kyndryl

# Data center energy optimization with Kyndryl Data Center Advisor

# Contents



In today's era of cloud, IOT, digitization, and social media, the importance of data centers has drastically increased. Once considered cost centers, data centers have transformed into nerve centers for today's enterprises—making managing data center operations a key function for enterprises worldwide.

While data centers have become increasingly sophisticated, managing operations has become exponentially more complex. Data center outages still happen despite the redundancies built into their design, and energy continues to be the biggest cost driver for data centers despite innovations in equipment efficiency. The inability to optimize equipment performance and manage capacities due to lack of skilled resources often leads to bottlenecks in meeting availability, efficiency, and compliance goals.

It's clear that conventional data center management tools are no longer enough. To be operationally and economically viable, data centers need to implement artificial intelligence (AI) and machine learning (ML) solutions—which is where Kyndryl™ Data Center Advisor can be a game changer.

## What is the Data Center Advisor?

Kyndryl Data Center Advisor is a fully managed advanced analytics service that uses AI and ML capabilities to optimize data center infrastructure operations, providing data-driven and AI-enabled insights to improve efficiency, mitigate operational risks, and optimize operational costs. Additionally, the service empowers data center operators with actionable insights to help reduce turnaround time.

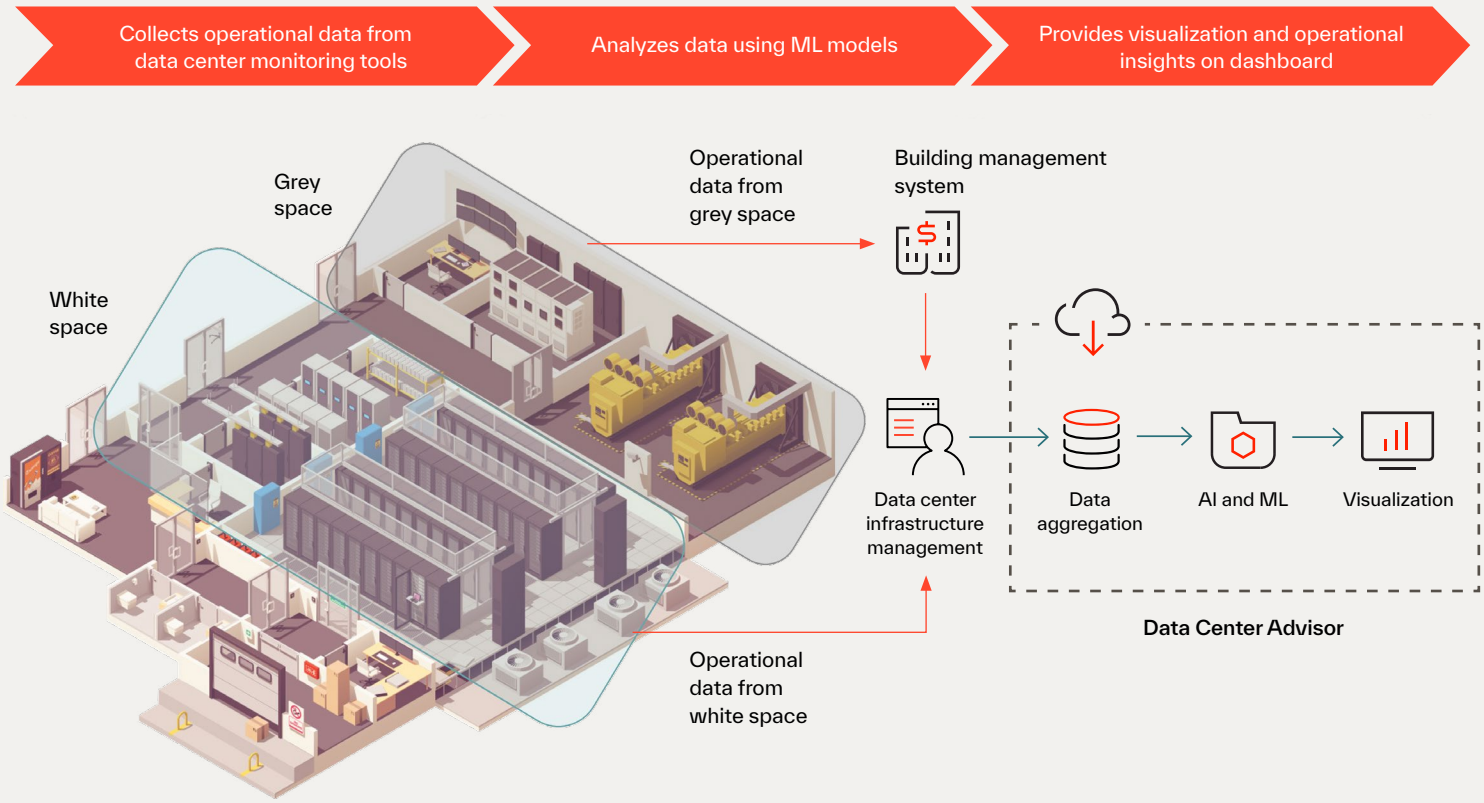


Figure 1. *Kyndryl Data Center Advisor in action*

## How does it work?

Data Center Advisor works by aggregating operational data from both data center facilities and IT endpoints. It analyzes the data using ML models and provides predictive operational insights.

## Advanced analytics capabilities

Data Center Advisor provides a full spectrum of advanced analytics capabilities for data center operations, designed to empower the user with data-driven insights.

