Life Is On | Schneider Electric

Transforming Data Center Services: AI-Driven Condition-Based Maintenance

Enhancing resilience, reducing costs, and boosting efficiency amid skilled talent shortages

By Sophie Brause, Cristene Gonzalez-Wertz, Natasha Nelson and Wendi Runyon

The business case for AI-driven Condition-Based Maintenance

In an era where data centers face tremendous growth and mounting operational complexity, the adoption of systemic Condition-Based Maintenance (CBM) is no longer optional – it is essential.

The shift from reactive, calendar-based maintenance to AI-driven Condition-Based Maintenance not only addresses the acute shortages of skilled technicians but also delivers measurable benefits across efficiency, resilience, and sustainability.

> “In 2024, half of operators (51%) reported difficulty in finding qualified candidates to fill their job openings, for the third year running. Given the industry’s growth and planned capacity, a persistent shortage of skilled staff may be another factor impeding data center construction of facilities because they do not have sufficient staff to operate them effectively.[1] ”

Source: Uptime Institute survey

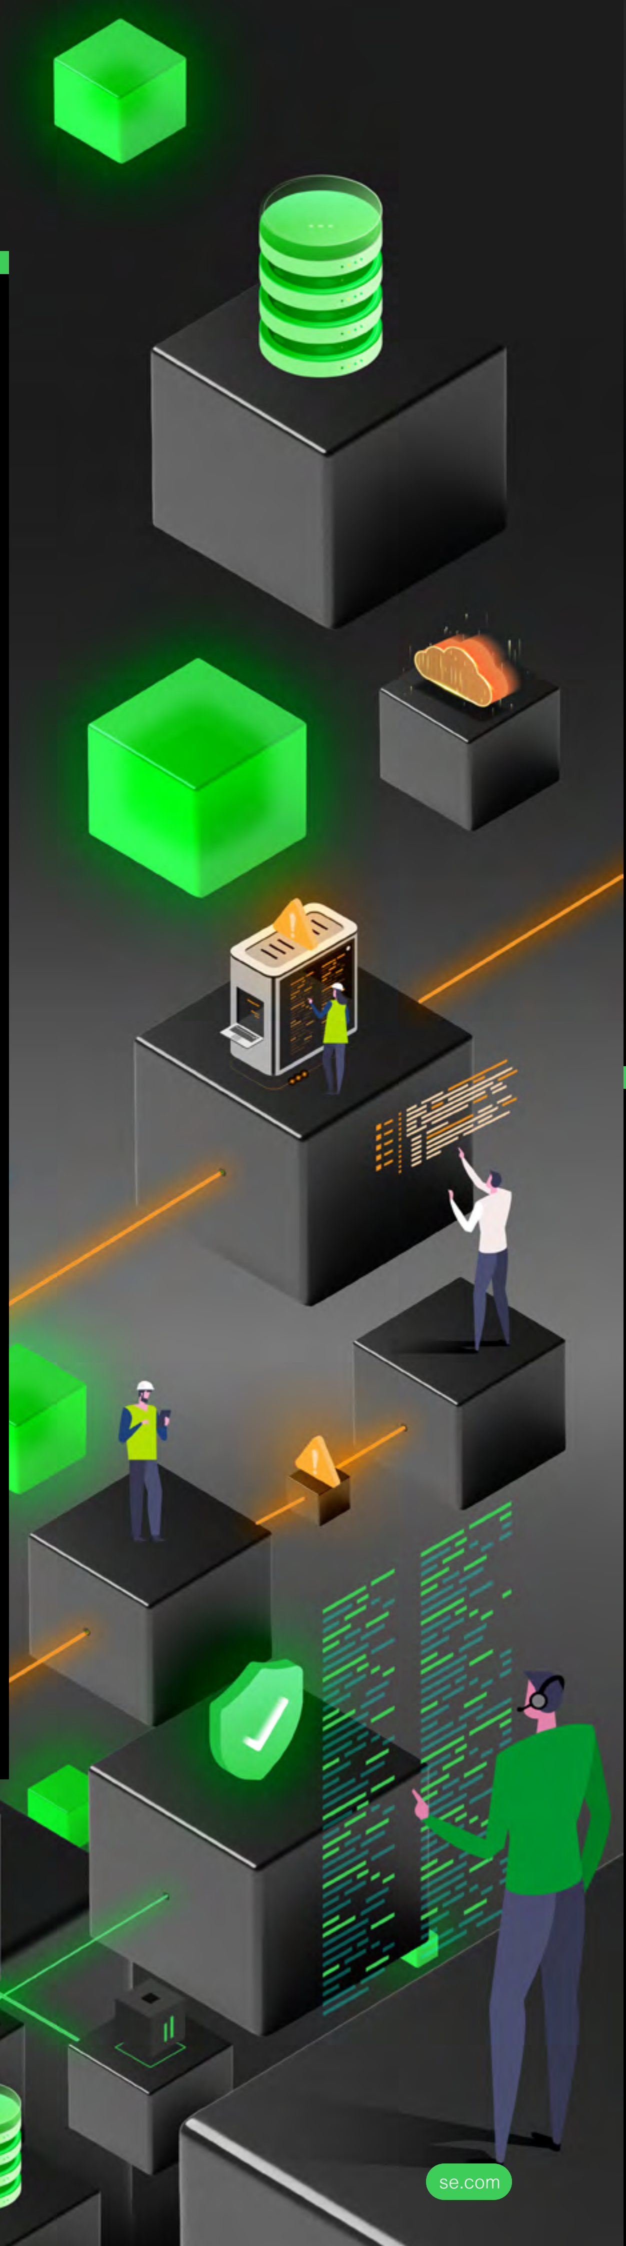

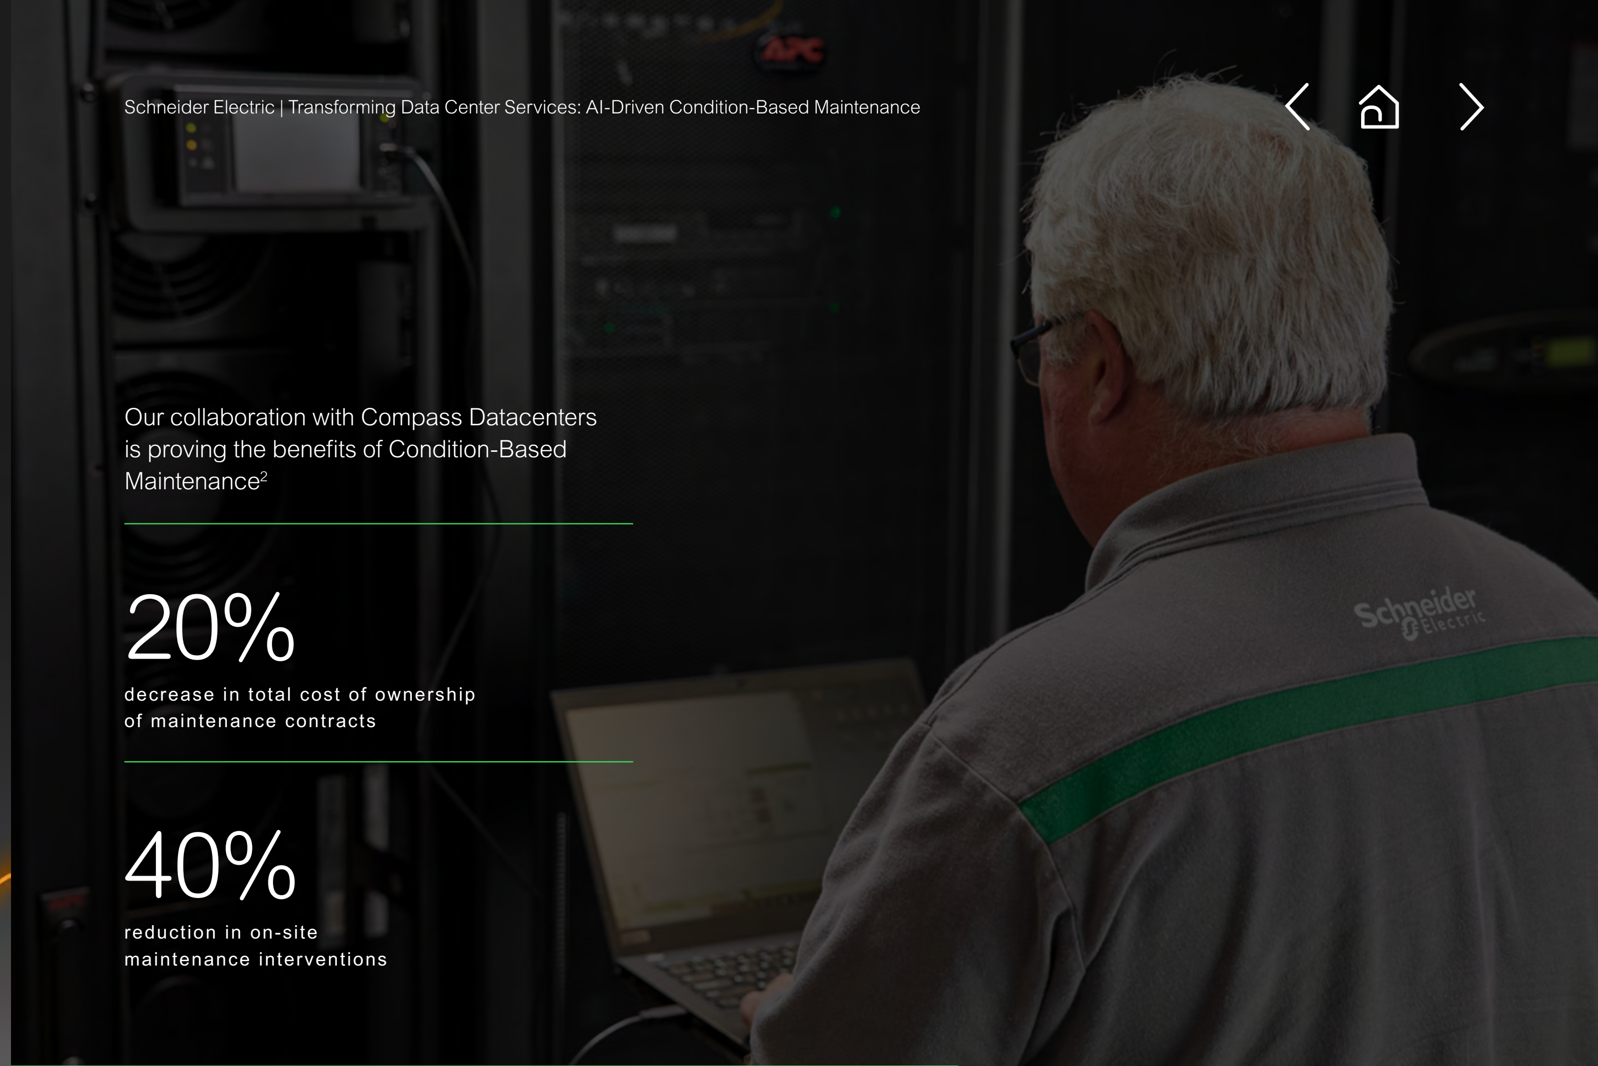

Our collaboration with Compass Datacenters is proving the benefits of Condition-Based Maintenance[2]

20%
decrease in total cost of ownership of maintenance contracts

40%
reduction in on-site maintenance interventions

AI-driven CBM

Systemic Condition-Based Maintenance attributes

- Integrates multiple data streams across electrical and mechanical infrastructures as well as external environmental data to examine overall data center systems health
- Embeds AI, advanced analytics and digital twin technology alongside human experience and ingenuity to service as an early warning system for the data centers
- Enhances risk detection, response and prioritization
- Migrates from calendar to Condition-Based Maintenance
- Generates continuous learning opportunities from captured insights
- Productivity, efficiency and cost savings

AI-driven CBM optimizes energy consumption, extends equipment life, and supports compliance with sustainability mandates.

CBM analyzes real-time and historical data from assets across power systems, cooling, and IT loads, providing a holistic view that enhances risk detection and response. Understanding the interactions and interdependencies across systems moves optimization from isolated fixes to a “systemic” approach. The strongest systemic service platforms embed **advanced analytics**, **AI**, and **digital twin technology** to generate real-time insights, driving smarter decisions while reducing the need for reactive interventions.

The numbers tell the story: integrating Condition-Based Maintenance early – from the design phase – delivers tangible benefits; up to a 40% reduction in on-site maintenance interventions and a 20% decrease in operational costs.[3] This approach minimizes human error, reduces downtime risks, and stabilizes **Operating Expenses** (OPEX) over time – critical factors for long-term business resilience.

As boards and executives plan for 2025–2030, transitioning to AI-drivenCondition-Based Maintenance isn't just an operational improvement – it's a strategic necessity. It strengthens resilience, supports sustainability goals, and optimizes resource allocation. The path forward is clear: plan smarter, operate more efficiently, and leverage AI to future-proof operations.

Boost efficiency, cut downtime – AI-driven systemic maintenance reduces costs and downtime by 20%[4]

Productivity across systems to reduce downtime and costs

Systemic service can reduce overall operational costs by shifting from calendar-based to Condition-Based Maintenance powered by AI. The result? Lower TOTEX by optimizing maintenance interventions.

Central intelligence to improve local service delivery

Operational precision through AI-driven maintenance

Minimizing impact to operations due to human error and on-site maintenance is critical, given a shortage of qualified service technicians. Leveraging a Network Operations Center[5] – (NOC) and AI can deliver a 40% reduction of intrusive on-site maintenance activities.

Smarter, leaner, stronger – Condition-Based Maintenance optimizes energy, extends asset life, and drives insights

Optimized energy extends asset life, and drives insights

Systemic Condition-Based Maintenance delivers a threefold benefit, providing human talent with technical strength. It delivers: sustainability by tracking and optimizing energy consumption while extending usable life. It can elongate asset useful life (to improve efficiency and reduce cost). It develops continuously evolving actionable insights including recommended actions and prioritization, that drive smarter decision-making.[6]

Schneider Electric | Transforming Data Center Services: AI-Driven Condition-Based Maintenance

Plan early, perform better – designing with Condition-Based Maintenance ensures seamless monitoring from day one

Begin with the end in mind

The most effective systemic service strategies are embedded from the start. Planning for Condition-Based Maintenance during the data center design phase drives seamless integration of sensors and tools, enabling real-time performance to be "plumbed in" for monitoring from day one.

OPEX management + risk mitigation = better operational control

Risk management and future proofing

Condition-Based Maintenance establishes a comprehensive and continuous learning framework for risk management *in and across* data centers. By enabling early detection of potential equipment failures, it strengthens resilience and stability. AI-driven Condition-Based Maintenance allows companies to stabilize OPEX over the life of the contract. **During transitions to new technologies or new energy sources, CBM can be a hedge against risk.**

