EDB POSTGRES® AI

Optimizing the Core

Quantifying the Transformation of Customers' Data Centers with EDB Postgres® AI

enterprisedb.com

TABLE OF CONTENTS

Abstract

As global pressure mounts to achieve net-zero emissions, organizations are increasingly seeking innovative solutions to reduce energy consumption and lower greenhouse gas (GHG) emissions. The sovereign data and AI platform EDB Postgres AI (EDB PG AI), from EnterpriseDB (EDB), presents a compelling opportunity for businesses to optimize IT infrastructure while reducing costs and making meaningful progress toward their sustainability goals.

This white paper explores the methodology behind quantifying the carbon reduction enabled by EDB PG AI—specifically through minimizing the architectural footprint required to run enterprise applications in data centers. Using robust methodology and case studies of three customers, we demonstrate how EDB PG AI can reduce core usage by up to 94% in certain scenarios, and thus can significantly reduce associated emissions for applications relying on EDB technology. In addition, the paper provides a transparent overview of the approach used to estimate avoided emissions, equipping organizations with a practical model for assessing the downstream climate benefits of more efficient software deployment.

Glossary of terms

Term	Definition
Avoided emissions	The amount of carbon emissions that have been or can be avoided using EDB PG AI, compared to the same application without EDB PG AI
Carbon emissions	The release of carbon dioxide (CO_2) into the atmosphere, primarily from the burning of fossil fuels; in a broader context, the term can also include other greenhouse gases
CO_2e	Carbon dioxide equivalent (CO_2e), a standard unit for measuring the impact of different greenhouse gases, including carbon dioxide, methane, and nitrous oxide
Emission factor	A value that quantifies the amount of greenhouse gases released into the atmosphere for a given activity or process, and used to calculate a business's carbon emissions
Environmental sustainability	Using natural resources responsibly to protect the environment; this includes using less energy and relying less on fossil fuels, thus reducing the amount of greenhouse gases released into the atmosphere
Greenhouse gases (GHG)	Gases in the atmosphere that trap heat, raising the Earth's surface temperature
$MtCO_2e$	Metric tons of carbon dioxide equivalent ($MtCO_2e$), a quantification of greenhouse gas emissions in terms of millions of grams of CO_2e, commonly used in emissions reporting

1. Introduction

The digital transformation has spurred an exponential increase in data and computing demands. This places immense pressure on data centers, which use tremendous quantities of energy. Their rapid growth is contributing significantly to global carbon emissions, even as the solutions they host improve quality of life. Therefore, opportunities to make data centers more energy efficient are paramount to fighting the worst effects of climate change. As businesses strive to meet ambitious environmental, social, and governance (ESG) targets, optimizing IT infrastructure for sustainability is now a critical imperative.

EDB offers advanced data and AI management solutions that not only enhance performance, scalability, and efficiency but also play a key role in helping organizations achieve their sustainability goals. By reducing the computational resources needed to run applications, EDB PG AI can significantly reduce the resulting carbon emissions released at data centers by applications running the software. This white paper lays out a methodology for quantifying these avoided emissions and presents compelling evidence related to three anonymous major global enterprises.

To quantify the environmental benefits of EDB Postgres AI, this analysis focused on three large global customers in the banking, financial services, and insurance (BFSI) sector—an industry with demanding infrastructure and sovereignty requirements. Together, these customers operate more than 120 data centers, placing them among the few companies worldwide with such large-scale footprints. EDB collaborated with Incendium Consulting, a global advisory firm specializing in sustainability strategy, GHG accounting, and ESG reporting, to assess the emissions impact of adopting EDB PG AI. The analysis was based on a combination of customer-provided infrastructure data, validated modeling assumptions informed by EDB technical experts, and extrapolation of those assumptions to two additional EDB customers who are currently implementing EDB solutions. The resulting findings demonstrate the scale at which EDB PG AI can drive meaningful efficiency and emissions reductions, with the methodology and results independently validated by Incendium.

2. Methodology and assumptions for avoided emissions calculation

To assess the environmental impact of EDB PG AI, we developed a methodology to calculate avoided emissions, which are defined as the difference between emissions generated by an application running without EDB PG AI (the "baseline") and those generated by the same application running with EDB PG AI (the "new emissions").

2.1 Avoided emissions formula

Avoided emissions = Old emissions (without EDB PG AI) – new emissions (with EDB PG AI)

2.2 Key components and derivations

- **Emission factor:** Derived from data provided by Customer 1's Tier 1 applications using EDB PG AI and Tier 2–4 applications with and without EDB PG AI, along with the data center, server, and core architecture. Customer 1's emissions per application are linearly tied to the total number of cores required to maintain the application, with old Tier 2–4 applications requiring 144 cores that generate 10.8 $MtCO_2e$ per application annually and Tier 2–4 applications with EDB PG AI requiring only 24 cores and generating 1.8 $MtCO_2e$. This corresponds to emissions of 0.075 $MtCO_2e$ per core per year.
- **Tier definitions:**
  - **Tier 1 applications:** High availability and replication requirements, typically using more servers to ensure maximum uptime
  - **Tier 2–4 applications:** Fewer servers required due to lower availability needs
- **New architecture:** EDB PG AI introduces modern replication patterns that significantly reduce the number of servers and cores required to run applications.

2.3 Assumptions

- **Location independent:** Differences in emission intensity of the grid were not considered. However, when comparing percentage reductions in electricity-based emissions in a before/after study, this is not impactful when percentages are reported. Comparing across organizations or data centers with different emissions intensities is meaningful and was not considered.
- **Market-based emissions mechanisms ignored:** Mechanisms such as renewable energy certificates (RECs) and other market-based methods for reducing emissions were excluded from the analysis for simplicity.
- **Linear emissions scaling:** Emissions are assumed to scale linearly with server count for each customer.
- **Constant server efficiency:** Server efficiency is assumed to remain constant within different application tiers for calculation purposes.
- **Core count validation:** Core counts are used as a primary metric to standardize server scaling across customers and validate computational resource utilization.
- **Equipment supporting applications:** Secondary equipment used in the support of applications was not considered. Therefore, total realized changes in electricity demand may differ.
- **Future adoption by Customers 2 and 3:** Core count changes were estimated for Customers 2 and 3 based on assumptions by EDB technical experts and sales staff with deep knowledge of the industry. Customers 2 and 3 were assumed to fully adopt the architecture for an estimated number of applications. At the time of this modeling exercise, only Customer 1 has fully implemented EDB's solutions for the subset of applications modeled.

2.4 Data gaps and pending information

Certain configurations were estimated from feedback provided by EDB's technical personnel. Other potentially meaningful variables, such as generative AI (GenAI) and analytics, were not assessed at this time. Customer 1 publicly reports on data center electricity usage and has fully implemented the solutions for the applications modeled, but Customer 2 and Customer 3 do not; therefore the impact of EDB PG AI on GHG emissions for Customers 2 and 3 were modeled as future expectations based on full implementation of EDB's solutions on a subset of their applications.

2.5 Equation to derive customers' avoided emissions

Using the equation cited in 2.1, we derived the following avoided emissions for each customer:

$$\text{Avoided emissions} = 3.6\ MtCO_2e * \text{sum}\left(_{\text{Tier 1 apps}}\right) * \text{sum}\left(\Sigma_{\text{Old Tier 1 apps}} (D \times S \times C) / \Sigma_{\text{New Tier 1 apps}} (D \times S \times C)\right) + 10.8\ MtCO_2e * \text{sum}\left(_{\text{Tier 2 apps}}\right) * \text{sum}\left(\Sigma_{\text{Old Tier 2–4 apps}} (D \times S \times C) / \Sigma_{\text{New Tier 2–4 apps}} (D \times S \times C)\right)$$

Key:
D = Data centers
S = Servers per data center
C = Cores per server

3. Results and discussion: Quantifying greenhouse gas impact across customers

The implementation of EDB Postgres AI has resulted in significant reductions in core usage and carbon emissions for Customer 1. Modeling expected impact for two additional customers similarly showed large potential future reductions. The following sections detail core and emissions reductions for the three enterprises studied.

3.1 Customer 1

Number of applications: 300 Tier 1, 900 Tier 2–4

- Core usage reduction:
    - Without EDB PG AI: 172,800 cores
    - With EDB PG AI: 36,000 cores
    - Total reduction: 136,800 cores, a **79% decrease in core usage for these applications in scope**
- Tier-specific impact:
    - Tier 1 applications: **Emissions reduced by 66%**, from 10.8 MtCO_2e per app to 3.6 MtCO_2e
    - Tier 2–4 applications: **Emissions decreased by 83%**, from 10.8 MtCO_2e per app to 1.8 MtCO_2e
- Emissions reduction:
    - Without EDB PG AI: 12,960 MtCO_2e
    - With EDB PG AI: 2,700 MtCO_2e
    - Avoided emissions: 10,260 MtCO_2e, a **79% reduction in emissions generated by these applications**

These results highlight the significant impact EDB PG AI could have on customers' data center efficiency and environmental sustainability. The analysis showed a substantial decrease in Tier 1 and Tier 2–4 applications' emissions, due to a reduction in the number of servers required to run Customer 1's applications.

3.1.1 Customer 1: Validation

Customer 1's public sustainability reporting includes data center energy emissions. Avoided emissions were also calculated as follows:

Avoided emissions = Expected emissions – observed emissions

Where:
Observed FY23 emissions = Customer 1's reported data center emissions
Expected FY23 emissions = FY22 emissions x FY23 revenue growth

Using the available data, we calculated avoided emissions:

Expected emissions = 32,067 MtCO_2e x 1.13 = ***36,191 MtCO_2e***

Observed emissions = ***31,555 MtCO_2e***

Avoided emissions *= 36,191 MtCO_2e - 31,555 MtCO_2e =* ***4,636 MtCO_2e***

Avoided emissions percentage realized *= avoided emissions realized / avoided emissions calculated*
4,636 MtCO_2e / 10,260 MtCO_2e = ***45%***

Note: Calculated values may appear slightly different due to rounding effects.

