PRACTICAL GREEN SOFTWARE

GUIDE FOR ENTERPRISES

January 2026

INFOCOMM MEDIA DEVELOPMENT AUTHORITY

Practical Green Software
Guide for Enterprises

Your cloud bill just doubled. Your AI workloads are running 24/7. And your CFO is asking why compute costs keep climbing while your sustainability targets slip further out of reach. If this sounds familiar, you're not alone and the problem is about to get bigger. Globally, data centres are increasingly consuming more power, with AI a major driver. For a business, the question is how to grow while keeping energy and carbon in check.

If energy supply cannot keep pace with this demand, plans to scale AI and other emerging technologies for business innovation could stall, costs could escalate, and sustainability goals could be severely compromised. Currently, only about 16% of the world's largest companies are on track to achieve net-zero emissions in their own operations by 2050.

The question is no longer whether to decarbonise, but how to do it while still growing. One practical lever is **green software that is designed and operated to minimise energy and compute, cutting emissions across its lifecycle.**

Cut waste, Cut cost

The win-win of green software

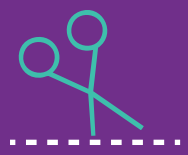

Unlocking savings with green software illustrates double materiality, where actions reduce both operational costs and environmental impact, showing the dual importance of financial and sustainability outcomes.

- **Lower costs:** Streamlining and optimising software stacks cuts energy bills and frees up computation capacity, directly improving the bottom line. In Singapore, where high energy costs amplify exposure to volatile prices and future carbon taxes, the financial upside is critical.
- **Meet sustainability goals:** At the same time, credible decarbonisation actions reduce carbon footprint and help meet regulatory, investor and customer expectations.

Your quick health check

Where does your enterprise stand?

Use this quick check to understand whether your infrastructure is broadly efficient or carrying avoidable energy and cost overhead. This is a fast signal, not a detailed assessment.

Start with Power Usage Effectiveness (PUE)

Think of PUE like a **BMI check for your data centre's energy efficiency.** The lower the value of the PUE, the more energy efficient the data centre, i.e., energy is more efficiently used and channeled towards the processing of and execution of compute workloads. Most cloud providers surface this metric directly in their sustainability dashboards.

While more advanced metrics such as **Software Carbon Intensity (SCI)** can be used for deeper software-level measurement, **PUE is the essential starting point.**

Actions to consider

- Move critical workloads to data centres or cloud regions with a **PUE of 1.3 or lower**
- Where possible, choose providers that offer **energy-efficient custom chips** for AI workloads
- For flexible or batch jobs, pick regions and times that run on **higher shares of renewable energy**

Get Started

with these three steps

As a starting point, focusing on three practical steps can help businesses deliver early results across cost and carbon reduction.