AI improvements for data center operations

Predictive, Machine Learning (ML) and Generative AI

Creates a digital twin for every asset and its components using classic AI techniques and advanced analytics. This allows us to compare theoretical performance with actual performance based on streaming data, identifying anomalies. Our technicians utilize Retrieval-Augmented Generation pre-trained Large Language Models (LLM) for precise recommendations on addressing system issues efficiently. They also provide feedback to enhance the models continuously.

Navigating operational complexity

It's the people. More specifically, it's the lack of available technicians to maintain complex, expensive and mission-critical data center equipment. While each geography has slightly different needs – mechanical, electrical, or design, junior or senior talent – rapid data center industry growth places a continued strain on resources. And those resources are supporting equipment that grows more integrated – representing connected, dependent precision systems across multi-million or even billion-dollar facilities. **All of this amplifies risk.**

Schneider Electric | Transforming Data Center Services: AI-Driven Condition-Based Maintenance

The problem continues as companies try to collect and maintain technical and industrial knowledge across the service teams. Given technician shortages and the speed of technology change, keeping technical knowledge updated is not a trivial task. When a technician isn't familiar with a piece of equipment, it introduces more risk. Just how big a risk is it? According to Uptime Institute estimates based on 25 years of data, human error plays a role in two-thirds to four-fifths of all outages.[7] And if it were only about technician shortages, **it might be manageable – but it's not.**

Complexity compounds as data centers grow in physical size, integrate new technologies, and expand across multiple regions, creating management challenges that escalate non-linearly. Consider the complexity of data center equipment and assets. Then there's the need for new (non-traditional) data center layouts. We need to consider the geographical spread of assets a technician might need to cover. Then we should add in power sourcing and location-specific challenges. Operators are trying to manage growth in complexity without increasing costs or vulnerabilities. As the amount of equipment types, variants (models and generations), location, servicing strategies, techs, and other elements are considered, the challenge quickly becomes unmanageable.

This complexity carries tangible **operational expenses**. It drives up **maintenance challenges**, often reactive, inefficient, and prone to human error. It increases **risk exposure**, hard to quantify but potentially catastrophic if unmanaged.

Yet we have a solution. **AI thrives in solving these competing optimization problems and some companies are pointing it at data center maintenance and operations.** Even more, some software and service teams are helping the industry shift data center maintenance and management from reactive to proactive, from isolated problem-solving to holistic, systems-wide optimization.

66%

of data center operators have experienced downtime due to human error

Source: Uptime Institute

The transformational role of AI

AI isn't just a tool – it's a catalyst for transformation.

AI powers continuous monitoring, predictive maintenance, and dynamic risk management, converting data overload into actionable insights.

In fact, AI helps data centers escape the "curse" of too much data and not enough problem-solving capacity. It supports root cause analysis, helps predict failures, diagnoses potential issues, and recommends proactive interventions while optimizing spare parts management. All to reduce future risk.

The conversation around AI is often about its use in customer-facing apps. In fact, the global data center market is projected to grow from $256 billion in 2024 to $775 billion[8] by 2034, with a CAGR of 11.72%.

But growth of this magnitude demands **infrastructure that's efficient, resilient, and future-ready**.

Change at this level can create cascading operational challenges – maximizing uptime, maintaining equipment addressing densities and equipment, executing on modernization needs – the list goes on. All while the industry is experiencing a workforce shortage. Each of these challenges shares dependencies with the others. Consider the cross-pollination problem in just a few competing challenges.

Replacing equipment early maximizes uptime but may increase overall costs

Modernization increases functionality and capacity but introduces new maintenance needs

Higher rack densities necessitate advanced cooling solutions and specialized technical skills

There are plenty of others. And that's where the problem sets in. Much of the maintenance in data centers is operating under three sub-optimal principles:

- **01.** Many maintenance organizations are still using calendar-based maintenance. Equipment has a service date and that is when things get done. Earlier, if there is a failure but by a specific date if not.
- **02.** Many equipment contracts cover a thin slice of the technology. The software and services used to manage equipment in the data center are still centered at the individual asset level, not at a systems-wide (systemic) level.
- **03.** Few vendors' services are equipment agnostic. Most companies and technicians are constrained by principle 2 (above). They lack the understanding, connectivity, analytics and manpower to cover the equipment from other providers or across the data center ecosystem.