We calculated that Customer 1 avoided 4,636 MtCO_2e from their data centers, which is approximately 45% of the total emissions we estimated would be avoided by the current subset of applications using EDB PG AI. It is not surprising that the avoided emissions are less than expected. Customer 1's data center usage may be increasing faster than revenue growth, for reasons such as an increase in the computational power required for AI or not realizing all the available data architectural gains that the use of EDB PG AI would allow. That we detected a reduction in data emissions intensity as a factor of revenue suggests that EDB PG AI reduces customers' data center energy usage in a real-world setting.

3.2 Customer 2 expected impact after implementation

Number of applications: 1,000 Tier 1; 3,000 Tier 2–4

- Core usage reduction:
  - Without EDB PG AI: 216,000 cores
  - With EDB PG AI: 96,000 cores
  - Expected total reduction: 120,000 cores, a **55.6% expected reduction in core usage for these applications**
- Tier-specific impact:
  - Tier 1 applications: **Core usage reduced by 66.7%**, from 144,000 cores to 48,000 cores
  - Tier 2–4 applications: **Core usage dropped by 33.3%**, from 72,000 cores to 48,000 cores
- Expected emissions reduction:
  - Without EDB PG AI: 18,900 $MtCO_2e$
  - With EDB PG AI: 9,000 $MtCO_2e$
  - Avoided emissions: 9,900 $MtCO_2e$, a **52% expected reduction in emissions generated by these applications**

Customer 2 could experience substantial core usage optimization with EDB PG AI, showing notable efficiency improvements across application tiers. EDB PG AI could enable Customer 2 to reduce the number of servers required for Tier 1 applications and the number of cores required for Tier 2–4 applications. Given the greater initial computational requirements of Tier 1, most of the impact (66.7% vs. 33.3% change in cores) would be seen in Customer 2's Tier 1 applications.

3.3 Customer 3 expected impact after implementation

Number of applications: 2,500 Tier 1; 7,500 Tier 2–4

- Core usage reduction:
  - Without EDB PG AI: 420,000 cores
  - With EDB PG AI: 80,000 cores
  - Expected total reduction: 340,000 cores, an **81% expected decrease in core usage for these applications**
- Tier-specific impact:
  - Tier 1 applications: A reduction of 94%, from 360,000 cores to 20,000 cores
  - Tier 2–4 applications: Core usage remained steady at 60,000 cores, reflecting existing efficiencies or operational requirements for these applications
- Expected emissions reduction:
  - Without EDB PG AI: 175,500 $MtCO_2e$
  - With EDB PG AI: 22,500 $MtCO_2e$
  - Avoided emissions: 153,000 $MtCO_2e$, an **87% expected reduction in emissions generated by these applications**

Customer 3's adoption of EDB PG AI could lead to a reduction in the cores required of up to 94% for Tier 1 applications, with no changes estimated in the cores required to run Tier 2–4 applications. Therefore, there may be specific additional opportunities to optimize Tier 2–4 applications. Given the large reduction in cores for Tier 1 applications and the total number of applications for Customer 3, validation of this estimate could make a compelling business case.

3.4. Discussion

The three case studies clearly demonstrate that customers have the potential to reduce their greenhouse gas emissions by more than 50% for applications using EDB PG AI. Customer 1's real-world validated data demonstrates that the potential for substantial avoided emissions remains even when other complexities are factored in. Customers who invest in EDB PG AI can reduce costs, energy consumption, and greenhouse gas emissions from data centers. Unlike many sustainability efforts, customers do not need to consider financial trade-offs, which indicates that adopting EDB PG AI should be a clear early step for organizations investing in operational efficiency or sustainability.

A number of important factors were not considered in this analysis. For example, the carbon intensity of electricity varies for data centers in different locations. Further, the cost of market-based mechanisms such as renewable energy credits (RECs), which are commonly leveraged by organizations to reduce market-based scope 2 emissions, were not considered. Therefore, there are substantial unexplored costs that EDB PG AI saves for customers. Collaboration with Customers 2 and 3 post-implementation to compare with the expected results would also further validate the model.

The operational and cost advantages of adopting EDB PG AI further strengthen its value proposition, supporting organizations in achieving their green IT objectives while optimizing database performance.

4. Acknowledgments

Customer data and EDB PG AI product details were provided by technical staff at EDB. The methodology for calculating avoided emissions was developed by Incendium Consulting.

Incendium Consulting is a global sustainability advisory firm specializing in net-zero strategy development, climate risk mitigation, greenhouse gas accounting, and ESG and sustainability reporting. Its team of experts brings diverse backgrounds and deep industry knowledge spanning a variety of markets including technology, financial services, real estate, pharmaceuticals, and media. Incendium has a proven track record of supporting companies as they navigate complex climate and sustainability challenges.