

SOLUTION BRIEF



RED DOT ANALYTICS

Red Dot Analytics Pte. Ltd. (RDA) in collaboration with NTU Singapore, develops DCWiz, an industrial AI platform that digitalizes, optimizes and automates datacenter management and operations for business continuity and energy efficiency

This industrial AI platform system represents a "zero to one"1 breakthrough in the application of AI for more stable and efficient data centre management

CHALLENGE

Electricity usage of data centres has sky-rocketed in recent years, fuelled by the increasing demand for this mission-critical ICT infrastructure. In particular, cooling has been identified as one of the major contributors of energy consumption in data centres, taking up about 38% of total energy usage. Concerns have been raised on how to sustain this rapid growth without compromising the carbon footprint. Moreover, the complexity of data centre system management has increased exponentially, while the economic losses brought about by unplanned data centre outage (UDO) have grown in severity over the years. Enterprise data centres in particular, face the two-fold challenge of trying to achieve increased productivity while keeping operational costs down.

Against this backdrop, recent progress in AI technologies has provided an unprecedented opportunity for data centres to enhance its energy

efficiency and optimize its system management. Whilst the concept of using AI to improve the control policy of IT systems and facilities, as well as in automating system management is not new, data centres face two key challenges in implementing AI-based approaches:

- **Data Scarcity**

In order to develop an accurate model, all AI algorithms need to be trained by substantial data. In the context of physical infrastructures like data centres however, the collection from and access to large numbers of IoT sensors remains challenging due to high cost (e.g. costly sensors and meters), the long time period involved (e.g. peak/off-peak and seasonal attributes), as well as concerns over compliance and security (e.g. emergency faults and anomalies).

- **Risk averse Mindset**

Due to historical reasons (e.g. limited data and safety

INDUSTRY

Data Centres

SOLUTION

DCWiz couples an industry-grade digital twin with AI to create an intuitive yet high accuracy platform for data centre operators to digitalize, optimize and automate data centre operations without compromising on the safety standards required in this mission-critical infrastructure.

BENEFITS

- Change validation
- PUE optimization
- OPEX cost reduction
- Uptime improvement
- Human error reduction
- User-friendly interface

PARTNER

National Supercomputing Centre (NSCC) Singapore



compliance), the management of data centres today still depend to a large extent on the decisions of human experts. This has rendered the data centre industry with a risk-averse mindset that hinders the adoption of AI technologies. Such risk aversion has also mandated data centres to maintain high standards of reliability and safety with sufficient levels of redundancy. Consequently, data centres today are often encumbered by poor energy efficiency and high maintenance costs.

Novel approaches are thus imperative for AI-based solutions to be successfully deployed into data centre management and optimization.

SOLUTION

To tackle these challenges, a team at NTU and RDA, led by Professor Wen Yonggang from the School of Computer Science and Engineering has developed an integrated industrial AI solution combining an industry-grade digital twin with Deep Reinforcement Learning (DRL) to optimize and automate data centre management and operations.

- **Industry-grade Digital Twin**

The industry-grade digital twin is able to achieve data quality within $\pm 0.5^{\circ}\text{C}$, providing an accurate yet intuitive 3D simulation platform that covers the layout, thermal dynamics, and airflow directions of the data centre. This allows human experts to better grasp the state information of the system, quickly pinpoint hot spots and anomalies as well as ascertain how any change introduced can affect the energy efficiency as well as safety of the data centre.

More than that, the high-fidelity digital twin can generate massive amounts of **synthesized data** with high quality and diversity (including “black swan” events such as emergencies and anomalies, peak/off-peak loads, and seasonal dynamics) within a relatively short time period. These can be used to augment existing available datasets for AI algorithms to solve the aforementioned data bottleneck issue.

- **Deep Reinforcement Learning**

DCWiz employs a deep reinforcement learning (DRL) approach whereby system behaviors are learnt by iteratively interacting with the “environment” (data center) to derive an optimized control policy. Before DRL (or any AI-based approach) can converge to a final optimized solution, it is often subjected to numerous uncertainties.