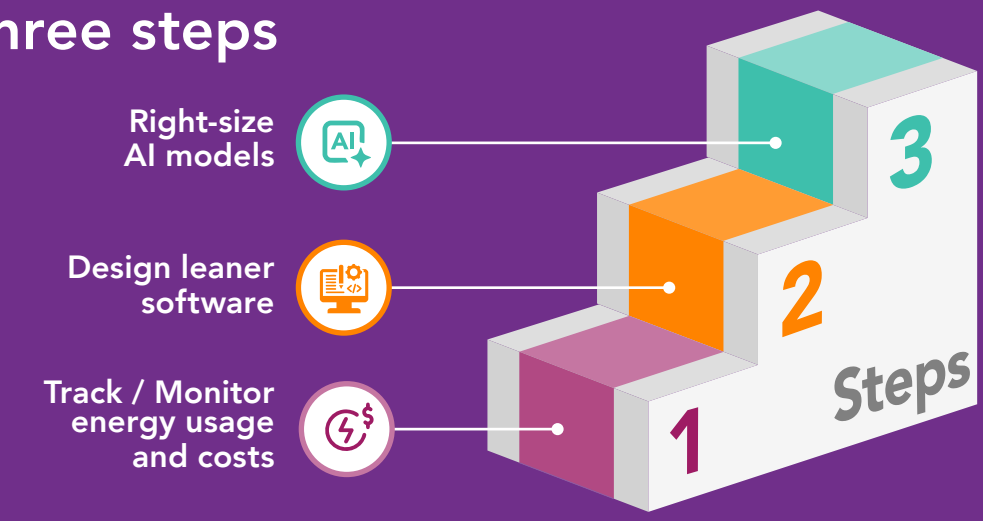

Track/Monitor energy usage and costs

- Know what you are paying? Know what is on your bill?
- From action to impact

Know what you are paying? Know what is on your bill?

Most businesses already have the data they need – they just do not know where to look. Here's how to read your cloud usage:

- **Access your provider's carbon dashboard:** Many cloud providers offer carbon footprint and energy usage dashboards that provide a starting point for tracking emissions by service and region.
- Use tools that provide daily or hourly data on energy use and emissions across your cloud environments, instead of relying only on delayed provider reports. Select tools like the open-source Cloud Carbon Footprint tool, to track emissions across multiple cloud providers in near real-time.
- For more detailed view, consider calculator that track IaaS emissions (ie virtual machines) based on actual operating environment. The Singapore's Infocomm Media Development Authority and Energy Studies Institute in NUS had developed a Cloud Carbon Calculator that estimates emissions based on vendor, instance type, utilisation and duration. You can now have an overall view of your carbon emission of a cloud service provider dashboard, as well as breakdown into workload or server instances through the Cloud Carbon Calculator.
- **Identify your biggest consumers:** Look for compute services (virtual machines), GPU usage for AI workloads, and storage. These are typically where energy use, emissions, and costs are concentrated.
- **Compare month-over-month:** Are costs and emissions trending up? This signals inefficiency that needs addressing.

From cloud usage to emission hotspots

Cloud emissions typically follow cloud spend. To understand where emissions originate, start with the same place most organisations already look – monthly cloud usage and cost.

Your provider's carbon dashboards help translate this usage data into emissions signals, making it easier to identify the services, regions, and workloads that contribute most to your cloud footprint.

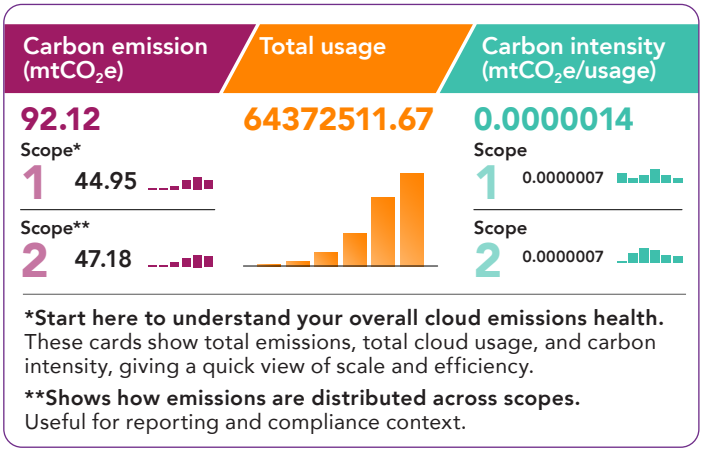

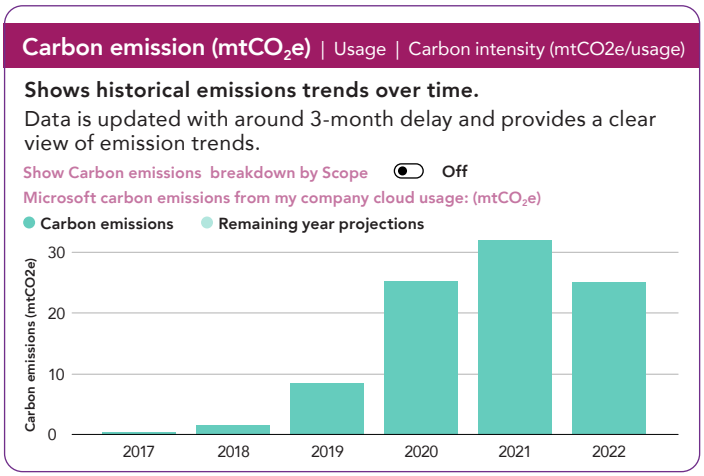

Illustrative examples of cloud provider dashboards used to understand emissions by service, region, and workload.

Illustrative examples are shown in the box stories included in this guide.

By subscription name | By region | By service

This view shows which subscriptions, regions, or services contribute most to emissions.

Subscription name	Carbon emissions ($mtCO_2e$)	Usage	Carbon intensity ($mtCO_2e$/usage)
Subscription 4	59.95*	19609118.32	0.0000031
Subscription 3	17.60	15477249.14	0.0000011
Subscription 1	8.49	15482378.89	0.0000005
Subscription 2	6.08	13803765.32	0.0000004
Total	**92.12**	**64372511.67**	**0.0000014**

*Subscriptions with higher usage contribute more to emissions.

Carbon emissions amount by location (mtCO2e)

Carbon intensity: Low High

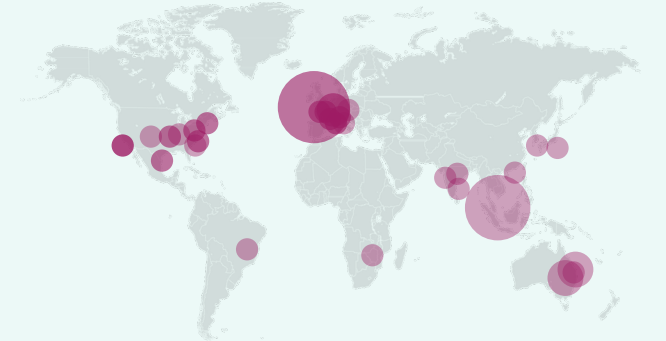

The larger the circle, the greater the emissions.

Illustrative cloud billing dashboard

Spend Summary

Cost Explorer

Welcome to the Account Billing console. Your last month, month-to-date, and month-end forecasted costs appear below.

Current month-to-date balance

$7,453.41

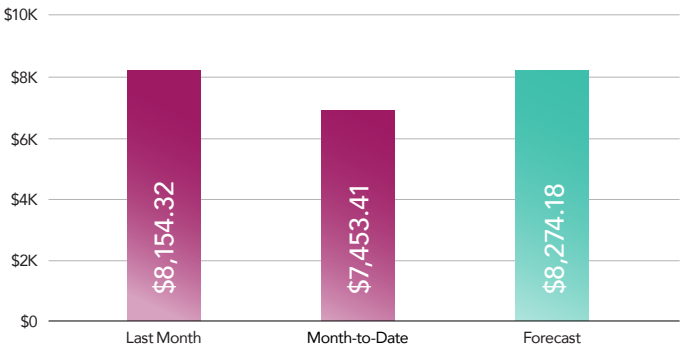

Month-to-Date Spend by Service

Bill Details

The chart below shows the proportion of costs spend for each service you use.

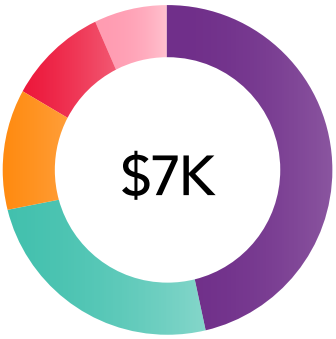

Cloud Compute Usage	$3,700.71
RDS	$1,876.35
ElastiCache	$938.18
DynamoDB	$625.44
Other Services	$312.57
Tax	$0.16
Total	**$7,453.41**

Illustrative examples of cloud provider dashboards used to understand emissions by service, region, and workload.