Calendar-based service can remove assets before expending their usable life. When scheduled maintenance occurs, it's relying on historical patterns, which is beneficial but not necessarily always reliable. Equipment wear may be influenced by many specific local conditions and the systems or environment within which that asset operates.

Even when considering specific equipment for targeted replacement, what criteria should be used? Replacing a UPS after 8 years instead of 10 could increase reliability – but at what cost? CBM helps determine the optimal replacement window based on real-world operating conditions, rather than generic lifespan estimates. For the batteries within a UPS, how should environment, backup times, and battery composition factor in? These sorts of optimization problems are where AI brings a distinct advantage to servicing.

To reduce risk, extend asset uptime, and optimize operations, a diagnostic in a data center shouldn't just assess one piece or group of equipment. It should expand systemically to include multiple assets. Consider how different servicing might look if you assessed the generator or utility source (including the generator power meter and breaker) plus the utility power meter and breaker as individual assets and as groups of connected assets. Then you could add additional equipment or asset classes as you extend into to the white space. AI thrives in solving these kinds of competing optimization problems.

The path forward

What's needed goes beyond a maintenance strategy – data centers need a blueprint and approach that delivers a reliable, resilient, and efficient operation. By embedding AI-driven Condition-Based Maintenance into the fabric of infrastructure management, organizations can reduce costs, manage risks, and enhance sustainability.

The time to act is now.
Plan smarter. Operate better.
Lead with resilience.

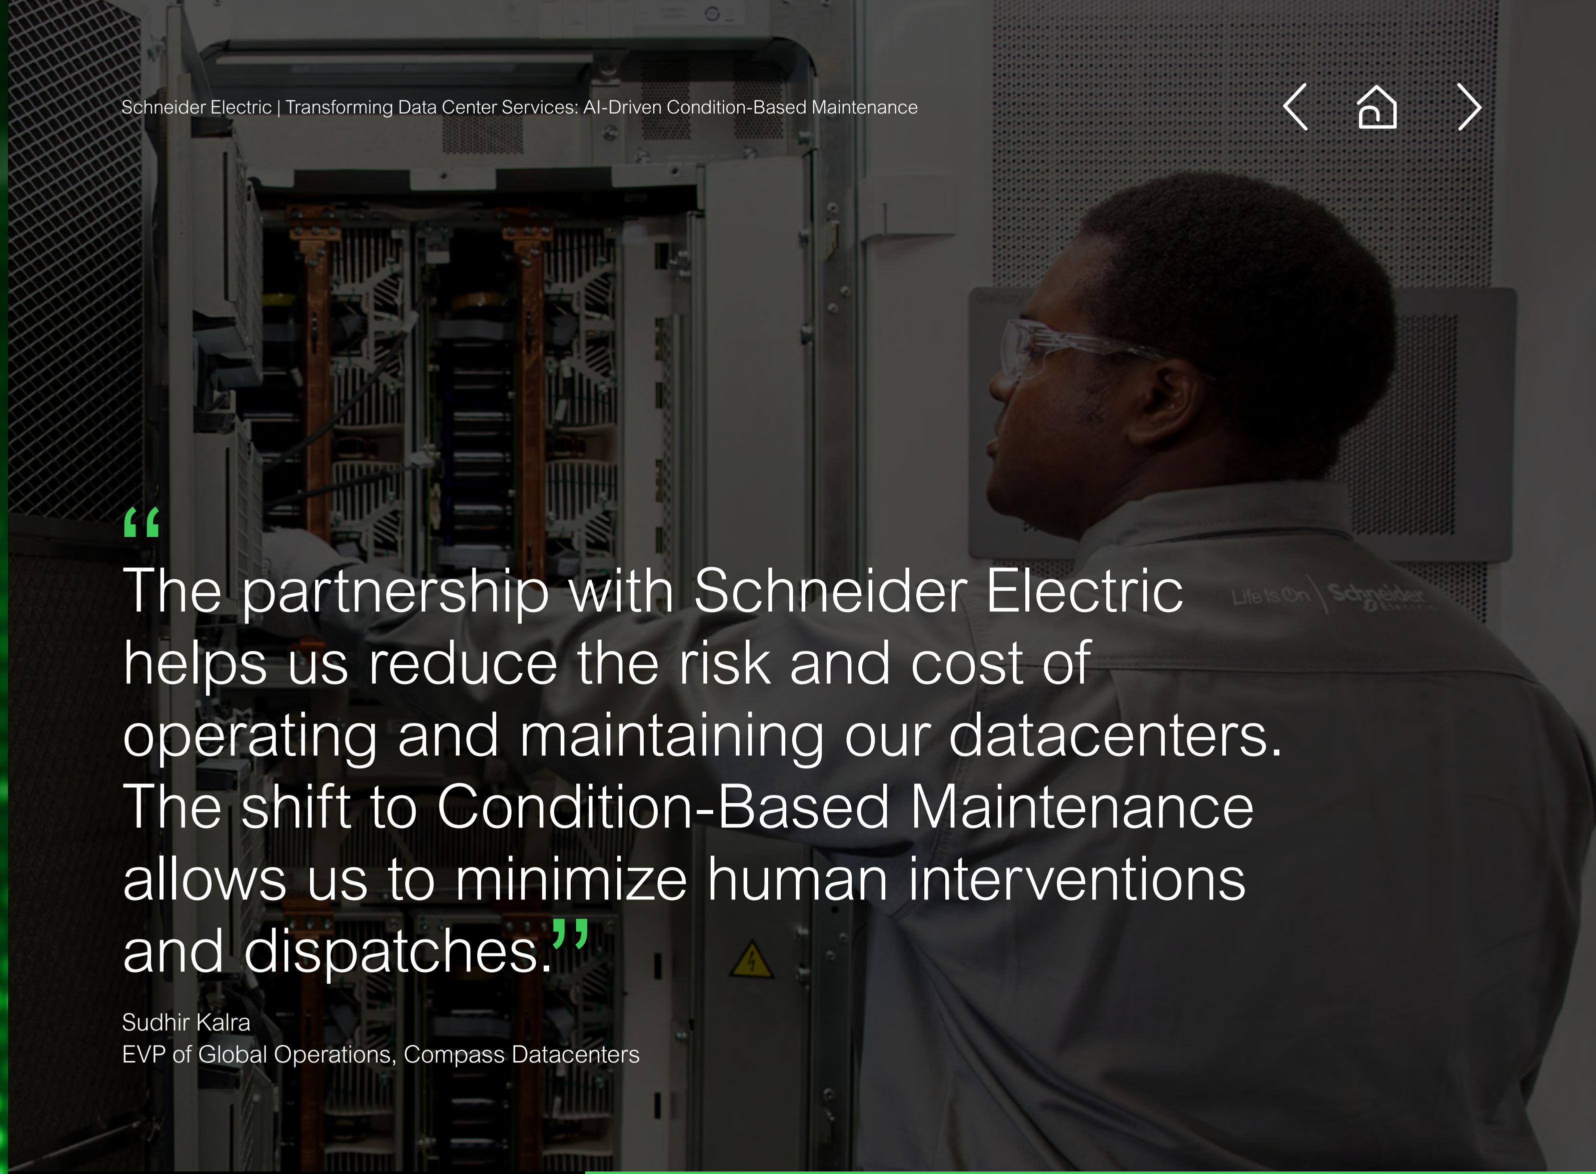

"The partnership with Schneider Electric helps us reduce the risk and cost of operating and maintaining our datacenters. The shift to Condition-Based Maintenance allows us to minimize human interventions and dispatches."

Sudhir Kalra
EVP of Global Operations, Compass Datacenters

From calendar-based to Condition-Based Maintenance – and into the age of systems-wide servicing

We're witnessing a major shift in data center operations as we move from calendar-based maintenance to a condition-based approach. The old model, where interventions are scheduled based on time rather than need, often results in wasted resources or, worse, missed warning signs that lead to unexpected failures.

Monitoring changes in conditions allows a data center to rely on real-time data from sensors to assess the health of critical assets and trigger action only when it's required. It's a significant leap forward.

But it only addresses part of the challenge. Most players in data center ecosystem will ask the original equipment manufacturer (OEM) to monitor and service their respective equipment: "we will monitor all of your liquid cooling assets to deliver optimal performance." This is good, as it places the responsibility for those assets with the organizations that know them best. However, it has a significant downside too.

When each vendor provides support for only their components – power and cooling, for instance – this maintains operating silos. It adds complexity and requires a data center to provide access to each vendor. When this occurs, the data center not only misses out on synergies across the systems or the data center, but also injects risk by allowing more humans into the environment.

At Schneider Electric, we know by experience that customers find value in a sole source relationship for service when possible.

That was the situation in our work with Compass Datacenters, with over 17 active data centers, growing to 15,000 assets across the world live. They saw operational benefits in a uniform contract and service across all of their sites (not a regional/local service per location). It allows them to derisk and increase operational efficiency – reducing the number vendors and SLAs (service level agreements) to onboard and manage.

With one point of contact to manage maintenance and service, centralize work orders, knowledge management, parts management and deployment, they can build better "systemic" understanding and expertise across their chosen equipment, and build deeper relationships with one provider who can prioritize them as a client.

Cybersecurity and Data Integrity in Condition-Based Maintenance (CBM)

With IoT and advanced monitoring in CBM, cybersecurity is non-negotiable. Focus on securing data streams and system integrity to maintain reliable operations.

Actionable steps:

1. Build security into CBM from the ground up and top down:
Apply security protocols from data collection to analysis.

2. Run continuous risk audits to test everything:
Identify vulnerabilities and adapt quickly. Data centers are critical infrastructure, down to each sensor.