Therefore, instead of solely interacting directly with the physical data centres, the DRL agent interfaces with the well-calibrated digital twin and is trained with massive data synthesized from the digital twin together with real data from the physical system. Such an approach can however involve overwhelming manual tuning which can be prohibitively costly. DCWiz thus adopts a “**Train the Trainer**” framework which involves two layers. First, domain knowledge from human experts is mapped into AI to calibrate the simulation engine; the resulting “AI trainer” then serves as an intelligent agent that is capable of “supervising” and deriving the optimal control policy. With such a framework, DCWiz exhibits competitive performance in its auto-tuning capability which significantly saves sampling costs for data centre operators.

BENEFITS

DCWiz’s integrated sandbox approach offers data centre operators a new paradigm for infrastructure management. Compared with the traditional system management in data centres, which relies solely on human expert knowledge and limited sensor readings from the data centre infrastructure management (DCIM) system, DCWiz offers the following benefits:

- **High precision with easy-to-understand user interface**

With its auto-calibration algorithm, calibration of the DCWiz digital twin requires only a small number of physical sensor readings, yet DCWiz is able to achieve a greater degree of accuracy (at least $\pm 0.5^{\circ}\text{C}$) compared to other commercial simulation software that are used for the same purpose. The fine-grained state predictions offered by DCWiz is coupled to an intuitive 3D graphical visualization interface, enabling one to for example, determine the temperature at any point within the data centre room space. This feature enables operators to accurately grasp the current state of the data centre using a much smaller number of physical sensors, without violating safety compliance.

- **High safety and efficient what-if analysis**

DCWiz allows data centre operators to conduct what-if analysis safely and efficiently to validate the effects of changes to the data centre (e.g. tolerance tests for peak hours). For instance, in order to estimate the effects of a proposed change to a data centre, operators can simply input an instruction to DCWiz and execute the simulation to visualize the outcome. Unlike testing on a physical data centre, such testing on the cyber digital twin requires no prior preparation (e.g. system maintenance), needs no consideration on time factors (e.g. peak versus off-peak hours), and does not pose any threat to the physical system.

- **Automated cyber-physical control loop**

DCWiz provides an automated cyber-physical control loop to reduce the extremely labour-intensive and error-prone manual operations in data centre management. The solution starts at the physical loop, where data are collected to construct and calibrate the digital twin. Subsequently, in the cyber loop, AI algorithms are iteratively trained and tested via the digital twin to obtain an intelligent model that is able to recommend optimized control actions according to real-time system readings. Subsequently, the recommended control actions can be validated via the industry-grade digital twin to prevent any risk during deployment. Only when these control measures are proven secure will they be deployed to the actual data centre.

1. Design Validation

A cognitive digital twin of the NSCC Data Centre Aspire 1 was built using the DCWiz platform to validate the data centre design as well as understand the gap between the expected performance (based on the design) and the actual performance observed during operation. Through DCWiz, inefficiencies in the design that affect the energy efficiency of the data centres were discovered and solutions to rectify these design flaws proposed. In this case, it was found that by some simple reconstructions of the data centre hall to restrain hot/cold air mixing and minimise the loss of cold air, an improvement of PUE from 1.35 to 1.3 could be achieved (for 40 racks), with an accompanying energy cost savings of S\$6,000 per month.

2. Optimization of Cooling Operations

Due to the intricacies of the interactions between the different components involved in data centre cooling, developing an optimal control policy for cooling is a challenging task faced by many data centre operators. Through “trace-based” AI optimisation², controlled optimization of air cooling and warm water cooling systems were achieved. With the help of DCWiz, NSCC achieved energy savings of 15% for an air-cooled system and 30% for a water-cooled system.

Widely recognized in industry and academia, DCWiz has won the 2020 IEEE TCCPS Industrial Technical Excellence Award, 2016 ASEAN ICT Award (Gold Medal) and 2015 DCD Award - APAC.

RESULTS

DCWiz recently has demonstrated a Proof of Concept (PoC)/Proof of Value (PoV) at the National Supercomputing Centre (NSCC) of Singapore. For enterprise data centres such as NSCC, improving productivity whilst keeping costs low are the main priorities. DCWiz enabled these two objectives to be met via:

