



## A SELF-DRIVING SUSTAINABLE CLOUD FOR GREENER BUSINESS

Businesses must make their energy-guzzling data centers more sustainable; one intelligent way is advanced AI.

Businesses [use the cloud](#) to manage scalability, agility, availability, and develop new capabilities. However, cloud data centers are huge energy consumers — they will soon account for 15% to 25% of scope 2 (indirect) carbon emissions.

As the world becomes ever more judicious toward greener energy, firms must consider their sustainability credentials while managing cloud deployments. Physical servers require huge amounts of power to cool energy-guzzling machines. Also, electric grid power is carbon intensive. No wonder CTOs and CSOs are working together to [create green or sustainable clouds](#). Infosys Cobalt, a suite of cloud services, keeps [sustainability at the core](#), which enhances ESG attractiveness for clients and investors.

Leading thinkers globally are developing dedicated research programs and technologies to address the issue. Among them are the University of Illinois Urbana-Champaign (UIUC) and Infosys, who have partnered to offer self-driving sustainable clouds and other innovations in the area.

Some firms are reducing resource-intensive legacy infrastructure, while decreasing the number of clouds and core data centers themselves. Other initiatives include moving edge locations to the cloud, embracing cloud-native operations, rationalizing resource utilization, and purchasing carbon offsets. It also means enhancing power usage effectiveness (PUE) and carbon usage efficiency (CUE).

Another way to go greener is to place data centers in cooler regions, though cooler land is scarce due to rising competition. Some also use water from rivers to cool circuitry, like those built in Denmark. Some are advocating for greener software, which, when written appropriately, will reduce the use of hardware to run a line of code.

## The self-driving cloud

This all amounts to a huge optimization problem — how should hardware, software, and carbon be dealt with to reduce the complete energy footprint of server farms globally?

Infosys and UIUC have partnered to create and implement self-driving sustainable clouds, at scale. Here the cloud learns for itself how to increase performance, reduce power, and meet extensive service level agreements (SLAs). This, in turn, requires advanced AI and machine learning (ML) algorithms.

Current experimental approaches have proven inadequate for optimization problems. Heuristics often provide quick short-term solutions, but they are not the best ones to meet multiple rigorous SLAs at once. Instead, we propose a new-generation cloud innovation with enhanced reinforcement learning (RL)-driven control technologies, which has the parallel power to transform the entire energy transition space and revolutionize future cloud programs.

So how does this intelligent “AI blanket” control layer work?

These ML models address the huge feature space of possible cloud inputs before allocating resources in the most

energy-efficient way. The RL reward function automatically senses the environmental situations of the data center and performs the best suitable management policy without specific domain knowledge.

Central to the AI formula are virtual machines (VMs) that abstract away the underlying compute and allow physical machines in the cloud to simultaneously run different applications. In this way, compute goes up, increasing performance, while RL balances the amount of power generated (or energy guzzled), enabling large cloud providers to still meet customer SLAs on availability, performance, energy usage, total cost of ownership (TCO), and other defined metrics.

This approach reduces the need to fully switch data centers to renewable energy, a topic of discussion in policy circles in the US, China, and Europe. For instance, the [Chinese government is urging CSPs in Beijing to reduce PUE below 1.5](#), advocating for greener energy and water saving in their facilities.

World data centers require \$4.5 trillion to move completely to renewable energy, exceeding US military spending since 2001

However, switching to renewable energy requires significant upfront investments. For instance, world data centers require [\\$4.5 trillion](#) to move completely to renewable energy, exceeding US military spending since 2001. While it is important to eliminate carbon footprint, it is of course more advisable to limit energy use in the first place.

## Keeping green costs low

There is another factor to this discussion.

Green data centers offer reputational advantages and support sustainability, but they are expensive. The price of running workloads on green energy is on average 10% higher than conventional fossil fuel, according to Infosys experts. Using RL and VMs, cloud computing will address the sustainability problem, while reducing the cost of green cloud services. The optimization problem addressed just has to include operational expenditure as a significant variable to be optimized, while meeting all other SLAs. One good example is Vapor IO, which builds modular data centers. It has created an AI system “[Synse](#)” that uses telemetry to dynamically place workloads for high availability, low costs, and low carbon emissions. It distributes applications across multiple clouds using a control plane overlay running on AI.

This sort of system will go a long way to optimize cloud business performance, juggling the need to increase sustainability while providing industry-specific solutions, at a faster time to market.

# The Infosys solution

Infosys is working with its partners to improve its sustainability index. Part of this work is in the use of multicloud, in which the “Right Cloud Framework” delivers accelerated cloud adoption, agility in operations, and flexibility with “as a service” for all cloud components. This framework includes the green data center, which leverages all the elements already discussed.

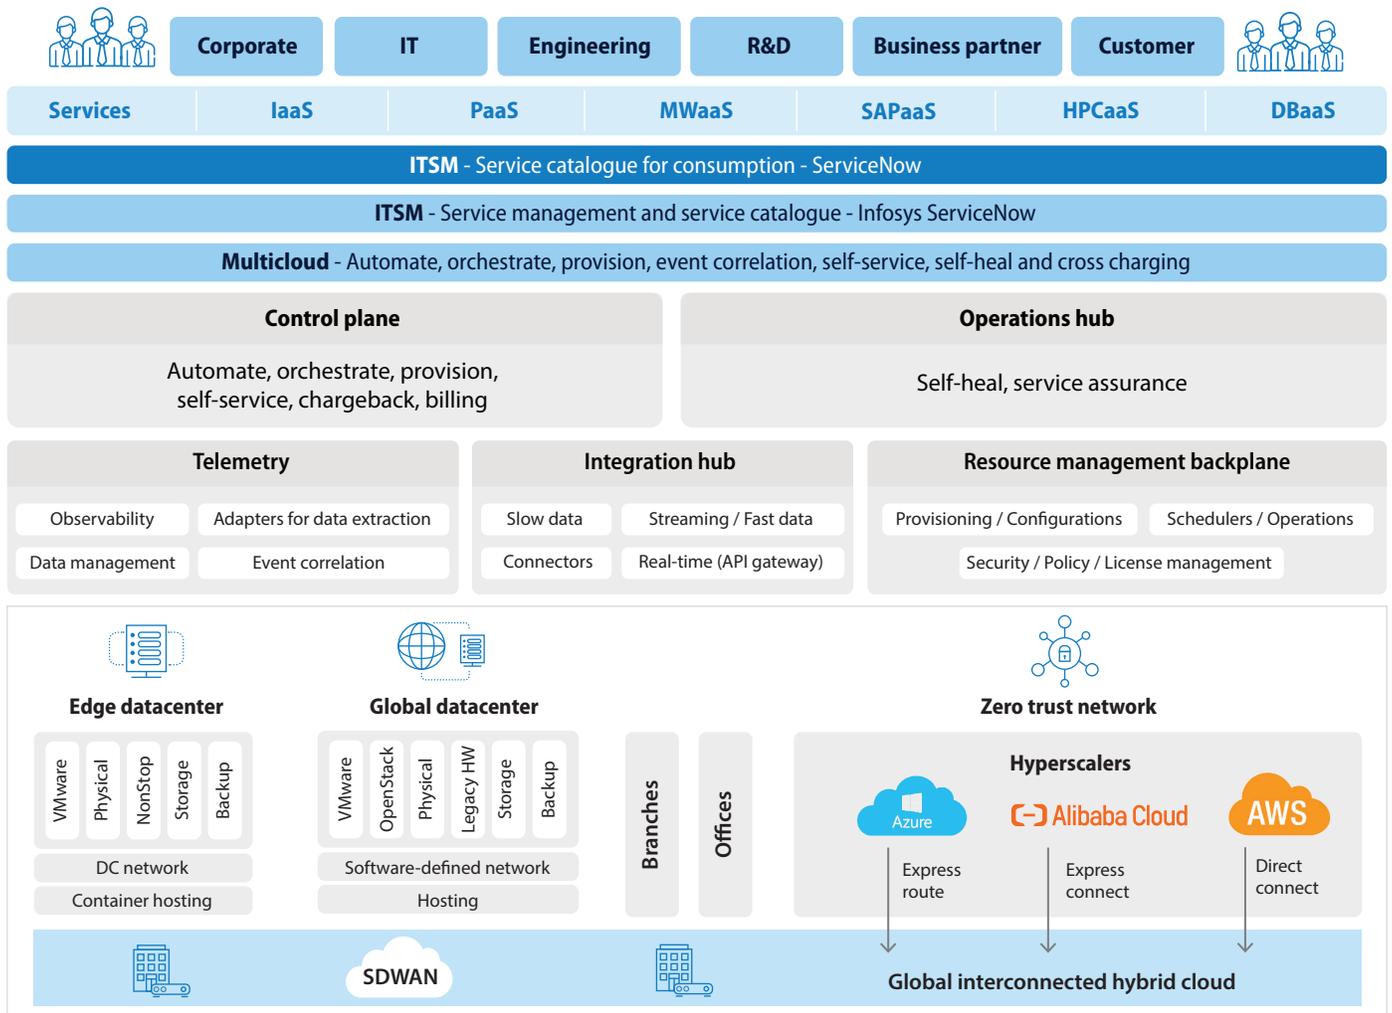
Using the control plane, telemetry, and resource management backplane (see Figure 1), goes a long way to reducing scope 2 emissions by 50% in just three years.

Working with Infosys, companies:

- Design data centers with high operational standards and energy efficiency in mind.
- Reduce over-provisioning of hardware. Our work finds that most customers over provision by 50% of peak requirements.
- Use consumption models that enable pay-as-you-go cost benefits and a massive improvement in energy efficiency.

Our work in private cloud brings many benefits. According to our Enterprise Infrastructure Transformation team, the solution results in 84% less rack space, for greener operations; 30% reduction in carbon footprint, with an eye toward reducing power and the amount of cooling needed (the goal is to utilize 80% less energy and 85% less water than traditional cooling systems); and ultimately better performance. This solution delivers an improved carbon footprint target of less than 1.4 PUE (which can be improved by a further 50%, based on our new research at UIUC), and lower product power consumption by using industry certified hardware products and data center design refinement. In our renewal centers, customers keep products and materials in use for longer. Further, by creating a plan between IT and the CSO, firms ensure data centers align to the principles of the Circular Economy through a CapEx model.

Figure 1. Infosys’ right cloud framework



Source: Infosys

## Moving forward

If global hyperscalers like Google, Microsoft, and Amazon, and system integrators like Infosys, create a hybrid cloud orchestration platform based on this “AI blanket” control layer technology, CIOs and CTOs will have a windfall. They can then together control VM operations, improve performance, and take advantage of renewable energy while meeting energy efficiency SLAs. This requires visible energy consumption and carbon-aware workloads. Placing and scheduling workloads through ML aims to minimize energy consumption and reduce emissions.

Beyond enterprise borders, the gains are likely to be huge. Telemetry and AI will enable a composable grid, wherein, for instance, a data center in Dallas, can talk in real time to a data center in Illinois or San Francisco, and trade energy supply with the lowest carbon emissions possible. This will ultimately prevent emissions for the ICT industry (alone) at a rate of [ten times its own footprint by 2030](#).

Work is already in progress. The Open Grid Alliance, founded by VMware and Vapor IO, is a multicloud framework that enables dynamically composed and optimized applications with the right type and number of resources for energy efficient operations. The aim is a globally distributed multidimensional fabric of “[heterogeneous compute, data, and intelligence to power on-demand, context-aware, immersive, and distributed modern applications anywhere](#).”

Until recently, energy efficiency was something spoken about in the context of cloud costs but didn't directly impact the sustainability criteria mandated by governments and industry bodies. But now, sustainability is a major factor in business decisions.

The future is where a self-driving, sustainable cloud solves complex optimization problems for the good of people, planet, and profit

Sustainable cloud optimization largely focuses on utilization, energy, and energy sources (which themselves will be dynamically optimized for cost and performance). Getting this right will lead to more governance on other stand-out issues such as security, quality of service, data provenance, and even data privacy. The future is where a self-driving, sustainable cloud solves complex optimization problems, benefiting CSPs and their clients by promoting people, planet, and profit.



## **Authors**

### **Rajeshwari Ganesan**

*Distinguished Technologist, Infosys*

### **Professor Ravishankar K Iyer**

*University of Illinois Urbana-Champaign*

### **Harry Keir Hughes**

*Infosys Knowledge Institute*

---

## About Infosys Knowledge Institute

The Infosys Knowledge Institute helps industry leaders develop a deeper understanding of business and technology trends through compelling thought leadership. Our researchers and subject matter experts provide a fact base that aids decision making on critical business and technology issues.

To view our research, visit Infosys Knowledge Institute at [infosys.com/IKI](https://infosys.com/IKI) or email us at [iki@infosys.com](mailto:iki@infosys.com).

---

For more information, contact [askus@infosys.com](mailto:askus@infosys.com)



© 2023 Infosys Limited, Bengaluru, India. All Rights Reserved. Infosys believes the information in this document is accurate as of its publication date; such information is subject to change without notice. Infosys acknowledges the proprietary rights of other companies to the trademarks, product names and such other intellectual property rights mentioned in this document. Except as expressly permitted, neither this documentation nor any part of it may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, printing, photocopying, recording or otherwise, without the prior permission of Infosys Limited and/ or any named intellectual property rights holders under this document.