3. Engage and prepare staff for a "security first" mindset:
People are both the biggest risk and strongest defense. Lack of awareness can expose vulnerabilities, while informed teams can prevent breaches.

4. Use real-time monitoring and AI to ward off cyber threats:
Detect anomalies and respond immediately. Speed to response is mission critical. Strong cybersecurity in CBM protects data accuracy, sharpens decision-making, and cuts downtime. The result: resilient, efficient and protected data center operations.

Schneider Electric | Transforming Data Center Services: AI-Driven Condition-Based Maintenance

Data center stakeholders may find themselves using different ways to extract data exposing cybersecurity vulnerabilities. Operators often wind up with a different dashboard from each company for individual assets; and they encounter a variety of service techs coming into the data center at different times, resulting in many more human interventions than necessary. And that increases cost and risk.

Operating is the mission critical word for data centers – all this equipment must run – and run as part of the data center's system. In truth, it's a set of systems within systems.

Often, technicians are educated on their own technology – so they can maximize its performance. Yet, a technician maximizing one component might not be delivering an optimized experience across all of them running together. These components and systems run in context to each other. They are responsive, integrated and "collaborative." Any solution needs to understand the performance of each component but also the performance of the collective components. And each component's performance is generating a veritable ton of data.

The real game-changer here is AI. By processing these enormous streams of data, AI identifies performance patterns humans can't, helps predict failures before they happen, and even automates parts of the decision-making process by delivering outcomes. Yet, to build these AI models, a solution needs more than access to the system level data, it needs subject matter expertise that can validate and contextualize the models.

An example may be helpful. Consider Switchgear Thermal Analysis. This AI analytics model determines the relationship between thermal data from UPS, switchgear, and container to model the performance of switchgear and compare across a population.

In the work Schneider Electric is doing for Compass Datacenters, we currently compare 10 switchgear in a location. Using this AI model, we can determine which switchgear will be likely to require service. Recently, we found that only 3 out of the 10 switchgear should be serviced. Using this model, we prevented shutdown of 7 Power Centers and a number of data centers, maintaining uptime with no loss of reliability.

In another example, we leverage our AI model for Compressor Input Power Predictor (CIPP). This CIPP model estimates the likelihood of issues including coil and filter problems, fan failure, refrigerant concerns, and more. This AI-based analytics model enables us to detect issues with cooling units before they happen, so we can maintain performance of the cooling, UPS, switchgear and batteries all at once.

FIGURE 1: SYSTEMIC CBM COVERS THE ENTIRE SINGLE LINE DIAGRAM

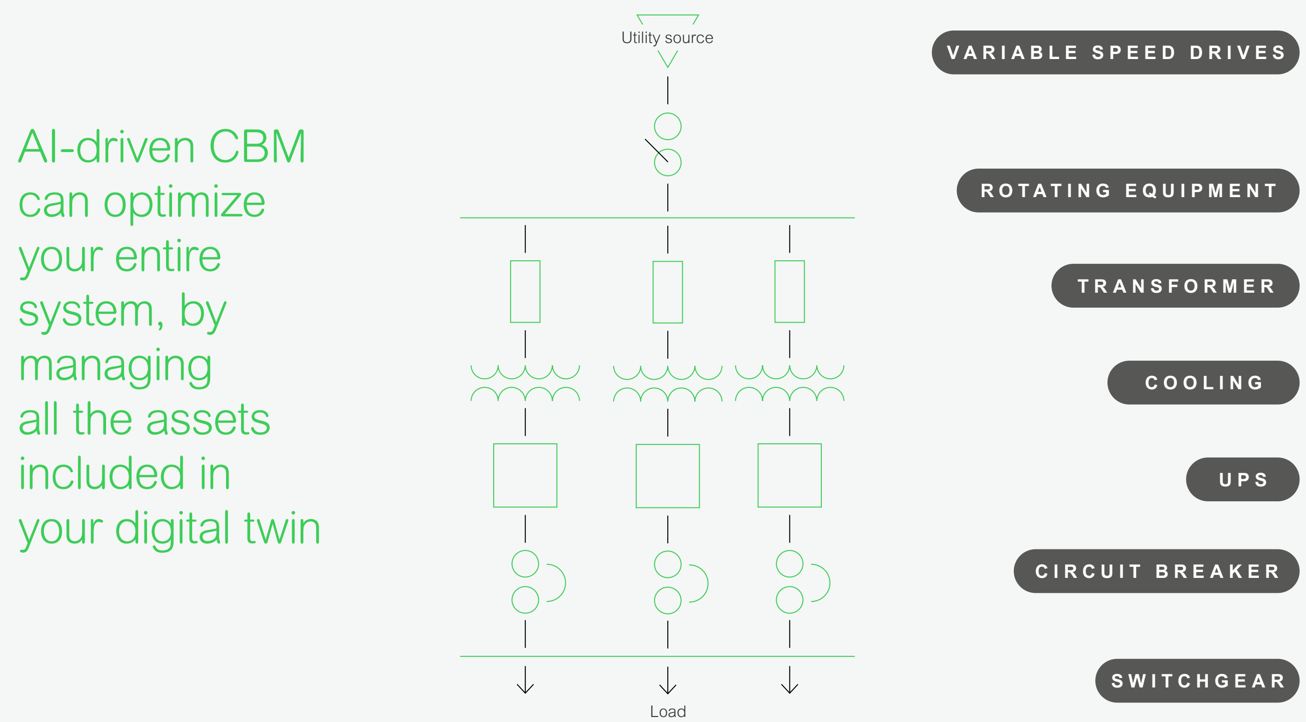

Figure 1: Systemic CBM covers the entire single-line diagram.

Condition-Based Maintenance harnesses the untapped potential of continuous learning

Systemic service and Condition-Based Maintenance don't just react to data – they learn from it. AI systems embedded in Condition-Based Maintenance platforms continuously analyze operational patterns, identifying inefficiencies, and offering actionable recommendations.

This means that over time, data center operations not only run more smoothly, they improve themselves. For example, by monitoring cooling systems, AI can suggest adjustments to airflow or temperature settings, optimizing energy consumption and prolonging equipment life. These insights are refined with every new data point, creating a virtuous cycle of operational efficiency.

There is also human feedback in the loop. It's vital that insights and feedback from the Subject Matter Experts (SMEs) and site management contribute human observations of the problems. They bring ingenuity into the solutions, and work alongside each other – bi-directionally.

AI combined with SME knowledge develops an accurate and precise AI model that can determine when events are normal vs. abnormal to a system. AI can detect an anomaly. Skilled SMEs, like the +6,000 Schneider Electric has, can tell you what the cause and condition of the equipment may be. Combining both AI and deep "cross-the-data-center" systems knowledge by the SMEs provides an automated anomaly detection with contextual knowledge on the detection for service and/or diagnosis support.

Costs and lack of qualified staff are 1 and 2 on the list of management concerns[9] in 2024, according to the Uptime Institute

71%

Lack of qualified staff

80%

Cost

Source: Uptime Institute

Two examples can bring clarity. The first[10] determined a missed configuration during install or prior servicing. The Network Operation Center (NOC) noted battery temperature swings and saw that the Heating, Ventilation, and Air Conditioning (HVAC) units were not properly sharing cooling. The system noted the lower heat load. But the AHA! moment was with the Services Representative on-site, who discovered that the HVAC controller was set to the wrong operating mode. This incorrect operating mode can lead to the overuse of certain HVAC units and poor air circulation in the data center.

Proactive Analytics Example

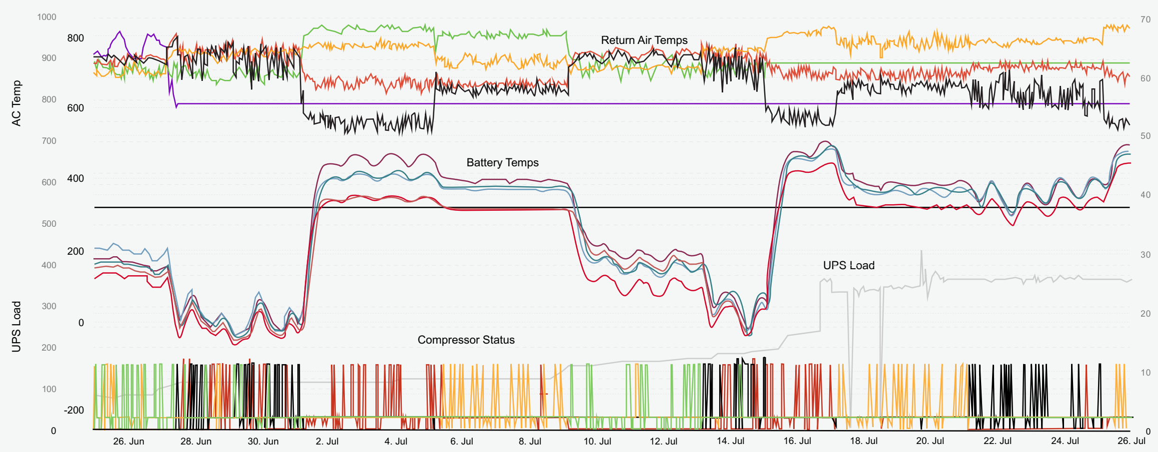

Data trends in condition

- Battery temperature swings
- HVAC units not properly sharing cooling
- Lower heat load /container lightly loaded

Conditions Identified

- HVAC controller's operating mode was improperly configured

Figure 2: Example of a proactive analytics battery diagnosis.

The SME used this event to develop an analysis that identifies improper configurations and is now included in the cooling analytics models for this client.