Descriptive and diagnostic analytics provide data exploration, anomaly detection, business intelligence, and visualization to answer questions about what happened and what is currently happening in the data center. Advanced predictive analytics techniques use ML algorithms to uncover near real-time insights and forecast and predict future events. Finally, the AI models for prescriptive analytics provide simulations and recommendations for optimal actions to achieve business objectives.

## Advanced analytics in action

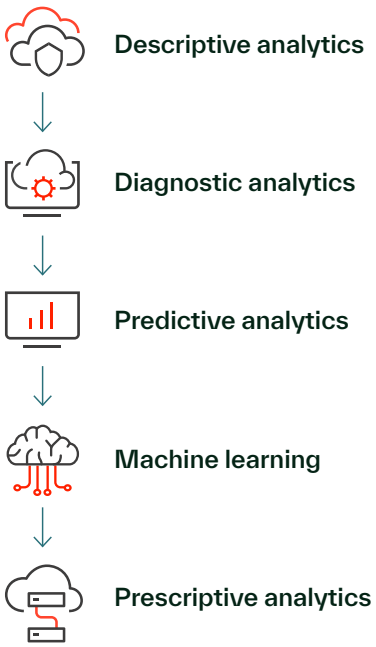

## Use cases and ML models

### Business Objectives

Optimize data center energy efficiency

### Use Cases

- Optimize data center power usage effectiveness (PUE)
- Optimize chiller plant efficiency and capacity deliverance

### AI Models

- Predict PUE and recommend target set points of mechanical and electrical components for PUE optimization
- Provide prescriptive insight on chiller plant operations, including sequence of operation, recommendations on mode of operations, and recommendations on optimum set points

---

### Business Objectives

Improve equipment performance and reliability

### Use Cases

- Predict equipment outages and degraded performance

### AI Models

- Enable performance and health anomaly detection from equipment metric data and system logs

---

### Business Objectives

Advise on appropriate equipment maintenance

### Use Cases

- Prescribe equipment maintenance based on actual operating condition

### AI Models

- Identify the root cause of a potential outage and recommend appropriate maintenance of equipment and components

## Energy optimization—how it is done?

In a typical data center, IT devices like servers, storage, and network account for the largest share of direct electricity use, with cooling systems next in line for total energy consumption. Data Center Advisor is designed to identify areas of optimization within these cooling systems and provide prescriptive, analytics-based insights and recommendations on how to minimize energy footprint and improve efficiency.

As chillers are a significant percentage of HVAC energy consumption, they need to be operated with optimum parameters to maximize their efficiency across data centers for an overall energy-efficient operation. Other subsystems like pumps, cooling towers, and air handlers also have considerable potential for optimization.

Data Center Advisor has AI models that analyze the operating behavior of these different devices and help maximize their efficiency under varied operating conditions. These AI models help to achieve the desired level of data center energy efficiency and PUE by optimizing and improving key performance indicators (KPIs) of the cooling system components.

The predictive and prescriptive AI models in Data Center Advisor are trained on huge amounts of data gathered from several disparate data centers. These models can identify potential areas of optimization within the complex landscape of a data center to generate actionable outcomes, which can drive significant improvement in energy efficiency.

Deep learning neural network models can dynamically manage HVAC equipment set points and optimize date center energy consumption and reduce cost by more than

**25 percent**

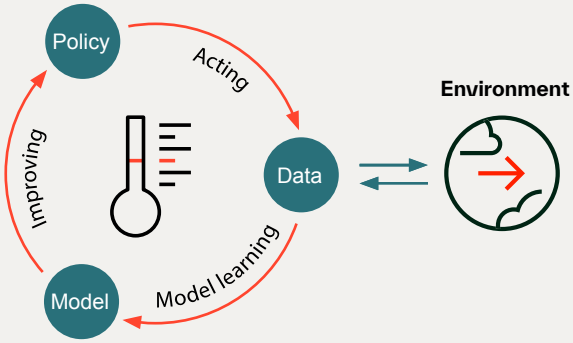


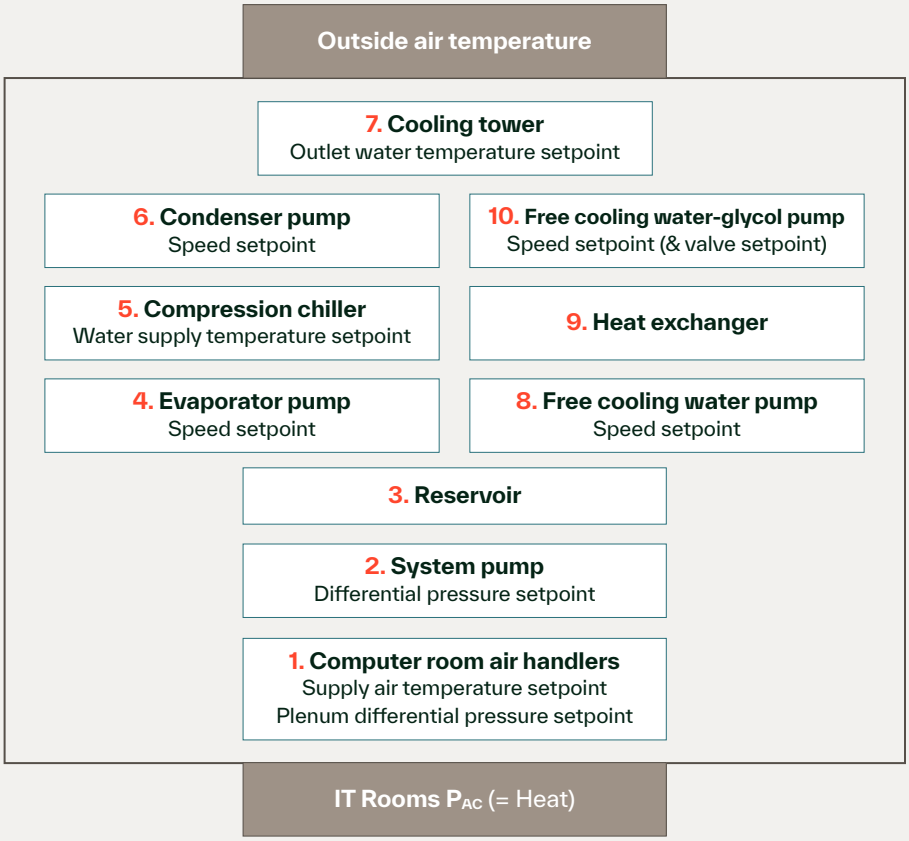


Figure 2:
*Optimizing cooling energy through setpoint modulation in a client environment*

## Target set point for operational parameters

Data Center Advisor optimization models take the following approaches to minimize cooling energy use and improve data center PUE:

- Dynamically prescribe set points based on IT load, cooling demand, ambient temperatures, and other changing conditions.
- Simulate the power consumption at various set points.
- Simulate scenarios for different device types, such as chillers with and without variable speed drive.

