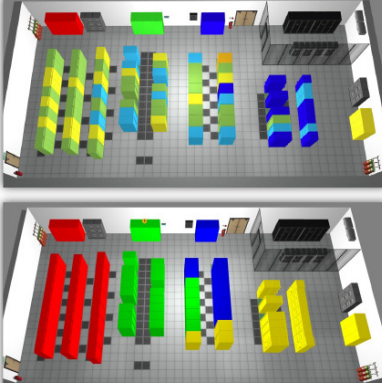
sensor technology it is also possible to deduce these patterns where the fluctuations are much smaller. In fact, one of the key values from this Zones of Influence approach is not so much how the absolute temperatures match up, but more in the relative movement of the sensor. And given that the latest digital sensors available don't require any calibration, they can provide a critical source of Machine Learning data for automation systems.



Adapting to operational changes



Adapt to Operational Changes – Before




- ✘ Full granularity of monitoring
- ✘ Cooling zones adapt to cooling operational changes
- ✘ In this example, the rack temperatures are within ASHRAE TC9.9, but the BLUE AHU is not affecting many racks
- ✘ An example of too much airflow in the space provided

Examples

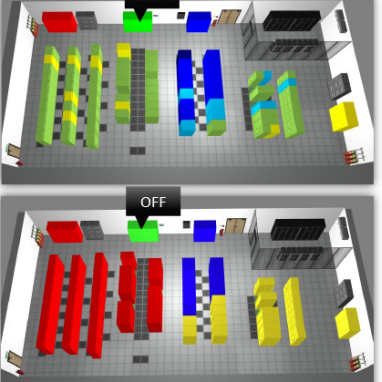
OPTIMIZATION ADVISED

taking advantage of EkkoSense’s Machine Learning algorithms and its Zones of Influence approach means that data center teams can be even more responsive to change. If, for example, a cooling unit is switched off, the EkkoSense model will recognise the change, adapt at the next iteration following the change, and then take a few

iterations to return a new cooling zone configuration. In this first example, while the rack temperatures are within ASHRAE guidelines, the blue cooling unit is not impacting many racks. It’s a clear example of too much airflow being provided into a space, hence the optimization advice.



Adapt to Operational Changes – After



- ✘ When a unit is switched off, the clustering profiles update for all AHUs
- ✘ Switching off the GREEN AHU has the following effects:
 - ✘ Less airflow allows the BLUE AHU to provide more cooling to the main #
 - ✘ Process takes 1-4 hours – shorter time would benefit resilience tests and automation

Examples


STAGE 1 OPTIMIZATION COMPLETE

In acting on this advice, the green cooling unit at the top is put into standby mode (see below) which has several interesting effects from a Zones of Influence perspective. Firstly, on the left-hand side, the temperature distribution of the IT equipment becomes more evenly spread. While on the right-hand side, average rack temperature has lifted for the smaller row and passive equipment – but in an area of the room that’s less thermally at risk. However, the most obvious change is in the center aisle, which now has a lower but more consistent inlet temperature. The cooling zones clearly show an advancement of the blue zone, as well as the red zone expanding to fill the area previously being cooled by the green cooling unit.

Air, being one of the poorest conductors of heat, gives a very long latency time to be able to validate the changes and align the racks to the appropriate zones – even in cases where the changes to a particular rack can be noticed quite early. Currently for the EkkoSense Machine Learning algorithm to adjust to this change takes between one and four hours – depending on the complexity of the room and the rate of change within the data center. It is this, combined with the impact that a cooling change makes when it has to ‘distribute’ a balance throughout the remaining operational cooling units that takes the most time – typically around one to two hours.