In fact, 2024 was the hottest year on record according to NASA[11], beating out 2023. The impacts of climate change impacts are evident. Prolonged exposure to high temperatures can wreak havoc on equipment. The data trend showed high outdoor temperatures, brought into the AI systems through a feed and combined these with frequent changes in the evaporator temperature. These were causing frequent freeze alarms.

The blower motor was considered suspicious, and could lead to frozen compressors. Combining the data architect's observations with the Services Representative's on-site experience allowed development of new analytics to detect issues before they occur, creating a tighter feedback loop for the use case.

Schneider Electric | Transforming Data Center Services: AI-Driven Condition-Based Maintenance

Analytics Example

Event Overview

- Customer lost power from utility causing its system to switch to the generator for power source
- Customer asked for analysis on the event
- Later determined that animal caused outage to occur

Data Analysis on Event

- The system was analyzed by subject matter experts before and after the transfer event to ensure it was operating normally once power switched back to the utility source
- Considered data from the batteries, UPS, switchgear, and cooling units
- Considered data including temperatures, power, breaker statuses, and time

Next steps

- Communicated to customer that the system performed a successful transfer and returned to normal operations after the event
- Develop new analytics to detect when a system returns to abnormal operation after a transfer event

Figure 3: Example of new use developed based on event that occurred on Customer site.

It takes time and knowledge to train the models. This is not something that happens overnight; it is a journey. Finding a partner who has the data AND domain expertise is critical to success.

Through our analytics, voice of customer, and analysis of problems occurring at sites since collecting data, we determine where our new analytics models can benefit the client and the equipment manufacturers. As an example: there were potential reports of a manufacturing defect in the HVAC units. Since there are currently over 1,000+ HVAC units on site, it would be costly to service and replace all. We developed an analytic that would determine the likelihood of encountering the defect to better identify when it might be present.

Another example: we learned that the design of the data center could affect the airflow and subsequently, the temperature of battery and UPS. Using AI and the Services Representative experience, we developed an analysis to identify air circulation problems, with a possibility to mediate (like changing HVAC configuration), and an alert to the NOC.

Responding to the events and data together – with both trends and insights – allow SMEs to identify where AI and prediction models can prove valuable in detecting issues with the data center system. It's not an either/or. It's definitely a both.

Future proof

Creating a future that reduces data center operations risk

As new technology comes into the data center, new analytical models are needed. These analytics need to be event driven – based on physical elements, indices of performance and composite health scores for individual equipment components across vendors and for specific use cases. To develop each analytical model, new properties are identified and incorporated. Then the model is tested on the live data.

Schneider Electric is advancing the use of AI for assets and systems in our CBM approach. For our emerging models, we build on both physics-based and AI-based models using historical and real time data. For new models we have in incubation, we build new properties and data points that are incorporated into new analytical threads. This effort is completed by an AI team, SMEs, and our analytics team to define the logic, data points, and development on the platform.

What makes this approach particularly powerful is its ability to adapt to new technologies and environmental demands. As data centers transition to next-generation equipment, such as liquid cooling or modular power systems, Condition-Based Maintenance frameworks evolve alongside them. Similarly, during shifts to renewable energy sources or distributed edge infrastructure, systemic Condition-Based Maintenance helps avoid operational risks are minimized while performance remains consistent. This flexibility isn't just a nice-to-have – it's essential for future-proofing.