For a frequently varying data center operating parameters, the prescriptive AI models can also be trained to perform cluster analysis of the changing operating patterns with any new dependent covariates and simulate set points against the significantly varying attributes for the identified cluster.

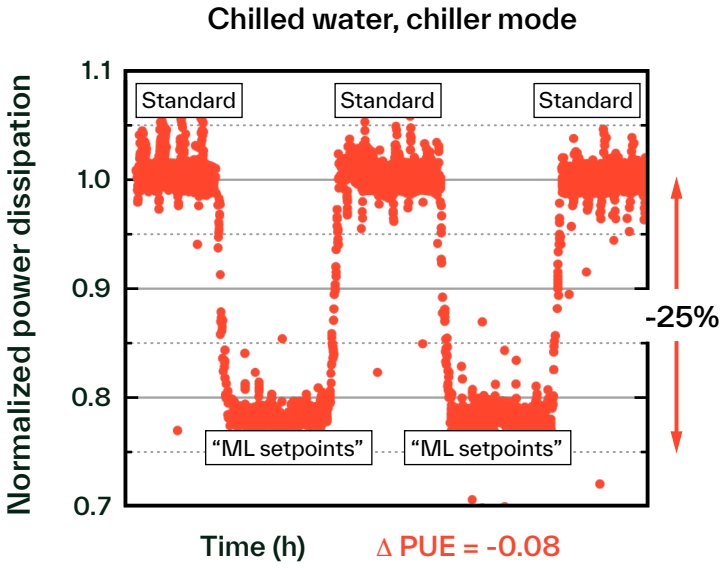

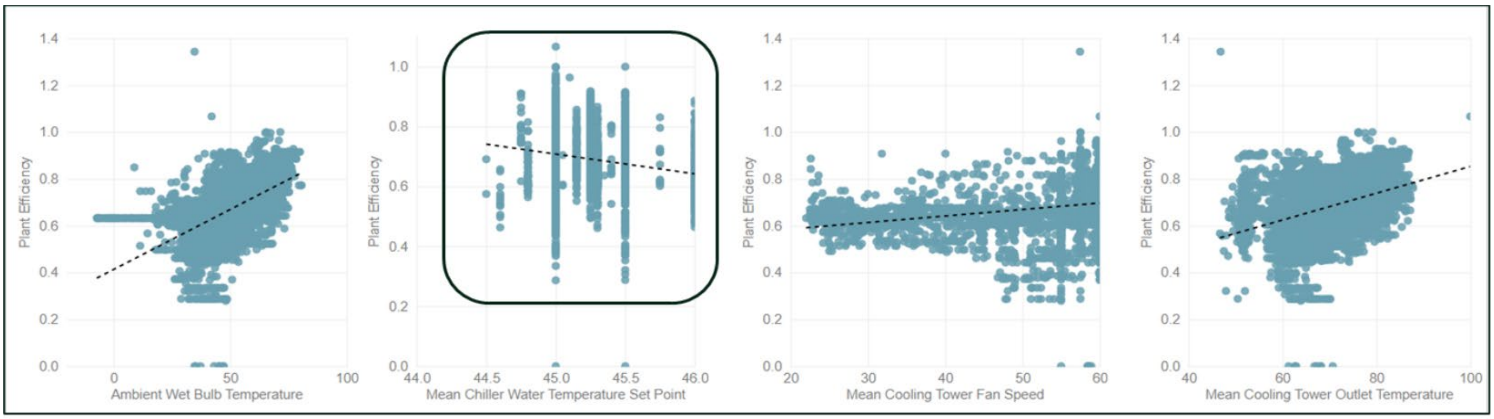


Figure 3. Client data shows how increasing chiller water temperature setpoint can improve efficiency

## The real-world power of Data Center Advisor – customer example

Using real client data, we performed sensitivity analysis of key influencers on plant efficiency, which indicated that increasing chiller water temperature setpoint can improve efficiency.