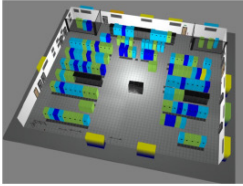


Additional cooling zone benefits that can be derived from cluster monitoring



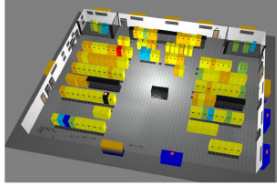
Benefit – Resilience Testing

Examples

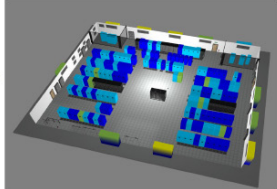


REDUNDANCY TEST

OPTIMIZATION




- ✂ Clustering can enable targeted resiliency testing
- ✂ Output for resilience test helps recommender systems to improve optimization config



Adopting the EkkoSense Zones of Influence approach can also deliver key benefits in terms of resilience testing, weighted resilience testing and enhanced capacity management.

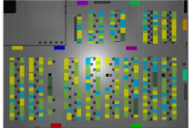
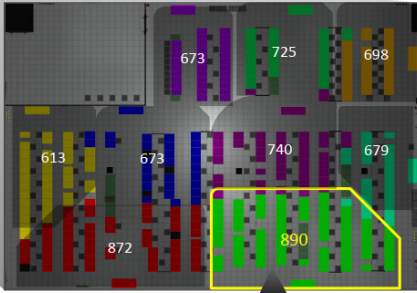
Cooling zones can provide very effective redundancy and resilience testing. These tests are time-consuming, and being able to identify and target key units provides a much higher degree of confidence in the ongoing resilience of the cooling plant within a facility.

Cooling units can also be given weighted scores to identify those cooling units that are carrying more load and/or serving a larger population of critical equipment. An N+1 classified data center can only confidently claim N+1 if it is able to run with any of the cooling units out of operation.




Benefits – Weighted Resilience Testing

Examples

- ✂ Adaptive cooling zones combined with weighted compliance, control and cooling optimization scores

RECOMMENDED RESILIENCE TEST





Supporting enhanced capacity planning



Having cooling zones in place also highlights the local cooling capacity within those zones – meaning that any new rack or equipment placement can be sense-checked prior to deployment.


In this example, a capacity request for a new rack requiring 3kW of power and cooling is initially rejected due to insufficient cooling within the zone. Providing the model is kept up to date with existing allocated loads and future reserved loads, then the cooling zones can be

used to ensure that capacity is not breached at some later date due to iterative but unconnected changes taking place over time. In the above example, you can see that the new rack was accepted in a new location within the room. However, without the control offered via cooling zones, the decision may have been to still place the new rack in the first location but with the support of additional cooling capacity – adding both significant costs as well as additional carbon loads.

Improving the data model

Applying Machine Learning algorithms and taking advantage of Zones of Influence within the data center optimization process not only allows teams to build the connections between their racks and cooling units, but also start to see how these relationships can impact airflow within spaces.

Applying clustering algorithms clearly helps to uncover previously hidden cooling configurations and relationships within the data, however these can only be as good as the feature sets they are provided with. Adding rack power data, for example, can considerably enhance the model – making it possible to represent the full data center cooling cycle with twinned supply and return cooling zones

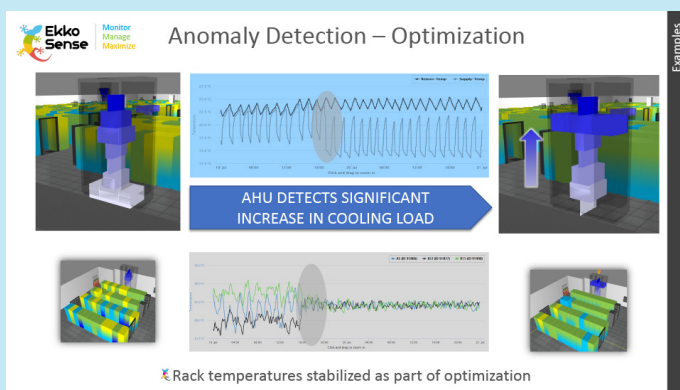




Enhance operating efficiencies by making the most of anomaly detection.