What to look for in dashboard

- Monthly emissions aligned to cloud usage and cost
- Services with the highest consumption (compute, GPU-intensive workloads, storage)
- Regions where emissions are concentrated
- Month-over-month trends that signal inefficiency

Illustrative examples are shown in the box stories included in this guide.

From action to impact

DEPLOY GREENOPS DASHBOARDS[1]

Near real-time visibility into energy and emissions

TRACK PROXY METRICS[2]

CPU hours, GPU hours, memory usage

(For example, sustained CPU utilisation below 30% to 40% indicates over-provisioning)

SET THRESHOLDS & ALERTS[3]

Detect abnormal CPU/GPU spikes, idle resources, and over-provisioned instances, and trigger rightsizing actions such as shutting down unused resources or shifting flexible workloads across regions or time windows

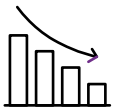

- **20 to 40% reduction** in compute waste
- **6 to 12% reduction** in overall cloud spend
- **10 to 20% reduction** in cloud cost and emissions from improved visibility and control

(Impact achieved when teams regularly act on alerts)

● Action (What teams do) ○ Impact (What happens when teams act on the signal)

Design leaner software

- Streamline the build process
- Use green code generation
- Prioritise efficiency for critical workloads
- Design for idle time

Streamline the build process[4]

Configure build and test pipelines so that only the affected modules are rebuilt and tested, instead of the whole application. Clean out unused code and components, like decluttering your digital workspace. Tools such as SonarQube, Codacy and Semgrep can highlight code smells, duplication, and other issues that bloat the codebase and can hurt performance. **Less code builds faster and more efficient systems.**

STREAMLINE THE BUILD PROCESS
(incremental builds, selective rebuilds, build and parallel execution)

OFTEN DELIVERS 20 to 50%
shorter build times in large codebases, with similar reductions in CPU and energy use where measured

● Action ○ Expected benefit

Use green code generation[5]

By using prompt engineering with AI coding assistants, developers can generate green code that **consumes less energy, emits less carbon, and may have fewer operating costs.**

USE GREEN PROMPT ENGINEERING
with AI coding assistants to auto-generate energy-efficient code

Example prompt: Optimise the following code for time, space, and energy efficiency, avoiding unnecessary computation, memory allocation, and I/O. Lightweight prompt frameworks such as PromptZero apply a similar idea by nudging AI models to produce more concise, lower-compute outputs.

15 to 30% lower runtime energy; | **5 to 12%** lower computer cost

● Action ○ Expected benefit

Prioritise efficiency for critical workloads

For your most important systems, start with quick wins: speed up database queries, remove unnecessary steps, and cut out wasteful processing. Once you have made those improvements, ask your IT team about using code efficiency tools. These tools help spot which parts of your application are working harder than they need to and wasting energy.

Design for idle time[6]

Rework applications so that they scale down or shut down when not in use instead of running 24/7. **Eliminating idle resources can reduce cloud costs by 20 to 30% on average.** On the front end, avoid unnecessarily heavy images and video that increase data transfer and device energy use.