Based on this finding, we simulated two scenarios for changing the chiller water temperature setpoint of a water-cooled chiller to ML-recommended values.

The simulation output showed that increasing the water temperature setpoint by a factor of 1 above the standard setpoint would improve plant efficiency 14%, while increasing by a factor of 3 would improves the plant efficiency by 4%.

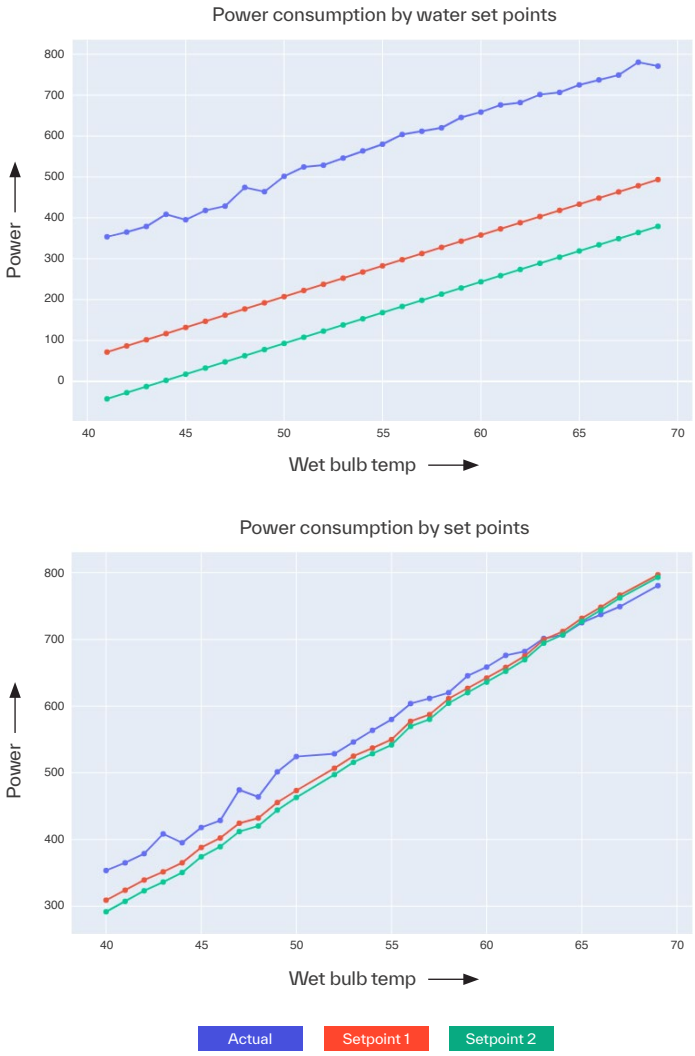


Figure 4.
Two scenarios for changing the chiller water temperature setpoint of a water-cooled chiller to ML-recommended values

## Prescriptive insights on optimal device combination

Different cooling devices—chillers, cooling towers, pumps—and the combination of such devices run in tandem to meet the cooling demand of a data center under various conditions. Each of these devices has a different level of power consumption based on factors like device health status, manufacturer, type of speed drives, and more. AI models use these differences to propose recommendations about which set of devices within a device category should be operated together to maximize efficiency.

For example, data center chillers may vary in design attributes like the type of chiller (cooled by air or water), capacity (tonnage), speed drives (constant speed drive or variable speed drive), device health score (how old the device is, run times), manufacturer, and more. These variations map to differences in power consumption when operated individually or in combination.

Data Center Advisor AI models take into consideration these parameters and respective operating behavior to come up with recommendations about which combination of devices should be run under what conditions to maximize efficiency.

## The real-world power of Data Center Advisor – another customer example

We used diagnostic analytics to run pattern recognition on a client's data set, identifying an optimal combination of chillers to run in pay cool mode. Recommendation by Data Center Advisor AI model on optimum chiller combination showed potential improvement in plant efficiency by 10%.