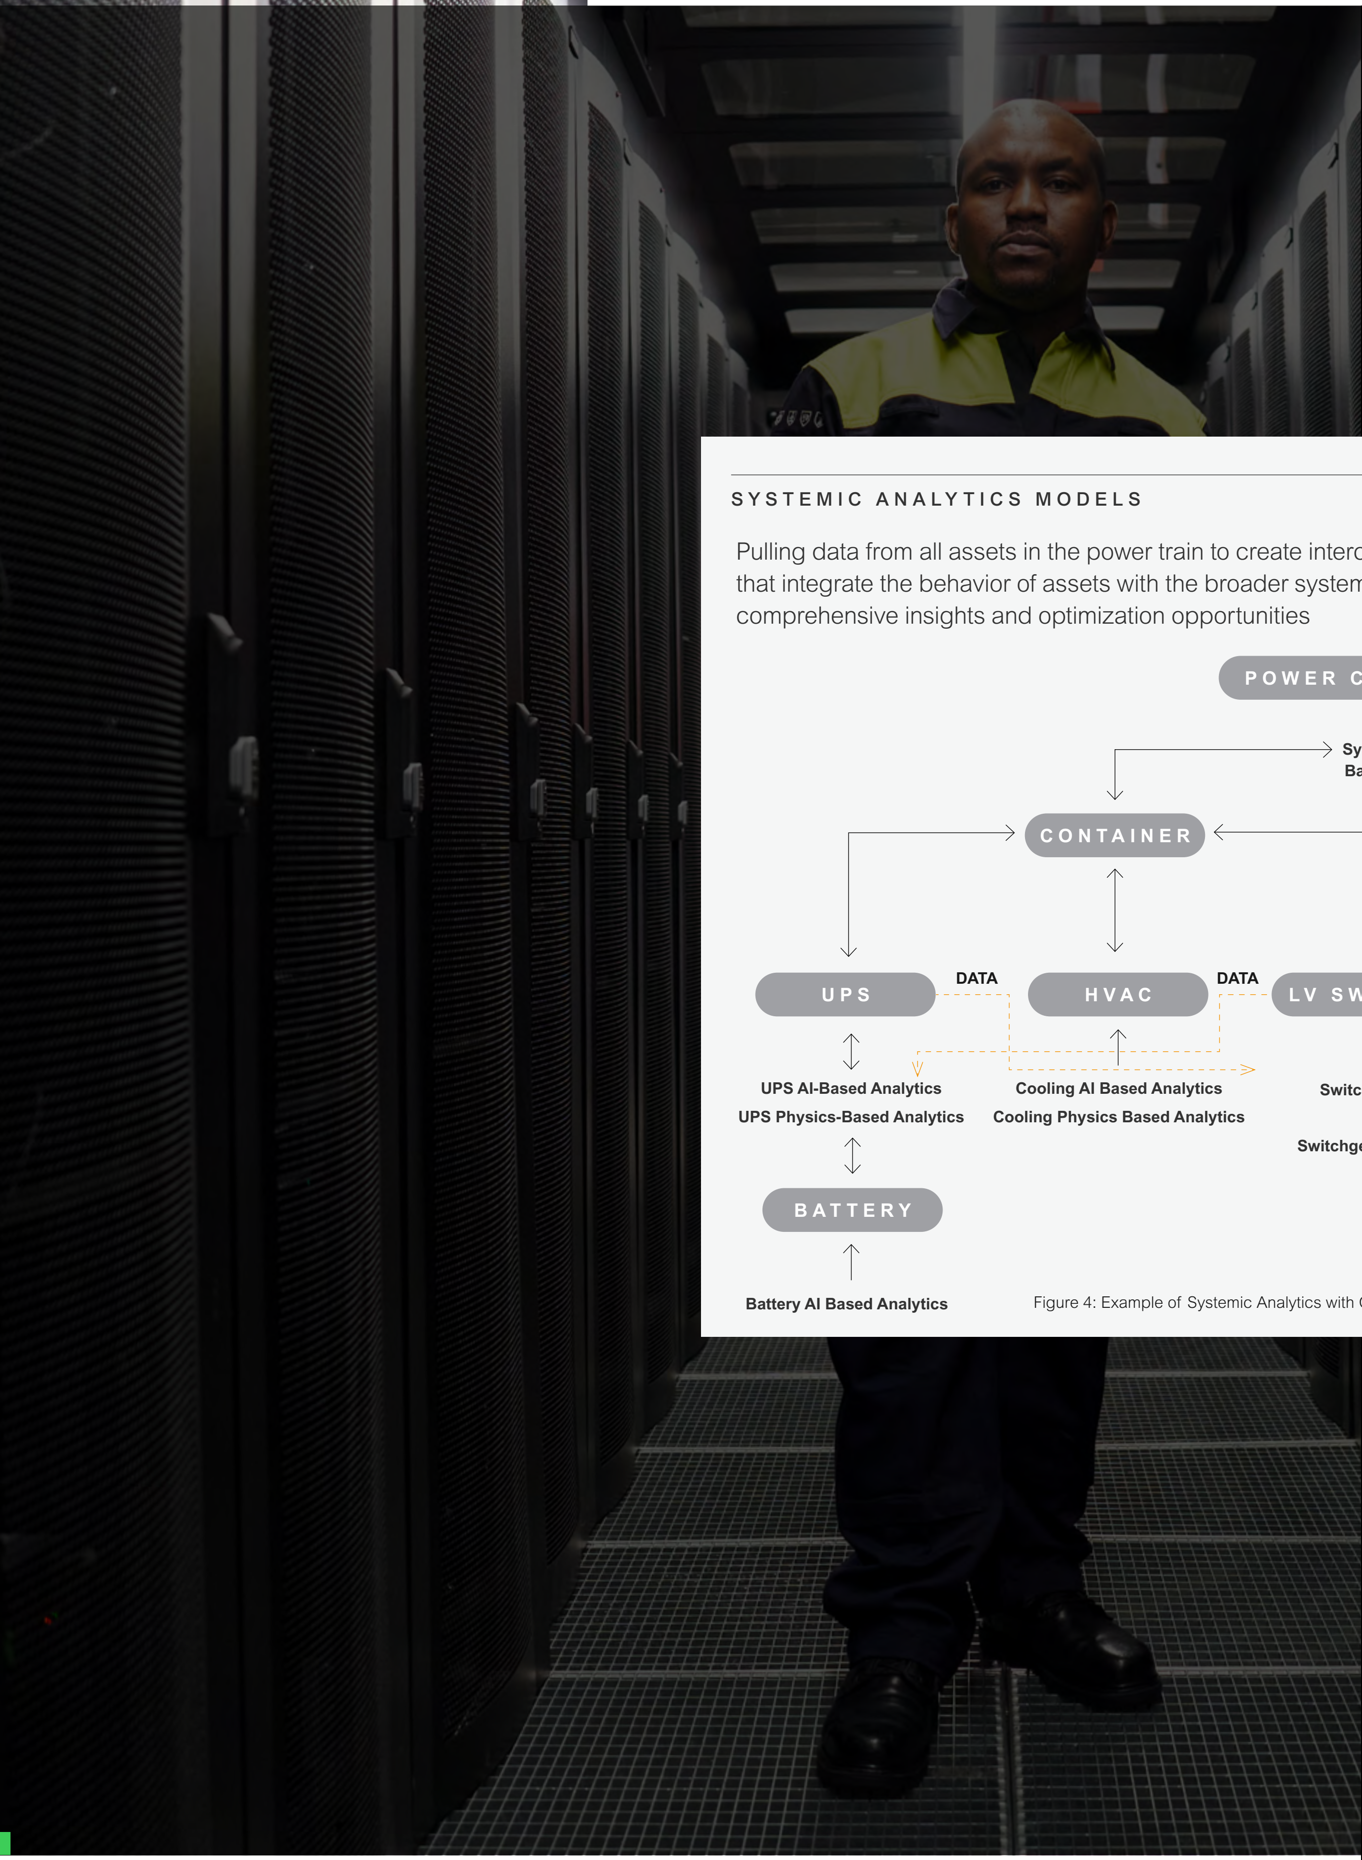

Systematic analytical models
We're not just optimizing performance – we're fundamentally rethinking how we ensure resilience and sustainability in data center operations. It's the difference between checking the calendar or having a full roadmap of the data center.

The latter is where the future lies. **This isn't just about saving costs, it's about precision, efficiency, and smarter operations.**

SYSTEMIC ANALYTICS MODELS

Pulling data from all assets in the power train to create interconnected models that integrate the behavior of assets with the broader system, leading to more comprehensive insights and optimization opportunities

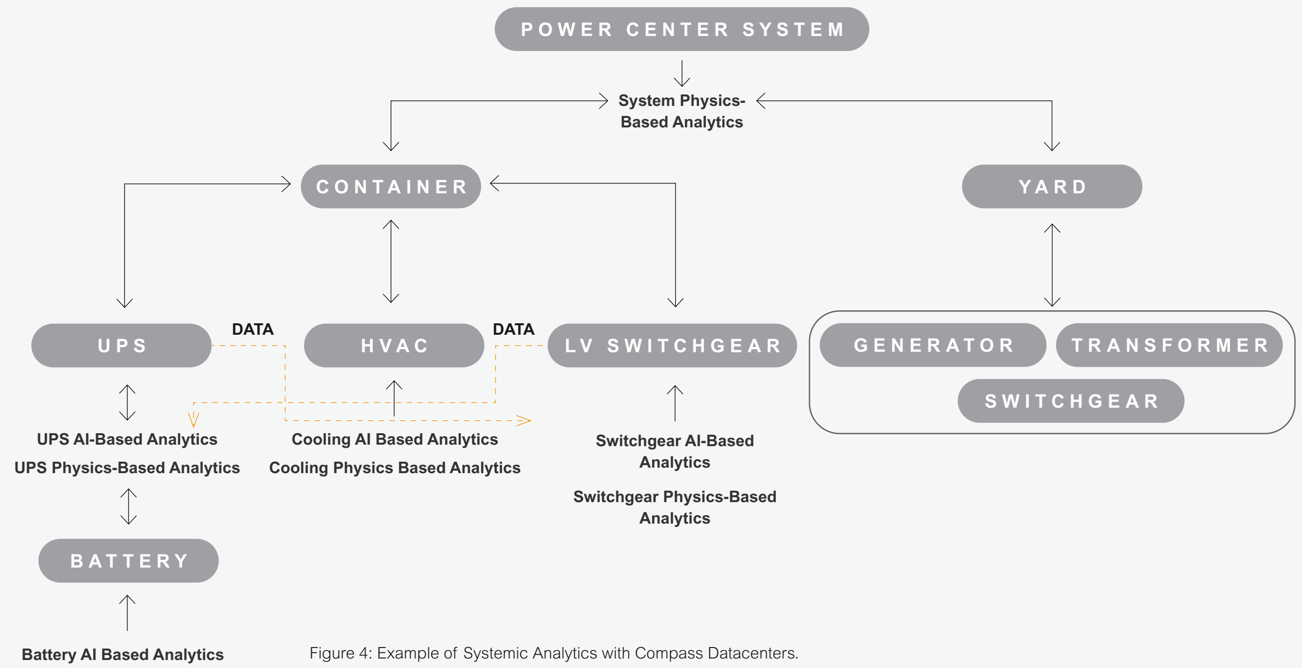

Figure 4: Example of Systemic Analytics with Compass Datacenters.

Data pulled and aggregated from all assets inside container and outside from generator enables **systemic analytics** and the development of both physics and AI-based **analytic models** that provide outcomes at both the **container and the generator**.

For example, we are pulling data from the output of the switchgear into the UPS model to provide better insights into UPS performance (and viceversa). Additionally, we are pulling data from all assets including generator into System models to provide better insights into performance of the entire system.

Schneider Electric | Transforming Data Center Services: AI-Driven Condition-Based Maintenance

Where AI makes a difference

Risk Management: Proactively identifying and mitigating threats

Real-time data + predictive analytics = proactive decision-making

At the core of Condition-Based Maintenance is the ability to **detect issues before they escalate**. AI-driven predictive analytics continuously monitor system health, identifying anomalies early to reduce downtime and prevent catastrophic failures.

With **up to 75% reduction in unplanned downtime**, significantly minimizing operational risk.

Resilience during technology transitions

Data centers are constantly evolving, whether adopting liquid cooling, edge deployments, or integrating renewables. Condition-Based Maintenance powered by AI helps manage risks during these transitions, keeping systems stable even as new technologies are introduced.

In industries like energy and manufacturing, AI has proven critical during tech shifts, reducing incidents tied to unfamiliar equipment or processes.

Insights from each asset, the whole system and people servicing it

Traditional maintenance focuses on individual assets but risks often arise from interactions between systems. AI enables a systems-wide view, identifying cross-system vulnerabilities that could lead to cascading failures.

Demonstrated success in smart grids and industrial operations, where system-level insights have led to significant risk reductions.

Operational Efficiency: Driving productivity & cost savings

Centralized operations enhance efficiency across multiple sites

Managing multiple data centers is resource-intensive. AI-powered **Network Operations Centers (NOCs) allow for remote monitoring and intervention**, reducing the need for on-site technicians and speeding up response times.

Schneider's Network Operations Center for Compass Datacenters has shown up to 40% reductions in on-site maintenance visits, directly improving OPEX efficiency.[12]

AI supports energy optimization, reducing both costs and carbon footprint

Energy is a major cost driver for data centers. AI-driven Condition-Based Maintenance optimizes energy consumption across systems, lowering both operational expenses and environmental impact.

Real-world results from Schneider's operations show up to 30% reductions in energy consumption, supporting both cost control and sustainability initiatives.[13]

Future-Proofing: Resilience for what's next

Digital twins enable scenario planning and stress testing

To future-proof operations, data centers need to model potential risks and performance scenarios. Digital twins create a virtual replica of physical systems, allowing for real-time simulations, predictive testing, and continuous improvement without disrupting live operations.

Widely adopted in aerospace, utilities, and manufacturing, digital twins have proven critical in anticipating failures and optimizing maintenance strategies before issues arise.

Expertise – from the Services Representatives and remote operations professionals – is powerful.
These teams are the unsung heroes of systemic service. Every model requires ground truth validation – comparing the data and analytics to real-time and real-world observation. They allow teams at the NOC to oversee, manage, and maintain data centers without being physically on-site. This is a game-changer in today's landscape of technician shortages and increasingly dispersed facilities. By centralizing expertise through advanced monitoring hubs and remote operations, skilled technicians can oversee multiple sites leveraging AI to handle the routine noise while focusing human effort on complex issues.

More technology at a broader scale can deliver innovation's benefits faster
This is also a place where extended reality innovations, such as EcoStruxure XR Operator Advisor, will increase in value looking ahead. Augmented, Virtual and Mixed Reality can visualize, interact, inform, train and empower operations and maintenance staff to help enhance safety, power availability, operational efficiency and sustainability, as well as improve the quality of operations and maintenance services.

Use cases include:

Training and simulation
Enhancing employee efficiency, understanding, and knowledge transfer. Allows new technicians to develop a comfort level with equipment, conditions and circumstances.

Maintenance and repair
Coordination from the NOC to the Service Representative via smart glasses, ipad or iphone, XR enables even remote service to be productive, even in low-touch or low-maintenance equipment. It increases safety and allows overlays to real-world equipment to bring "best of both" capabilities to the data center.

Quality control
Increasingly cameras are joining sensors – whether through drones or periodic monitoring to provide inspections and contribute to digital twin models. They can enhance safety and identify dangerous situations in advance, improving the data center's ability to react.

Collaboration
Improving communications between the NOC, the Services Representative and intelligent systems is mission critical for geographically distributed teams.