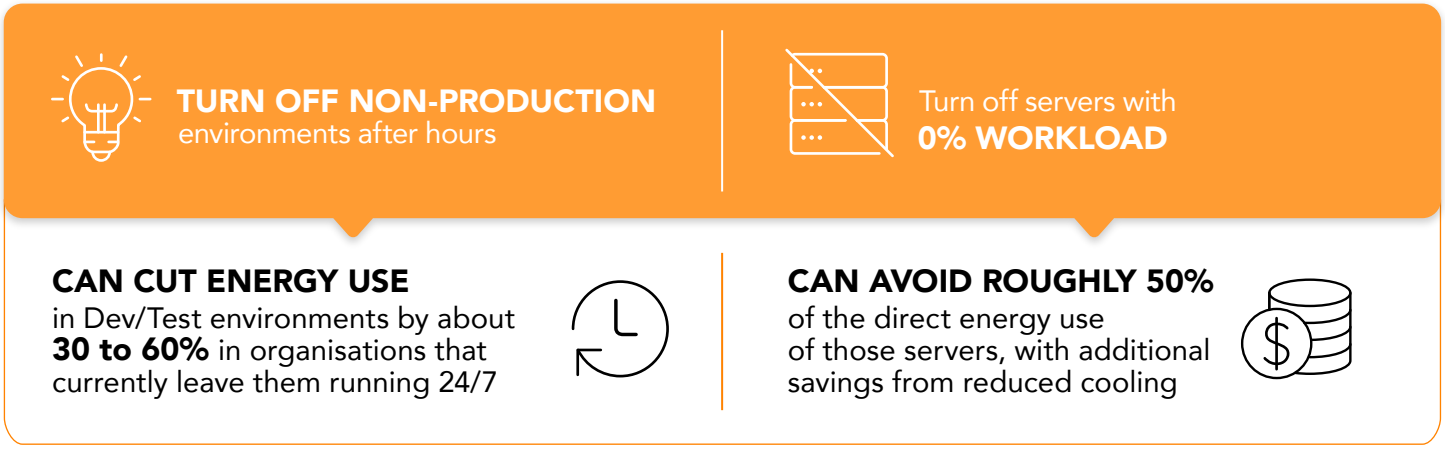

Right-size AI models

3

- Adopt a fit-for-purpose AI model strategy
- Data and prompt optimisation

Adopt a fit-for-purpose AI model strategy[7]

For most organisations, the real "carbon liability" lies not in training foundational models from scratch, but in the ongoing inference, and in the selection, fine-tuning and operationalisation of vendor models. Avoid deploying the largest model for every problem. Smaller models, often in the sub- to few-billion parameter range, are sufficient for many enterprise tasks such as classification, extraction, and summarisation. **Use pre-trained or smaller models that are right-sized for each use case.** Where possible, choose services that use energy-efficient architectures that only activate the parts of the model needed for each query.

Data and prompt optimisation[8]

When adapting models with your own data, focus on smaller, well-curated datasets that are relevant and diverse. This usually leads to better performance with fewer training runs and lower compute cost. Optimise prompts to reduce compute. **Well-designed prompts can significantly cut inference cost and latency.** Shorter, clearer prompts reduce token processing, while structured patterns such as retrieval-augmented generation (RAG) or constrained prompting minimise unnecessary computation. Even modest prompt improvements can **yield measurable energy and cost savings at scale.**

USE PRE-TRAINED OR SMALLER MODELS
that are right-sized for each use case

CAN DELIVER ROUGHLY 30 TO 70% LOWER ENERGY USE
per inference for many workloads, depending on the baseline model and hardware

● Action ○ Expected benefit

Data and prompt **OPTIMISATION**

PROMPT OPTIMISATION
has delivered double-digit energy and cost reductions, but the effect size is highly context-dependent

● Action ○ Expected benefit

The Singapore model: Growth with efficiency

Singapore is expanding data centre and AI capacity while driving energy efficiency improvements.

- Launched **Tropical DC standard (SS697:2023), IT Energy Efficiency standard (SS715:2025)** and **Green Mark for DCs 2024** to raise the bar for energy efficiency.
- At least **200 MW** of DC capacity will be made available by **EDB and IMDA**, with potentially more through the adoption of new and innovative green energy pathways.
- **Standards leadership:** Singapore is the **first government organisation to join the Green Software Foundation**, contributing to **SCI** and driving adoption.
- **Call-to-action for operators:** Improve **PUE**, align to updated standards, and leverage grant support to drive best-in-class energy efficiency.

Simple execution plan to drive results

A practical way to move from policy intent to measurable efficiency gains.