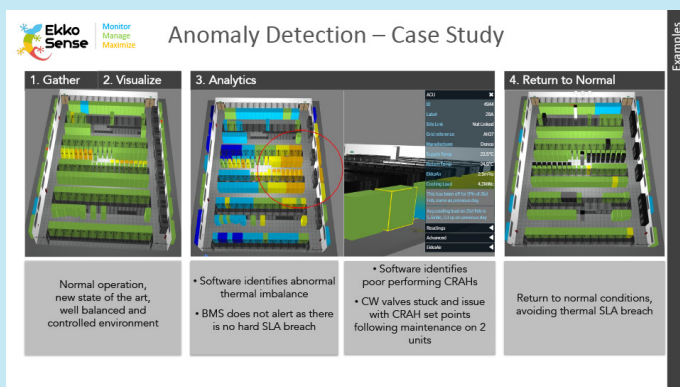
Machine Learning also proves particularly valuable in terms of uncovering potential anomalies that are hidden within more complex data sets.

In the example below, it's much easier to pick up an outlier when presented in a 2D plot and 3D representation that is completely masked in data that's presented as raw time-series charts.



In the above example we see a cooling unit and a subsection of a data center across a period of two consecutive days. At first glance things look relatively similar – some racks are hotter than others and some are colder. None of the racks are outside of compliance thresholds and the cooling unit would appear to the

engineers to be operating the same as before. Overlaying the images on to each other, however, and it starts to become apparent that the thermal 'profile' of the room has changed with a significant change in performance over the past 24 hours.



Another example demonstrates the ability to pick up on cooling unit failures. Consider the consequences of a change when an anomaly implies an inefficient or failing cooling unit that could lead to a thermal outage. Below we see how a unit is continuing with a normal, steady cooling duty profile. At some stage, the unit starts operating outside its normal parameters. Fortunately, the EkkoSense Machine Learning algorithm detects this instability as it becomes unstable and increased cooling load is detected on nearby cooling units before the original unit fails completely.

In this case the BMS would not have alerted as the cooling unit failure was not component-based and wouldn't have triggered a hard SLA breach. While the nearby cooling unit continues to cool, it would ultimately be unable to prevent a localised hot spot. This failure was picked up thanks to a combination of the anomaly detection on the cooling units combined with granular level of sensors at rack level, allowing the faulty unit to be repaired quickly and brought back online under normal operating conditions.



Turning Machine Learning insights into AI-powered actions – Cooling Advisor.

We've seen how comprehensive monitoring combined with Machine Learning algorithms can uncover and continually track the cooling zones within a data center – presenting a core foundation for effective automated cooling within critical facilities.

However, there's still considerable concern and reluctance when it comes to handing over control to what is essentially a Machine Learning powered black box.

To avoid this concern, EkkoSense believes it's important to concentrate on making the expertise behind our Machine Learning algorithms as intuitive and accessible as possible. Clear 3D visualizations and digital twin representations play a key role here, but it's also necessary to provide data center teams with the human auditability that's so important to them.

The 1st embedded advisory solution powered by AI and Machine Learning analytics.


That's where Cooling Advisor fits in, providing data center operations teams with focused power, space and cooling performance recommendations and advisory actions. Built right into the heart of our EkkoSoft® Critical software, Cooling Advisor is the industry's first embedded advisory solution. It helps organisations to unlock significant data center optimization benefits – including significant cooling energy savings – just by acting on the software's advice alone.

Cooling Advisor's actionable changes – such as suggested optimum cooling unit setpoint adjustments, fan speed points and standby settings, changes to floor grille layouts, checking that cooling units are running to specification as well as advice on optimum rack locations – are based on EkkoSoft Critical's powerful AI and Machine Learning analytics.