It's not just about cutting travel time – it's about amplifying impact. With remote operations, interventions are more targeted, decisions are faster, and the insights generated are actionable across an entire portfolio of data centers. This kind of scalability and precision isn't just efficient, it's transformative, enabling operators to do more with less while maintaining the resiliency and uptime their stakeholders demand.

In addition to generating insights used in the field, Virtual Agents are assisting the intake and diagnosis process. At Schneider Electric's Connected Services Hub (CSH), AI-powered virtual agents now assist technicians by diagnosing issues, translating service reports, and generating step-by-step maintenance recommendations – reducing service time and human error.

Application 01 – Bespoke solutions

Using generative AI, a tool can produce bespoke recommendations as a baseline for each report, based on specific site and conditions.

A specific site can generate detailed, step-by-step instructions, for example, to help a customer investigate a concern about a particular piece of equipment in their power system. The recommendations are based on monitoring information from our IoT-connected sensors, and insights from our digital analytics. And crucially, they're underpinned by a powerful Schneider Electric knowledge database, developed by our experts. To uphold accuracy, all the recommendations are reviewed and endorsed by our engineers before we send any report to a customer.

Application 02 – A new virtual expert

We've begun using a generative AI virtual expert for our staff experts to consult supporting their work. This "internal expert resource" is focused specifically on electrical asset management.

It can validate technical standards in a specific part of the world or check the specifications of a circuit breaker. The tool is trained using deep technical know-how, product documentation, white papers, instruction manuals and industry understanding, to make sure it provides useful, accurate answers.

The NOC is providing feedback for fine-tuning the virtual expert. And it cites sources to further transparency.

Application 03 – Summarizing insights more efficiently

The NOC also uses AI-powered technology to automatically summarize maintenance early warnings coming from our sensors and analytics into a digestible format, and translate this into different languages where necessary. This helps engineers get straight to the heart of relevant issues in their reports – offering authoritative insights and actionable advice.

It distills measurements and alerts into a few clear paragraphs, specific to the customer's individual context. Once the summary is produced, an expert can choose to translate, shorten or expand it at the click of a button. They then review and approve the content before incorporating it into the final report.

Data centers and the technology in them can't leverage off-the-shelf solutions – the content produced even by the latest AI would be too generic to be useful. So, we've drawn on our expertise both in data science and energy management to train equally powerful algorithms, but with a focus on our particular area of work. This means the output is accurate and specific enough to make a real difference.

Our iterative design process installs formal safeguards for accuracy as well. Last, we work out the computational power required for any use of AI before deploying it – because what's the point in a solution that uses more energy than it saves?

Our own experience running an AI-enabled NOC for Compass Datacenters using Schneider Electric's EcoCare membership services and EcoConsult Electrical Digital Twin has elevated remote operations to a whole new level. By providing a virtual representation of the entire data center ecosystem, our expert Consultants help simulate, predict, and optimize performance in real-time – without ever setting foot on-site. Thanks to the ETAP software and its seamless integration with the models, our Consultants provide actionable insights and recommended actions based on live data from power systems and cooling infrastructures. This combination transforms remote operations from monitoring to proactive management, enabling teams to test scenarios, predict equipment behavior, and make adjustments before issues escalate.

For example, using EcoConsult Electrical Digital Twin to model and fine-tune cooling strategies or preemptively address power distribution inefficiencies, reduces energy consumption and extend asset lifespans. Together, Schneider Electric EcoCare and EcoConsult Electrical Digital Twin redefine what's possible in systemic service, delivering unparalleled control and insight from anywhere. It also allows executive teams to model and control for costs, plan for equipment changeouts and develop better **Capital Expenditure** (CAPEX) and **Total Expenditure** (TOTEX) estimates.

Systemic service starts at design

The design phase is the perfect moment to bake systemic service and Condition-Based Maintenance (CBM) into your data center strategy.

It's here where you set the foundation to handle the growing complexity of modern operations – where resilience, scalability, and efficiency aren't just goals, they're non-negotiables. CBM relies on sensors, data streams, and advanced analytics to optimize performance in real time. Retrofitting these tools into an existing data center is a hassle – it is not only expensive, disruptive, and often leaves gaps in coverage but also requires equipment shutdown and human intervention. But when it's planned for during design, those sensors and tools become part of the architecture. The result? A seamless system ready to deliver real-time insights from day one.

Designing with systemic service in mind also avoids the pitfalls of fragmented monitoring. Instead of focusing on individual components, the data center obtains a complete view of how the systems interact, optimizing maintenance schedules, extending asset lifespans, and even reducing energy costs. It's not just a smarter way to operate – it's a more sustainable one.

Bottom line: planning for systemic service at the design stage means the data center is built for the future. It's proactive, cost-effective, and allows a data center operator to handle what's next with confidence.

In a greenfield vs brownfield comparison: our experience has found using just the cost of labor alone to install sensors and IoT for connectivity[14] – there is an estimated 7,000 USD in additional CAPEX per Power Center to install sensors on-site vs. in factory. This does not include the cost of equipment shutdown. Service ready kits, the sensors that go into new-build data centers, add approximately 1% in additional cost to the hardware. Adding the sensors up front prevents shutdowns and retrofits. It's a mindset shift to think about services BEFORE you need them.

What's a power center?

Compass Datacenters has Prefabricated Data Center Power Modules, also known as Power Centers, which are standardized, pre-engineered weatherproof containers that include switchgear, UPS, monitoring, security, and more. Power Centers are factory assembled allowing them to be simple, scalable and fast to deploy.

Compass plans to deploy 1,600 Power Centers. Saving $7K at the design phase per power center saves $11.2M[15] of installation.

Schneider Electric | Transforming Data Center Services: AI-Driven Condition-Based Maintenance

Mitigating risks in the transition to AI-driven Condition-Based Maintenance

Adopting systemic Condition-Based Maintenance (CBM) represents a fundamental shift in how data centers manage reliability, costs, and long-term operations. While the benefits are clear – reduced downtime, optimized maintenance, and enhanced sustainability – the transition requires careful risk mitigation. Executives must focus on capital cost control, TOTEX optimization, and change management to make Condition-Based Maintenance a strategic success.

One of the primary concerns for executives is the upfront investment required to deploy systemic CBM. Shifting from calendar-based maintenance to an AI-driven model means integrating **real-time monitoring sensors, advanced analytics, and digital twin technologies into the design** – all of which require limited capital investment compared to the cost of equipment. The financial case for CBM is rooted in TOTEX optimization, as it systematically reduces unplanned OPEX through predictive insights that **extend asset life and reduce costly emergency interventions**.

Systemic CBM requires a cultural shift from reactive maintenance to proactive optimization. **Services teams, technicians, and operations personnel may initially resist AI-driven decision-making, fearing job displacement or over-reliance on automation.**

As a risk mitigation strategy, leadership must invest in **structured onboarding programs** that show technicians how Condition-Based Maintenance enhances their expertise rather than replaces it. AI-driven maintenance should be positioned as an augmentation tool, improving safety, efficiency, and decision-making by reducing time-consuming manual inspections and guesswork. Engaging key stakeholders early – from service teams to procurement and finance – ensures that **Condition-Based Maintenance adoption is seen as a strategic enabler, not just a cost-cutting measure.**

Systemic Condition-Based Maintenance concentrates effort to targeted interventions from **scheduled replacements to condition-based interventions**, optimizing the **useful life of mission-critical assets**. Data-driven insights allow organizations to prioritize asset upgrades **based on performance degradation rather than lifecycle assumptions**. It targets operational investments and ongoing costs that can inflate budgets. However, each asset must be considered based on the data and the actual performance. **Executive teams should choose partners who bring enough experience across data center asset portfolios to increase overall systemic intelligence.**

While IoT sensoring of equipment should be in place when a new data center is established (greenfield), Systemic Condition-Based Maintenance implementation should follow a structured, scalable deployment model of testing and roll-on into the software and systems.

This is no different than any other safe software stand-up and risk mitigation. Systemic Condition-Based Maintenance doesn't remove these established best practices. Unit and systems testing must still be completed as assets are added.

SLAs and service imperatives also shift in AI-driven Condition-Based Maintenance. Calendar-based maintenance has an advantage in that it can be scheduled far in advance whereas systemic CBM may push servicing early or later in the cycle. This means gaining approvals for technician entry, and prioritizing service must be part of the de facto operating approach. Doing so maintains uptime and allows effective support of the customer and their equipment.

Additionally, using digital twins, systems and performance data, AI-driven insights, and exogenous data (consider sources such as data from manufacturing insights to weather) will enable predictive scenario planning that further refines uptime and performance. Scheduling regular reviews and evaluations not only engages stakeholders – putting technicians alongside executives in understanding the systems, their performance and safety – it provides the opportunity for continuous improvement, improved compliance and future scalability.

Revenue loss	Lack of competitiveness	Loss of customer trust
Power outages can have a major impact on production time and lead to revenue losses. A one minute outage costs 89,000 USD based on 1,477 USD per second, according to EPRI.[16]	The main value gains from electrification will be reaped by fast movers; laggards will fall behind. Grid saturation will add time, cost, and uncertainty to electrification projects.	Lack of operation reliability and control over emissions impacts customer satisfaction and trust. In a Accenture study, 54% of the 7,000 companies studied indicated losses of customer trust equated to $180bn in missed opportunities[17].

6 key questions executives should consider when evaluating Systemic Condition-Based Maintenance

Transitioning to Systemic, AI-driven Condition-Based Maintenance (CBM) is a strategic decision that impacts cost, resilience, and efficiency.