1. **Assess**
   - Establish a baseline for energy, cost, and utilisation
   - Identify high-impact workloads and inefficiencies
2. **Optimise**
   - Move workloads to lower-PUE regions and efficient hardware
   - Apply green code practices and remove wasteful execution paths
   - Right-size AI models and optimise prompts
3. **Validate**
   - Measure changes in energy use, emissions, and cost
   - Confirm improvements on priority workloads
4. **Scale**
   - Embed efficiency checks into delivery and operations
   - Roll proven actions across teams and platforms

Sources:

[1] **Deploy GreenOps Dashboards**
- FinOps Meets GreenOps: Cloud Cost & Carbon Optimisation on Azure, https://www.linkedin.com/pulse/finops-meets-greenops-cloud-cost-carbon-optimisation-azure-anil-kumar-n6qhf/
- MAIZX: A Carbon-Aware Framework for Optimising Cloud Computing Emissions, presented at LOCO '24 (Dec 2024). https://arxiv.org/abs/2506.19972

[2] **Track proxy metrics**
- A Systematic Review of Energy Efficiency Metrics for Optimising Cloud Data Center Operations and Management. Electronics 2025, 14, 2214. https://doi.org/10.3390/electronics14112214

[3] **Set thresholds and alerts**
- A systematic review on effective energy utilisation management strategies in cloud data centres. 2022. https://link.springer.com/content/pdf/10.1186/s13677-022-00368-5.pdf
- Investigating cloud instances to achieve optimal trade-offs between performance-cost efficiency. Feb 2025. https://link.springer.com/article/10.1007/s00607-025-01444-9

[4] **Streamline the build process**
- Build optimisation: A systematic literature review." ACM Computing Surveys 58.1 (2025): 1-38.
- What helped, and what did not? An evaluation of the strategies to improve continuous integration. 2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE). IEEE, 2021. https://arxiv.org/abs/2102.06666

[5] **Use green code generation**
- An Empirical Study: Leveraging Prompt Engineering with AI Coding Assistants to Develop Energy-Efficient Code (2025), https://www.techrxiv.org/doi/full/10.36227/techrxiv.175339126.69681777

[6] **Design for idle time**
- Maximising Energy Efficiency: Optimising Light-Off Scenarios During Idle Hours in Data Centres, https://www.caeled.com/blog/data-center-lighting/maximizing-energy-efficiency-optimizing-light-off-scenarios-during-idle-hours-in-data-centers/
- Joint Optimisation of Idle and Cooling Power in Data Centres While Maintaining Response Time. https://engineering.purdue.edu/~puma/powertrade.pdf
- 2024 United States Data Centre Energy Usage Report." Lawrence Berkeley National Laboratory (US Dept. of Energy), Dec. 2024. lbnl-2024-united-states-data-center-energy-usage-report.pdf

[7] **Adopt a fit-for-purpose AI model strategy**
- Small is Sufficient: Reducing the World AI Energy Consumption Through Model Selection. Oct 2025. https://arxiv.org/abs/2510.01889v1

[8] **Data and prompt optimisation**
- Energy-Aware Prompt Optimisation for Large Language Models: Balancing Accuracy, Cost, and Sustainability. TechRxiv. Aug 2025. https://doi.org/10.36227/techrxiv.175615629.90190202/v1

Acknowledgements:

The Infocomm Media Development Authority of Singapore would like to thank all our industry partners, agencies and individuals who have contributed to the green software movement in Singapore. We would also like to thank Accenture, Green Software Foundation and The National Climate Change Secretariat who have supported the development of this Practical Green Software Guide for Enterprises.

About Infocomm Media Development Authority

IMDA leads Singapore's digital transformation by developing a vibrant digital economy and an inclusive digital society. As Architects of Singapore's Digital Future, we foster growth in Infocomm Technology and Media sectors in concert with progressive regulations, harnessing frontier technologies, and developing local talent and digital infrastructure ecosystems to establish Singapore as a digital metropolis.

For more news and information, visit www.imda.gov.sg or follow IMDA on LinkedIn (IMDAsg), Facebook (IMDAsg) and Instagram (@imdasg).