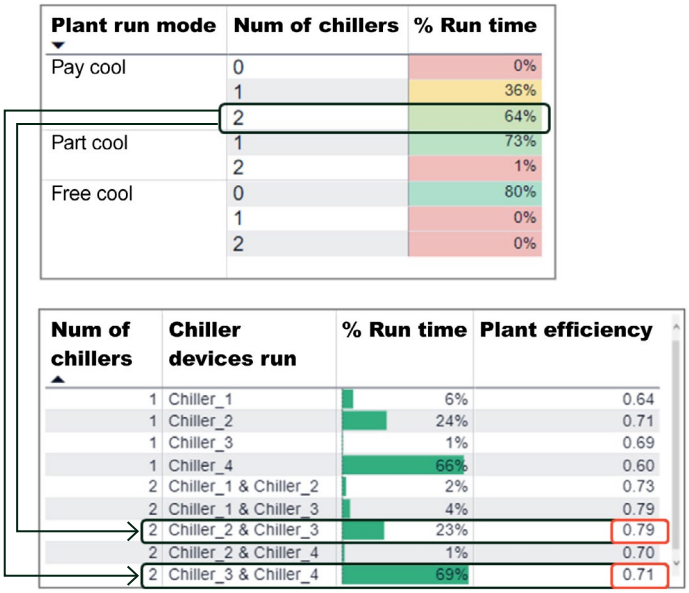

| Plant run mode | Num of chillers | % Run time |
|---|---|---|
| Pay cool | 0 | 0% |
| | 1 | 36% |
| | 2 | 64% |
| Part cool | 1 | 73% |
| | 2 | 1% |
| Free cool | 0 | 80% |
| | 1 | 0% |
| | 2 | 0% |

| Num of chillers | Chiller devices run | % Run time | Plant efficiency |
|---|---|---|---|
| 1 | Chiller_1 | 6% | 0.64 |
| 1 | Chiller_2 | 24% | 0.71 |
| 1 | Chiller_3 | 1% | 0.69 |
| 1 | Chiller_4 | 66% | 0.60 |
| 2 | Chiller_1 & Chiller_2 | 2% | 0.73 |
| 2 | Chiller_1 & Chiller_3 | 4% | 0.79 |
| 2 | Chiller_2 & Chiller_3 | 23% | 0.79 |
| 2 | Chiller_2 & Chiller_4 | 1% | 0.70 |
| 2 | Chiller_3 & Chiller_4 | 69% | 0.71 |

Figure 5:
*Data Center Advisor recommendation on optimal chiller combination shows potential for improvement in plant efficiency*


Figure 6: *Simulation of new setpoint values not based on historical data*

## Prescriptive insights for appropriate times to switch between cooling modes

A data center operates different cooling modes throughout the year: pay cool, when maximum numbers of chillers are run; free cool, when chillers are cut off to utilize ambient cold conditions and save on electricity; and part free cool, a transitional mode between pay cool to free cool modes.

These modes are operated depending on several factors, such as ambient temperature, pressure, relative humidity, entropy, location, load, seasonal variations, and more. Data Center Advisor AI models take into consideration all these variables and recommend when the operating mode should be switched from one mode to another.

Each of these modes have a set of operating parameters that are dynamically recommended by the AI models. Switching from one mode to another means different sets of devices and combinations may have to be operated.

Data Center Advisor AI models also learn from historical patterns and generate multi-variate forecasts, which then feed into our custom AI algorithm to predict when to switch the operating mode, enabling data center operators to plan and apply these actions in advance.

## Predict and simulate cooling system efficiency based on recommended values

Data Center Advisor AI models learn from the complex relationships between operating attributes and create generalized models used to produce results for simulations. This approach uses multiple models to predict a simulation scenario for a given set of operating parameters.

The AI models also help circumvent the problem of missing historical data with custom AI algorithms, enabling prediction that can help data center operators run new scenarios that have not been previously encountered. For example, if a certain setpoint value was never run historically, the operator can still predict efficiency by simulating it.

## Customer story

A Kyndryl client was looking for an AI-based data center infrastructure management solution to get insights on data center energy use and further optimize the energy footprint of the data center.