Recommendations are presented each time for human auditability before data center operations team members make the suggested changes. They can then use EkkoSense's data center performance optimization solution to confirm that Cooling Advisor recommendations are delivering the expected results. Cooling Advisor operation is particularly intuitive, with potential risk mitigated by the defining of clear action steps, the provision of obvious back-out mechanisms, as well as the ability to flag and unflag items so that optimization suggestions aren't repeatedly given for changes that cannot be implemented.

With the latest EkkoSoft Critical software-based data center optimization solution already unlocking data center cooling energy savings of up to 30% per annum, Cooling Advisor goes one step further, with clear recommended actions that take advantage of EkkoSense's embedded PhD-level optimization expertise in the form of cooling unit algorithms updated by learnings to equip data center teams with a powerful self-optimization capability.

By adopting analytics and Machine Learning, EkkoSense ensures that Cooling Advisor delivers advice that is specific to each data center room. This approach recognises that data centers never stay the same. Using Cooling Advisor and following its recommendations will allow operations teams to unlock up to 20% savings on their cooling costs.





Five key takeaways for Machine Learning and AI-powered software-based optimization.

The last three months have seen an increased focus on Machine Learning and AI, particularly as large language model generative AI tools such as ChatGPT have made the power of AI much more accessible.

Generative AI opens up huge potential, but its potential benefits remain largely inaccessible for a data center market that has still to take advantage of the latest Machine Learning and AI-powered data center optimization approaches.

The good news, however, is that applying Machine Learning and AI-powered software-based optimization can provide significant benefits and a rapid ROI for data center teams that need new levels of optimization support if they are to deliver against their reliability and efficiency goals.

AI-powered data center optimization approaches can certainly make a difference by cutting cooling energy and usage – with results achievable within just weeks. Bringing granular data to the forefront of their optimization plans, data center teams have already been able to not only remove thermal and power risk, but also secure reductions in cooling energy consumption costs and carbon emissions of up to 20%. It's hard to ignore the impact these kind of savings can have.

EkkoSense's 5 key takeaways for Machine Learning and AI-powered optimization:

Low-cost IOT sensors, a scalable cloud infrastructure and powerful communications all now combine to facilitate the crunching of complex data sets to support real-time optimization decisions.

Tangible benefits of risk reduction, energy savings and capacity release are areas today where AI & ML technology are already delivering.

Granularity of data for thermal and power is crucial to effectively drive out thermal risk and optimize for efficiency, placing an increased emphasis on rack-level sensing across critical facilities

With Zones of Influence and Machine Learning algorithms in place, it's now possible to predict worrying trends and highlight poor performance – well before hard faults are able to trigger traditional BMS alerts.

As Machine Learning and AI-enabled optimization platforms become the norm – and even more pervasive – it's important that organizations look for solutions that not only offer tangible optimization recommendations, but that also provide data center operations teams with the ability to both support increased workloads and deliver quantifiable carbon reductions.



Optimizing data center performance with EkkoSense.

EkkoSense is a global leader in the provision of software-driven thermal optimization solutions for critical live environments.

With its powerful EkkoSoft Critical SaaS 3D visualization and analytics solution for data centers, EkkoSense is making it even easier for data center operations teams to collect granular real-time data, visualize airflow management improvements, manage complex capacity decisions, and quickly highlight any worrying trends in cooling performance.

The key difference with the EkkoSense approach is that the solutions not only pick up the problems or underlying negative trends but also suggest best practice solutions based on EkkoSoft Critical's extensive knowledge base and deep analytics capability. This effectively removes data center thermal risks and provides 100% rack-level ASHRAE thermal compliance.

All this comes at a fraction of the cost of more expensive and complex legacy data center DCIM or CFD solutions and offers a genuine ROI of less than 12 months in most cases.

EkkoSense has already helped its clients to reduce their cooling power-related emissions with quantifiable carbon reductions contributing directly to corporate ESG and net zero sustainability targets.

EkkoSense has already helped its clients to reduce their cooling power-related emissions with quantifiable carbon reductions contributing directly to corporate ESG and net zero sustainability targets.