To secure board buy-in, an executive must frame Condition-Based Maintenance as a business-critical transformation that optimizes operations, reduces downtime, and drives financial stability. Here are questions that can help you make the case to your board for funding:

A strong board-level argument should:

- Showcase financial efficiency: **How CBM** stabilizes OPEX and reduces emergency maintenance costs.
- Prove reliability gains: **Use real-world benchmarks to highlight** reduced downtime and improved workforce productivity.
- Align with cost control priorities: **Frame Condition-Based Maintenance as a** TOTEX strategy, **balancing** CAPEX investment with sustained OPEX reduction.
- Provide a clear roadmap: **Outline a** phased deployment to de-risk investment while proving efficiency gains. **Condition-Based Maintenance** isn't a nice-to-have – it's essential infrastructure intelligence that **ensures** efficient, reliable, and cost-controlled **data center operations.**

01

How does AI-driven Condition-Based Maintenance improve operational efficiency and reduce downtime risk?
A one minute outage can cost 89,000 USD based on 1,477 USD per second, according to EPRI,[18] and the complexity of modern data centers makes reactive maintenance increasingly untenable. Traditional calendar-based maintenance leads to inefficiencies – servicing assets too soon or too late. Systemic Condition-Based Maintenance provides precise, data-driven intervention timing, reducing failures by up to 40%.

For the business case:

Highlight efficiency gains: Fewer unnecessary interventions, reduced human error, and lower downtime exposure.

Capture current capabilities and gaps in existing monitoring systems and determine where these might lead to unnecessary risks.

As a means to stabilize maintenance costs over time rather than as an isolated capital expense.

02

How does AI-driven Condition-Based Maintenance optimize resource allocation and reduce total operational costs (TOTEX)?
Focus on cost efficiency. Systemic Condition-Based Maintenance stabilizes OPEX by shifting from high-variability emergency repairs to planned, data-driven interventions. Organizations that embed Systemic Condition-Based Maintenance at the design phase experience an accelerated path to reduction in operational costs while avoiding costly retrofits.

To make the point resonate:

Quantify cost avoidance: AI-driven Condition-Based Maintenance prevents expensive emergency interventions and optimizes labor allocation, reducing total maintenance spending.

Evaluate where potential technology partners and solutions interface and how these might encounter operational conflicts that would not be spotted in isolation

Show how AI extends asset lifespans, delaying expensive replacements and spreading CAPEX more predictably.

03

Can AI-driven Condition-Based Maintenance make our workforce more effective without increasing headcount?
With ongoing technician shortages, it's unsustainable to scale operations through hiring alone. Systemic Condition-Based Maintenance allows a leaner team to manage a larger footprint by enabling remote triage and AI-assisted diagnostics.

To make a Board-level case:

AI-driven Condition-Based Maintenance reduces on-site interventions. allowing existing staff to manage more assets with less physical service demand.

AI-enhanced agents and workflows streamline technician productivity, ensuring field teams only deploy when necessary – eliminating wasted truck rolls and redundant service calls.

04

Can AI-driven Condition-Based Maintenance improve power and cooling efficiency for better performance?
AI-driven Condition-Based Maintenance doesn't just reduce energy use for Environmental Social Governance (ESG) reasons – it optimizes power and cooling performance to improve uptime, reduce equipment strain, and lower unnecessary capacity overhead.

Consider these points:

Better cooling efficiency prevents overuse of HVAC units, reducing premature wear and extending asset life.

AI-driven Condition-Based Maintenance prevents power distribution failures before they happen, avoiding cascading disruptions that compromise efficiency.

Less emergency intervention means less operational disruption, keeping systems running at peak efficiency.

05

Can AI-driven Condition-Based Maintenance integrate across multiple vendors and existing infrastructure?
Most data centers are a mix of legacy and next-gen systems, and the system's greatest success is in unifying maintenance across all assets – not just one vendor's ecosystem.

Be prepared to:

Highlight interoperability: the best Condition-Based Maintenance solutions aggregate data across power and cooling infrastructures – eliminating silos and creating operational efficiencies. They are at the heart, systemic.

Show how AI-enabled Condition-Based Maintenance makes vendor-specific support more agile, and delivers a better maintenance strategy across the whole data center.

06

How do we measure Condition-Based Maintenance success and scale it efficiently?
The board needs clear KPIs that show AI-driven Condition-Based Maintenance direct impact on efficiency, uptime, and cost savings before scaling further.

You'll want to:

Define KPIs, including failure reduction by rate, maintenance cost savings, efficiency gains, and uptime improvements.

Propose a phased rollout: start with a high-value subset of assets, prove efficiency gains, and expand based on ROI results.

Use pilot data or industry benchmarks to show that AI-driven Condition-Based Maintenance offers a compressed time to value.

Appendix

[1] Uptime Institute GlobalData Center Survey 2024, Uptime Institute

[2] Schneider Electric analysis with Compass Datacenters

[3] IBID

[4] IBID, page 3

[5] Note: at Schneider Electric the NOC for these datacenters is known as the Connected Services Hub (CSH). These can be used interchangeably.

[6] Schneider Electric analysis with Compass Datacenters

[7] Annual outages analysis 2023, Uptime Institute

[8] Uptime Institute Global Data Center Survey Results 2024

[9] IBID, page 8

[10] Schneider Electric analysis with Compass Datacenters

[11] Temperatures Rising: NASA Confirms 2024 Warmest Year on Record, NASA

[12] Schneider Electric analysis with Compass Datacenters

[13] IBID

[14] Schneider Electric analysis with Compass Datacenters

[15] IBID

[16] Ubicquia - Power Quality is Vital in the Changing Commercial and Industrial Landscape

[17] Accenture the bottom line on trust

[18] Ubicquia - Power Quality is Vital in the Changing Commercial and Industrial Landscape

Special thanks to Andrew Nix adding his cybersecurity expertise. To Natasha Nelson for sponsoring the development of this work. To Kevin Castel, Nathalie Delforge and the tireless work of Anna Russo and Emperor for bringing this story to life.

About the authors

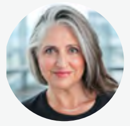

Cristene Gonzalez-Wertz
Director and Managing Editor of the Schneider Electric Data Center Research and Strategy Group (DCRS)

Cristene Gonzalez-Wertz is the Director and Managing Editor of the Schneider Electric Data Center Research and Strategy Group (DCRS). A veteran strategist and storyteller, she bridges the insight-to- action gap between complex technologies and strategic outcomes. She combines deep delivery expertise in artificial intelligence, business model redesign, and modernization, with the practical application of technology to drive measurable business value.

Her reports combine primary research with more than 25 years of global blue-chip client work. Readers benefit from distilled technical complexity and actionable, targeted strategies.

Before joining Schneider Electric, she was the Editor in Chief of the Leading Edge at DXC Technologies and a long-time veteran of the preeminent IBM Institute for Business Value.

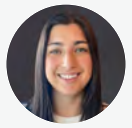

Sophie Brause
Offer Designer for Predictive Analytics

Sophie is currently serving as the Offer Designer of System Analytics for Schneider Electric's Services. She leads a team of AI data scientists, subject matter experts, and application engineers, driving the development of advanced analytics for Data Center systems to enable the transition from traditional calendar-based maintenance services to analytics driven Condition-Based Maintenance services. With extensive experience in Energy Management Software and Artificial Intelligence, Sophie is recognized for her contributions to digital innovation and transformation, consistently delivering forward-thinking solutions that optimize operational efficiency.

She holds a Bachelor of Science degree in Computer Science and Mathematics from James Madison University. Sophie is also a leader within Schneider's Leadership Development Program, serves as a mentor to new recruits into program, and is the Financial Officer for the Emerging Professional Network for Schneider Electric's Boston City Hub.

Natasha Nelson
CTO Services & Vice President EcoStruxure Power for End User

Natasha Nelson is a Vice President at Schneider Electric, and leads the EcoStruxure Power for End User, driving value and growth of Solutions business across all three levels of EcoStruxure for the Power & Grid, Mission Critical and Electro-Intensive market segments, with IoT connected products, edge and cloud solutions, using AI and Digital Twins. Natasha's focus on value creation, driving best possible outcomes for customers with digital services combining IoT, AI, analytics, expertise, cybersecurity, and field services.

Prior to joining Schneider Electric, Natasha was Chief Information Officer at Recipe, managing Information Systems, Analytics, e-Commerce/Mobile platforms, Data Center, Call Center Operations and IT Assets across 700 locations across five restaurant brands in Canada. Natasha has a Master of Business Administration from the Massachusetts Institute of Technology and a MIT Sloan Fellows Program, specialized focus on Digital Platform Business, Innovation and Cybersecurity management.

Wendi Runyon
Vice President, Global Services Incubation

With over 30 years of experience in the industry, Wendi currently serves as Vice President of Global Incubation for Schneider Electric's Global Services. Her team is responsible for understanding the market, developing innovation strategies, driving the innovation process from ideation through incubation, and pursuing Strategic Alliances for co-innovation.

Prior, Wendi spent many years inside Schneider Electric, and former acquisition APC, in Strategy, Innovation, Business Development, Product Management, and Strategic Alliances. Additionally, Wendi was Co-Founder and VP of Corporate strategy for Netcertainty as well as Intelligenxia.

Wendi holds a Masters in Business Administration (MBA) degree from Emory University and a Bachelor of Science (BS) in Industrial Engineering from Pennsylvania State University. She currently sits on Penn State's Industrial & Manufacturing Engineering IPAC (Industrial & Professional Advisory Council) Board, Penn State's Master in Service Systems Engineering Board, and SMU's Lyle School of Engineering Masters in Datacenter Systems Engineering (MDSE) Advisory Board.

Life Is On | Schneider Electric

Designed and produced by emperor
Visit us at emperor.works

se.com