Data Center Advisor energy optimizations AI models were deployed for the customer's chiller plant system. 36 months of historical operational data was collected from chillers, pumps, cooling towers, and other HVAC components. Similarly, 36 months of electricity consumption data was collected from IT and cooling systems. The data was analyzed using the Data Center Advisor models, the output of which was made available to the client through the Data Center Advisor dashboard.

With Data Center Advisor energy optimization AI models, the customer was able to realize 7.5% average improvement in plant efficiency on pay cool mode and 13% average improvement in plant efficiency on free cool model—which translates to 108,000 kWh of energy saving per month or 1.3 million kWh of energy savings annually.

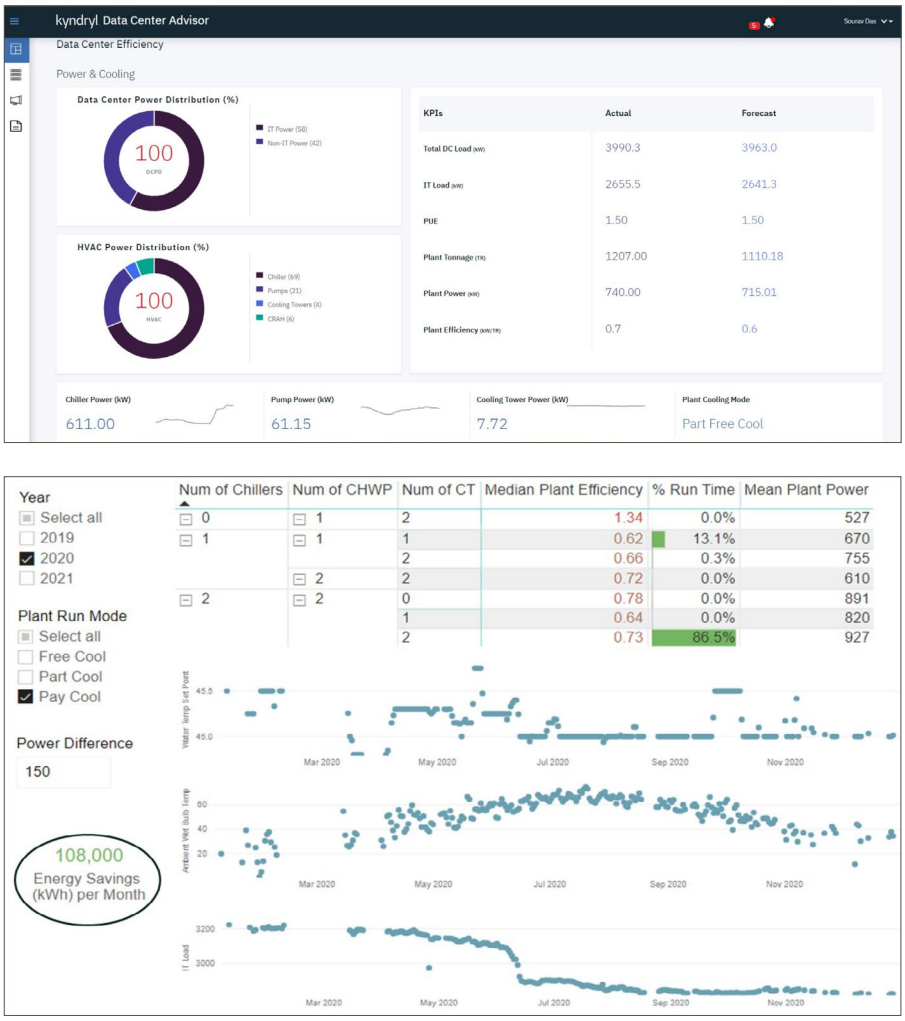





Figure 7: *Kyndryl Data Center Advisor dashboards.*

## Conclusion

Kyndryl Data Center Advisor provides data-driven and AI-enabled insights that to help you make proactive choices for improving data center infrastructure reliability and efficiency while driving down operational costs and empowering data center managers with the operational insights necessary to reduce turnaround time.

Learn more about how Kyndryl Data Center Services can help you transform your data center to respond with agility to dynamic business demands. Read the paper

Chat with a Kyndryl expert for a 30-minute strategy session. Schedule a consultation

kyndryl